

CLOUD INFRASTRUCTURE SERVICES

Final decision report

31 July 2025

© Crown copyright 2025

You may reuse this information (not including logos) free of charge in any format or medium, under the terms of the Open Government Licence.

To view this licence, visit www.nationalarchives.gov.uk/doc/open-government-licence/ or write to the Information Policy Team, The National Archives, Kew, London TW9 4DU, or email: psi@nationalarchives.gov.uk.

Website: www.gov.uk/cma

**Members of the Competition and Markets Authority
who conducted this inquiry**

Kip Meek (*Chair of the Group*)

Robin Foster

Colleen Keck

Paul Hughes

Chief Executive of the Competition and Markets Authority

Sarah Cardell

The Competition and Markets Authority has excluded from this published version of the final report information which the inquiry group considers should be excluded having regard to the three considerations set out in section 244 of the Enterprise Act 2002 (specified information: considerations relevant to disclosure). The omissions are indicated by [X]. Some numbers have been replaced by a range. These are shown in square brackets. Non-sensitive wording is also indicated in square brackets.

Contents

Summary	12
Introduction.....	12
The nature of competition in cloud services markets.....	12
The impact of AI on competition in cloud services.....	13
Entry and expansion in cloud services	14
Customers' ability to switch cloud provider and multi-cloud.....	15
Microsoft's software licensing practices	15
Our decision on competition	17
Our decision on remedies.....	18
Findings.....	19
1. Our task.....	19
Introduction.....	19
Nature and purpose of market investigations	19
This market investigation reference.....	20
The role of public cloud infrastructure services in the UK today	21
Scope and focus of the investigation	21
Our approach to evidence gathering	23
Information requests and customer calls.....	24
Consultations	24
Engagement with main parties	24
Requests for disclosure of certain documents and information	25
Jigsaw research	26
Approach to our assessment and the structure of this document.....	26
Next steps	27
2. Industry background.....	28
The demand for cloud services	29
Customers of cloud services	29
Growth in demand for cloud services	30
Providers' market positions	31
AWS 31	
Microsoft	32
Google 33	
IBM 34	
Oracle 35	
Other providers	35
Our assessment of the main providers.....	36
How customers purchase cloud services	37
Parameters of competition.....	38
Evidence from providers	38
Evidence from customers.....	40
Switching and multi-cloud.....	40

Switching process	41
Types of multi-cloud	42
Benefits and disadvantages of using multi-cloud for customers.....	45
Conclusions.....	46
3. Competitive landscape	48
Market definition	49
Product market.....	50
Geographic market.....	88
Market structure and concentration	92
Framework for our assessment.....	92
Shares of supply by revenue.....	94
Shares of supply by capacity.....	104
Shares of supply by flows of new business	108
Conclusions on market shares and concentration.....	117
Market outcomes	119
Introduction	119
How we use evidence on market outcomes	119
Profitability.....	121
Pricing trends	140
Quality and innovation.....	145
Prevalence of switching and multi-cloud	149
Conclusions.....	171
4. The impact of AI on cloud services.....	173
Introduction and background	173
Cloud provider revenue related to AI	176
Cloud services' AI-related revenue	176
Breakdown of cloud services AI-related revenue by service and customer	178
Our assessment of cloud provider revenue related to AI	179
The impact of competition in the provision of IaaS based on accelerated compute on cloud services markets	180
Providers of IaaS based on accelerated compute and their relative strengths....	180
Barriers to entry and expansion	184
Cloud providers' supply of access to FMs to their customers.....	193
How cloud providers compete to supply customers with access to FMs.....	194
Impact of the supply of access to FMs on competition in the supply of cloud services	195
Impact of cloud customers' use of FMs on switching and multi-cloud	198
Cloud providers' relative strength in the provision of access to FMs.....	200
Our assessment of cloud provider' relative strengths in the provision of access to FMs	202
The impact of AI capabilities on competition between cloud providers.....	203
The impact of AI on how customers choose their cloud provider	204

	Cloud providers' relative strengths in AI and the impact on their provision of other cloud services.....	208
	Our assessment of cloud providers' relative strengths in AI and the impact on their provision of other cloud services	211
	Conclusions.....	213
5.	Barriers to entry and expansion.....	214
	Sunk investment costs	215
	Stakeholder views on investment.....	216
	Our assessment.....	217
	Our conclusions on entry and expansion in the supply of IaaS.....	224
	Economies of scale	225
	Stakeholder views on economies of scale	226
	Purchasing efficiencies	227
	Operating efficiencies.....	228
	Investment in research and development	231
	Our conclusions on economies of scale in the provision of cloud services	233
	Size of cloud providers' product portfolio.....	233
	Importance of range of services	234
	Economies of scope.....	235
	Network effects	237
	Other potential barriers to entry and expansion.....	239
	Cloud credits	239
	Reputational barriers.....	245
	Regulatory barriers.....	246
	Public sector procurement	248
	Conclusions on barriers to entry and expansion.....	249
6.	Barriers to switching and multi-cloud.....	252
	Introduction	253
	Providers' incentives to facilitate switching and multi-cloud.....	254
	Cloud provider views.....	254
	Our assessment.....	257
	Technical barriers.....	260
	Introduction	260
	Conceptual framework	263
	Technical barriers to switching and multi-cloud.....	266
	Core services	280
	Ancillary services and tools.....	286
	Other technical factors	303
	Technical mitigations.....	310
	Conclusions on technical barriers	326
	Egress fees	327
	Introduction	327
	Background.....	327

Conceptual framework	329
Relevance of egress fees to customers' switching and multi-cloud decisions.....	330
Our assessment of the relevance of egress fees to customers' switching and multi-cloud decisions	365
Potential benefits of egress fees	368
Our assessment of the potential benefits of egress fees.....	387
Conclusions on egress fees	387
Conclusions on barriers to switching and multi-cloud	388
7. Licensing	390
Background	391
Description of the software licensing practices.....	393
Timeline of licensing practices	393
Using Microsoft software products on public cloud	395
Differences between using Microsoft software products on Azure compared to on non-Azure clouds via SPLA.....	398
Conceptual framework	400
Inclusion of other cloud providers' conduct in this assessment	402
Additional cloud provider submissions	403
Microsoft's market power in related software markets.....	405
Introduction	405
Framework	405
Product background	408
Market definition and market power assessments	414
Conclusions on Microsoft's market power in related software markets	453
The importance of Microsoft software inputs to AWS and Google	454
Distribution of cloud use by spend	455
Significance in the cost base of AWS and Google	456
Significance in shaping downstream competition.....	482
Our assessment.....	500
Conclusions on the importance of Microsoft software inputs to AWS and Google	502
Microsoft's supply of software: price and non-price differences	502
Our comparison of the price that Microsoft charges Azure customers for Windows Server and SQL Server and the price it charges AWS and Google	502
A cloud provider's analysis of Windows Server licensing cost versus infrastructure price.....	511
List prices	513
Price and non-price differences for the client-side products.....	513
Our assessment.....	520
Conclusions on price and non-price differences.....	520
The impact on AWS and Google's competitiveness from Microsoft's licensing practices	521
AWS' and Google's submissions on pass-through of input costs.....	521

Profit margins and mark-ups	523
Relative usage	535
Customers' choice of cloud providers	543
Our assessment.....	558
Conclusions on the impact on AWS and Google's competitiveness from Microsoft's licensing practices	559
Summary of our assessment and conclusions	559
Ability 559	
Incentive.....	561
Effect 562	
Conclusion	563
8. Committed spend agreements	564
Introduction.....	564
Conceptual framework and analysis structure.....	565
Stakeholder submissions.....	571
Cloud providers and ISVs	571
Industry bodies.....	573
Customers.....	573
Other parties	574
Our assessment	574
Prevalence of CSDs.....	574
Impact of CSDs on customers' choices.....	577
Impact of CSDs on competition.....	580
Potential benefits of CSDs	593
Conclusions.....	594
9. Our decision on competition in cloud services markets.....	596
Our view of well-functioning cloud services markets	596
Competitive assessment	598
Market concentration	598
Market outcomes	598
Impact of AI on cloud services.....	600
Barriers to entry and expansion.....	601
Barriers to switching and multi-cloud	602
Microsoft's licensing practices	604
Our decision on adverse effects on competition.....	606
Customer detriment.....	607
10. Our remedy decision	608
Our decision on remedies.....	609
Structure and context	610
Framework for our remedies assessment	610
Our duties	610
The CMA's digital markets regulatory powers	614
Stakeholder views on our remedies	615

Our assessment of effectiveness of these remedies	617
How the recommendations are effective in addressing the harm.....	617
Likelihood of the recommendations being acted upon	622
Proportionality assessment of these remedies	625
Effective in achieving its legitimate aim	625
No more onerous than needed.....	625
Least onerous if there is a choice between several effective measures.....	626
Does not produce disadvantages which are disproportionate to the aim	626
Other remedy options that we have decided not to pursue	627
Measures we have decided not to pursue under the remedy-making powers provided by EA02	628
Structural remedies	636
Utility-style regulation	636

Tables

Table 2.1 : UK IaaS and PaaS revenues, 2020-2024 (£bn).....	30
Table 3.1 : UK+EEA shares of supply for IaaS, 2020 – 2024	95
Table 3.2 : UK+EEA shares of supply for PaaS, 2020 – 2024	97
Table 3.3 : UK+EEA shares of supply for IaaS and PaaS combined, 2020 – 2024	98
Table 3.4 : UK + EEA shares of supply by data centre capacity, 2020 – 2026	106
Table 3.5 : Global shares of supply by data centre capacity, 2020 – 2026	107
Table 3.6 : UK+EEA shares of year-on-year IaaS revenue growth, 2021 – 2024	111
Table 3.7 : UK+EEA shares of year-on-year IaaS and PaaS revenue growth, 2021 – 2024	112
Table 3.8 : UK shares of supply by revenue from newly acquired customers, 2021-2024	114
Table 3.9 : prevalence of multi-cloud, unweighted and weighted by spend, 2020-2023 ..	168
Table 3.10 : Average proportion of spend on primary cloud for customers that multi-cloud, by spend band, 2020-2023 (%).....	169
Table 7.1 Microsoft's market share in server OS, variety of measures	417
Table 7.2 : Market shares for RDBMS, all deployment types, 2019-2023	423
Table 7.3 : Share of Azure UK customers with CSAs active in 2024 by revenue bracket	466
Table 7.4 : Windows Server input costs as a proportion of customer spend on all Azure services and Windows Server VMs on Azure, Google SPLA prices	478
Table 7.5 : SQL Server input costs as a proportion of spend on all Azure services and spend on Windows Server VMs and SQL Server IP on Azure, Google SPLA prices	478
Table 7.6 : Median SQL Server input costs as a proportion of customer spend on all Azure services and spend on Windows Server VMs and SQL Server IP on Azure, Google SPLA prices	479
Table 7.7 : Combined Windows Server and SQL Server input costs as a proportion of customer spend on all Azure services and spend on Windows Server VMs and SQL Server IP on Azure, Google SPLA prices	479

Table 7.8 : Percentage of revenue from sales of Server OS deployed on the public cloud	496
Table 7.9 : Azure customers' relative spend on Windows Server VMs	497
Table 7.10 : Average percentage difference between the input prices that Google pays for Windows Server and Microsoft's customer-facing prices	506
Table 7.11 : Median percentage difference between the input prices that Google pays for Windows Server and Microsoft's customer-facing prices	506
Table 7.12 : Average difference between the input prices that Google pays for Windows Server and Microsoft's customer-facing prices, as a proportion of total spend on Azure services and Windows Server VM spend on Azure	507
Table 7.13 : Median difference between the input prices that Google pays for Windows Server and Microsoft's customer-facing prices, as a proportion of total spend on Azure services and Windows Server VM spend on Azure	507
Table 7.14 : Average percentage difference between the input prices that Google pays for SQL Server and Microsoft's customer-facing prices relating to SQL Server PAYG usage	507
Table 7.15 : Median percentage difference between the input prices that Google pays for SQL Server and Microsoft's customer-facing prices relating to SQL Server PAYG usage	508
Table 7.16 : Average difference between the input prices that Google pays for SQL Server and Microsoft's customer-facing prices relating to SQL Server PAYG usage, as a proportion of total Azure spend and spend on Windows Server VMs and SQL Server IP on Azure	508
Table 7.17 : Median difference between the input prices that Google pays for SQL Server and Microsoft's customer-facing prices as a proportion of total Azure spend and spend on Windows Server VMs and SQL Server IP on Azure	509
Table 7.18 : Average combined difference between the input prices that Google pays for Windows Server and SQL Server and Microsoft's customer-facing prices relating to Windows Server AHB and SQL Server PAYG usage, as a proportion of total Azure spend and spend on Windows Server VMs on Azure	509
Table 7.19 : Median combined difference between the input prices that Google pays for Windows Server and SQL Server and Microsoft's customer-facing prices relating to Windows Server AHB and SQL Server PAYG usage, as a proportion of total Azure spend and spend on	510
Table 7.20 : The cloud provider's analysis of differences in licensing costs paid by that cloud provider compared to customer facing price for Azure customers	511
Table 7.21 Availability of the Microsoft products for use in the public cloud	514
Table 7.22 : Estimates of the margin Microsoft would earn if subject the SPLA licensing costs by customer group in the UK (2024)	531

Figures

Figure 3.1 : Prevalence of multi-cloud, weighted by spend, split by spend band, 2020-2023	169

Figure 6.1 : The costs associated with different customer activities	262
Figure 6.2 : Types of data transfers	328
Figure 7.1 : [REDACTED]	455
Figure 7.2 : Shares of Azure VM usage by operating system 2020-2023.....	489
Figure 7.3 : [REDACTED].....	537
Figure 7.4 : [REDACTED].....	537
Figure 7.5 : [REDACTED].....	537
Figure 8.1 : Discount that the rival must apply to win units of contestable demand	567
Figure 8.2 : Total profit to rival for different units of contestable demand it wins.....	569
Figure 8.3 : Total profit to rival for different units of contestable demand it wins for a representative ([REDACTED]) AWS' CSD	590

Appendix

- A. Demand for cloud services and how customers purchase cloud services
- B. Parameters of competition
- C. Quantitative surveys
- D. Market structure and concentration methodology and UK share of supply by revenue
- E. Profitability analysis
- F. Pricing
- G. Quality and innovation
- H. Prevalence of switching methodology
- I. Prevalence of multi-cloud
- J. Barriers to entry and expansion [Redacted]
- K. Public sector procurement
- L. Egress fees – background
- M. Egress fees – analysis of customers' cost scenarios
- N. Egress fees – free switching programmes
- O. Egress fees – customer views on egress fees
- P. Egress fees internal documents
- Q. Egress fees – analysis of cloud providers' cost of egress
- R. Market power – supplementary framework and evidence
- S. Methodology for a cloud provider's analysis of licensing costs
- T. Licensing analysis
- U. The impact of CSAs on competition, a quantitative analysis
- V. The impact of CSAs on competition, a qualitative analysis
- W. Remedies not taken forward in this market investigation

Summary

Introduction

1. Our task in this investigation has been to decide whether competition is working well in cloud services markets and, if it is not, to decide what action should be taken to remedy any harms to competition. Our final decision is that competition is not working well, so we have made recommendations to address this.
2. Public cloud infrastructure services (called cloud services in this report) give customers access to shared computing resources on demand so that the customers do not need to own the underlying hardware and software. These services include processing, storage, networking and other raw computing resources (often referred to as infrastructure as a service or IaaS) as well as services that can be used to develop, test, run and manage applications in the cloud (often referred to as platform as a service or PaaS).
3. Cloud services have transformed the way UK businesses and public services run their operations and are a critical input to the digital services that UK consumers use each day. Customers spent £10.5 billion on cloud services in 2024, with spending growing by nearly 30% each year since 2020 as they rely increasingly on cloud, rather than on-premises or 'traditional' IT. Cloud services also underpin AI model development and deployment.
4. Cloud services support many sectors' contribution to the UK's economic growth. It is therefore vital that competition works well in these markets to support innovation, investment and improved productivity for the benefit of people, businesses and the UK economy.

The nature of competition in cloud services markets

5. The IaaS market is highly concentrated and the two largest providers, Microsoft and AWS, each has a high share of supply at [30-40]% in 2024. Microsoft has grown its share while AWS' share has decreased since 2020. The PaaS market is also concentrated; Microsoft's share of it has slightly increased to [20-30]% and AWS' share has remained stable at [10-20]% since 2020.¹
6. The third largest provider, Google, has a much lower share of supply in both IaaS and PaaS markets at [5-10]% although this has grown since 2020. Other providers, including Oracle and IBM, have shares that are even smaller and do not supply as wide a range of cloud services as Microsoft, AWS and Google. There are a larger number of PaaS providers, including Independent Software Vendors

¹ The geographic scope of our markets is UK+EEA.

(ISVs), but their services are reliant on IaaS and we consider that they are a limited competitive constraint on IaaS providers.

7. Other data on cloud providers' capacity, shares of revenue flow and new customers indicate that these levels of market concentration and the leading positions of Microsoft and AWS are likely to endure.
8. Microsoft and AWS have been generating sustained returns from their cloud services substantially above their cost of capital in cloud services for a number of years.
9. Prices paid by cloud customers for different cloud services have moved in different directions for different services in recent years, with some services increasing in price and others falling. Customers say that cloud providers offer both quality and innovation to them, as well as advantages over on-premises IT.
10. We consider that a more competitive market would show better market outcomes, including more consistently competitive prices, as well as further improvements in quality and innovation.

The impact of AI on competition in cloud services

11. Some cloud providers have told us that AI is changing how competition works in cloud services. We have considered the impact on cloud services of both the provision of IaaS based on accelerated compute and customers' purchase of AI-based cloud services.
12. Microsoft, AWS and Google are all strong vertically integrated providers of accelerated compute and AI-related cloud services, including through partnerships with AI model developers. Other cloud providers are significantly weaker in this respect.
13. Whilst AI-related cloud services are likely to be an increasingly important part of the offer from cloud service providers, current evidence from customers and providers suggests that AI is yet to materially alter the competitive dynamics in cloud services.
14. Providers' supply of IaaS based on accelerated compute is not currently significantly impacting competition in cloud infrastructure services because the services are not substitutable and there are significant barriers to entry and expansion in IaaS. Any impact is further limited by the fact that IaaS based on accelerated compute is being used by a small number of customers who are primarily AI model developers.
15. Revenue from the provision to customers of access to AI models is rising for Microsoft, AWS and Google but it is still a relatively small part of their businesses.

The scope for demand for access to AI models to translate into greater demand for other cloud services is currently limited, although we expect it to grow over time.

16. AI capabilities currently play a limited role in driving customers' choice of cloud provider. However, these capabilities are likely to become more important to customers over time and so AI could impact competition in cloud services to a greater degree in the future. The nature and timing of this change and how it changes the relative strength of Microsoft, AWS and Google is dependent on a number of factors, including how customers' demand for AI capabilities evolves and the extent to which this leads to increased demand for other cloud services from the same provider.

Entry and expansion in cloud services

17. There are substantial barriers to entry and expansion in cloud services, particularly in IaaS as this requires significant capital investment in fixed assets such as data centres, networks and servers and components which become largely a sunk cost.
18. There are also economies of scale, meaning that larger cloud providers have lower average costs compared to smaller providers. The largest cloud providers are making substantial investments in IaaS and AI capabilities to expand their services in coming years, and while this investment can have pro-competitive effects and benefit cloud customers, it may also deter market entry or expansion by potential rivals.
19. The broad product portfolios of Microsoft, AWS and Google in both IaaS and PaaS are also likely to contribute to barriers to entry and expansion as range of services is an important consideration for customers when selecting a cloud provider.
20. We have considered whether procurement of cloud services by public sector customers affects competition in cloud services markets. Microsoft and AWS appear to be the largest providers to the public sector and this is consistent with their overall position in cloud services markets.
21. Public sector procurement policy aims to maintain competition from public sector customers, including by requiring competitive tendering of contracts, and greater competition in cloud services would create greater choice for public sector customers. We suggest that the UK government should continue to monitor the outcomes of public procurement of cloud services and drive best practice in the application of procurement frameworks.
22. Some stakeholders raised concerns about the impact of cloud credits on customers' incentives and the ability of smaller cloud providers to compete. We recognise that cloud credits could have a negative impact on smaller cloud providers that are unable to match the credits offered by Microsoft, AWS and

Google but we have assessed how cloud credits work and we do not consider that they are a feature which is harming competition in cloud services at present, although this harm could materialise in future if their value grew or scope changed.

Customers' ability to switch cloud provider and multi-cloud

23. We have looked at the extent to which customers can switch cloud provider and/or use multiple clouds, as their ability to exercise choice can incentivise providers to improve their offerings and reduce barriers to entry and expansion.
24. Very few customers switch between clouds: less than 1% of customers switch provider each year. Use of multiple cloud providers is more prevalent than switching but it is still uncommon for small and medium sized customers. Large cloud customers are more likely than smaller ones to use multiple cloud providers, although their spending generally remains concentrated with one main provider.
25. Customers face both commercial and technical barriers when seeking to multi-cloud or switch their cloud provider and many currently think that the costs outweigh the benefits. Barriers to multi-cloud negatively affect many customers' ability to use and integrate multiple public clouds and this limits customers' ability and incentive to exercise choice over their cloud provider.
26. A key commercial barrier is the presence and magnitude of egress fees required to transfer data between cloud providers for the purposes of switching and/or multi-cloud. These reduce the ability and incentives for customers to switch or multi-cloud, particularly smaller customers and those with large proportions of stored data for whom fees are highest relative to their total cloud spending. They also reduce the incentives of suppliers to compete for their rivals' customers.
27. Technical barriers stem from the differentiation of features and interfaces in cloud services, meaning that customers are inhibited from switching or multi-cloud as they cannot easily compare, substitute, or integrate services from different providers. Further disincentives include latency between clouds, a lack of transferable skills amongst customers' staff and insufficient transparency from providers on how to mitigate or overcome technical barriers.
28. Use of committed spend agreements by providers is widespread, and these can influence customers' choices in relation to workload allocation. But providers can profitably compete against these and, accordingly, in their current form and application, they do not harm competition in cloud services markets.

Microsoft's software licensing practices

29. Microsoft provides a range of software that customers use on cloud services, either on Microsoft's own cloud Azure or through licensing to other cloud

providers. We have examined whether Microsoft's licensing practices are influencing customers' choice of cloud provider, and whether this harms competition by preventing AWS and Google from competing as effectively for customers that use Microsoft software on cloud.²

30. Microsoft has significant market power in relation to several software products because customers are unable or unwilling to switch away from them and Microsoft has moderate to high market shares in respect of each of them. These software products are important inputs to cloud services, such that Microsoft has the potential to harm AWS and Google when customers purchase cloud services that incorporate them.
31. We found significant differences relating to price and/or quality when customers use these software products on Microsoft's cloud rather than AWS' or Google's. For example, the input price paid to Microsoft by AWS and Google for some of these products can be higher than Microsoft's customer-facing price for some cloud customers. AWS and Google pass through at least some of the input costs of Microsoft software, and customers generally perceive them to be more expensive than Microsoft.
32. While profit margin levels are not determinative of foreclosure, estimates of AWS' and Google's profit margins on bundles of cloud services that incorporate the relevant Microsoft software products are consistent with their partial foreclosure.
33. Microsoft does not make certain products available to AWS and Google through their licensing agreements, and customers with existing licences cannot bring these to AWS and Google in most instances. These restrictions mean that AWS' and Google's competitive offerings are directly affected by Microsoft's licensing practices.
34. Certain software products are used disproportionately on Azure compared to on AWS and Google. We consider that, given the very large difference, this is at least partly because some customers' choice of cloud is influenced by Microsoft's licensing practices.
35. Microsoft's licensing practices are adversely impacting the competitiveness of AWS and Google in the supply of cloud services, particularly in competing for customers that purchase cloud services which use the relevant Microsoft software as an input. As a result, Microsoft faces weaker competitive constraints from AWS and Google, its most significant competitors, which is reducing competition in cloud services markets.

² The relevant Microsoft software products are: Microsoft's Windows Server (which includes Active Directory functionality), SQL Server, Windows 10/11, Visual Studio and the productivity suites.

Our decision on competition

36. We have found that certain features of the UK cloud services markets lead to adverse effects on competition (AECs).
- (a) Market concentration, which is particularly high in the case of IaaS, and barriers to entry and expansion are features that have enabled each of the two largest providers, Microsoft and AWS, to hold significant unilateral market power in the cloud services markets and to earn returns above the cost of their capital over a sustained period. Forward-looking metrics indicate that these levels of concentration are likely to endure. These features give rise to an AEC in cloud services in the UK because it is harder for alternative cloud suppliers to enter and grow in these markets and customers face a limited choice of suppliers and products. This harm is exacerbated by the features arising from technical and commercial barriers to switching and multi-cloud.
 - (b) These technical and commercial barriers are features that give rise to an AEC in cloud services in the UK. These barriers lock customers into their initial choice of provider, which may not reflect their evolving needs and limits their ability to exercise choice of cloud provider. These barriers can restrict customers from responding to attractive offers or accessing innovative new services from another provider, leading to weaker competition between providers.
 - (c) Microsoft's licensing practices have the effect of reducing competition in cloud services markets by adversely impacting the competitiveness of AWS and Google in the supply of cloud services, particularly in competing for customers that purchase cloud services which use the relevant Microsoft software as an input. This gives rise to an AEC in cloud services in the UK. These licensing practices are a feature that, in combination with the other features we have identified, including Microsoft's large and increasing market share in these markets, further restricts the already limited choice and attractiveness of alternative products and suppliers.
37. We think outcomes for customers would be better if cloud markets were more competitive. These outcomes would include more consistently competitive prices,³ greater prevalence of switching and multi-cloud use and potentially higher quality and innovation.

³ By way of illustration, if prices were on average 5% above those in well-functioning markets, this would in aggregate result in UK customers paying around £500 million more per year for these services than they would in more competitive markets. The value of IaaS and PaaS UK revenue in 2024 was £10.5bn. Calculations are as follows: $10.5\text{bn} - (10.5\text{bn}/1.05) = 10.5\text{bn} - 10\text{bn} = 500\text{m}$.

Our decision on remedies

38. Cloud services provide critical infrastructure for the UK's economy and it is therefore vital that competition works well in these markets. In order to remedy the harms to competition that we have found, we recommend that the CMA Board use its digital markets powers to prioritise commencing SMS investigations to consider designating the two largest providers Microsoft and AWS with strategic market status (SMS) in relation to their respective digital activities in cloud services. This would enable the CMA to impose targeted and bespoke interventions to address the concerns we have identified, including with respect to features where there are specification risks around the design of effective market interventions.
39. The CMA's powers under the DMCC Act have been specifically designed to be effective in digital markets and, if the CMA was to designate Microsoft and AWS with SMS, it would be able to address the concerns we have identified with conduct requirements or other interventions. Measures aimed at Microsoft and AWS would address market-wide concerns by directly benefitting most UK customers and producing wider indirect effects by altering the competitive conditions for other providers.
40. The new digital markets powers have scope for iteration and effective mechanisms for monitoring and oversight. Any interventions could be adapted in response to market developments and would reduce the risk of regulatory divergence with interventions in other jurisdictions, including those in Europe under the Data Act.
41. The CMA has said that it will keep under review possible options for further SMS designation investigations, and it anticipates that options will be considered by the CMA Board in early 2026. We expect that our findings and recommendations will be taken into account as part of its decision.
42. While we recognise that there is some uncertainty around the implementation of our recommendations, we believe they are the only effective remedy capable of comprehensively addressing the AECs we have found in this investigation. We suggest that, pending the CMA's decision on future designation investigations, developments in these markets, including any actions taken by cloud providers, are kept under review and that the CMA Board factors these into its prioritisation decision.

Findings

1. Our task

Introduction

- 1.1 On 5 October 2023, the Office of Communications (Ofcom), in exercise of its powers,⁴ made a reference to the Competition and Markets Authority (CMA) for a market investigation into the supply of public cloud infrastructure services in the United Kingdom (UK).
- 1.2 On 5 October 2023, the CMA appointed from its panel a group of four independent members to lead the investigation (the inquiry group).⁵
- 1.3 On 28 January 2025, we published our provisional decision report (PDR) and proposed remedies for consultation with interested parties.⁶
- 1.4 This report sets out the findings of the inquiry and our final decisions.

Nature and purpose of market investigations

- 1.5 A market investigation allows the CMA the opportunity to assess whether competition in a market is working effectively with the focus on the functioning of the market as a whole rather than on a single aspect of it or the conduct of a particular firm within it.⁷
- 1.6 Where a market investigation reference is made, Part 4 of the Enterprise Act 2002 (EA02) empowers the CMA to examine whether any relevant feature of each relevant market prevents, restricts or distorts competition in connection with the supply or acquisition of any goods or services in the United Kingdom or a part of the United Kingdom.⁸ These powers enable the CMA to determine whether there are market-wide features, or a combination of features, that have an adverse effect on competition (AEC).⁹

⁴ Sections 131 and 133 of the Enterprise Act 2002 (EA02), as provided for by the Communications Act 2003.

⁵ Details on the members of the Group are on our case page: [Cloud services market investigation](#).

⁶ CMA, Cloud market investigation provisional decision report, January 2025, [Cloud services market investigation - GOV.UK](#).

⁷ See [CC3 \(Revised\)](#), [Guidelines for market investigations: Their role, procedures, assessment and remedies](#) (CC3(Revised)) paragraph 18.

⁸ EA02, sections 134(1) and 134(2B). For present purposes, 'relevant market' means a market in the United Kingdom for goods or services of a description specified in the reference (EA02, section 134(3)(b)) and 'relevant feature' means any feature of a relevant market which is not excluded from investigation by virtue of section 133(1A), or a combination of such features (EA02, section 134(2B)). As per section 131(2) EA02, a feature can be construed as a 'reference to the structure of the market; any conduct (whether or not in the market concerned) of one or more than one person who supplies or acquires goods or services in the market concerned; or any conduct relating to the market concerned of customers of any person who supplies or acquires goods or services'.

⁹ EA02, section 134(2).

- 1.7 The analysis focuses on market outcomes and the features that produce those outcomes, rather than evaluating the appropriateness of individual business decisions within existing market frameworks.¹⁰ The process is ‘investigative and inquisitorial, not accusatorial’, so involvement in a market investigation or being subject to remedies does not imply that participants have infringed the law or are suspected of wrongdoing.¹¹
- 1.8 Where an AEC and its causes are identified, we are required to decide whether we should take action, or whether we should recommend the taking of action by others, for the purpose of remedying, mitigating or preventing the AEC concerned or any detrimental effect on customers so far as it has resulted from, or may be expected to result from, the AEC.¹² Such remedies may require changes across entire sectors or recommendations for modifications to regulatory frameworks, reflecting the market improvement mandate of the investigation process.¹³
- 1.9 In choosing the appropriate remedial action, we are required to have regard to ‘the need to achieve as comprehensive a solution as is reasonable and practicable to the adverse effect on competition and any detrimental effects on customers so far as resulting from the adverse effect on competition’¹⁴ and may have particular regard to the effect of any action on any relevant customer benefits of the feature or features of the market(s) concerned.¹⁵

This market investigation reference

- 1.10 Prior to its reference to the CMA, Ofcom had undertaken a market study into cloud services with the intention to gain a better understanding of what it considered a critical component of the digital economy and to gather evidence to inform an assessment of whether competition is working well for consumers and citizens in the UK.¹⁶ Ofcom focused its study on the market for public cloud infrastructure services.
- 1.11 In April 2023 alongside the publication of its interim report, Ofcom published a notice and consultation on a proposal to make a market investigation reference to the CMA into the supply of public cloud infrastructure services in the UK.¹⁷

¹⁰ Per section 134(5) EA02, there is a detrimental effect on customers if there is a detrimental effect on customers or future customers in the form of: (a) higher prices, lower quality or less choice of goods or services in any market in the UK (whether or not the market(s) to which the feature or features concerned relate); or (b) less innovation in relation to such goods or services.

¹¹ CC3 Revised, paragraph 21.

¹² EA02, section 138(2).

¹³ CC3 Revised, paragraph 20.

¹⁴ EA02, section 134(6).

¹⁵ EA02, section 134(7).

¹⁶ Ofcom, [Cloud services market study final report](#), paragraph 2.4.

¹⁷ Ofcom, [Notice of a proposal to make a market investigation reference under section 131 of the Enterprise Act 2002](#) and Ofcom, [Consultation: Proposal to make a market investigation reference into the supply of public cloud infrastructure services in the UK](#).

- 1.12 In October 2023 Ofcom published its market study final report and its decision to refer the UK public cloud infrastructure services market to the CMA for a market investigation. Ofcom found that it had reasonable grounds to suspect that a feature or a combination of features of the markets for the supply of public cloud infrastructure services in the UK prevents, restricts or distorts competition.¹⁸

The role of public cloud infrastructure services in the UK today

- 1.13 Public cloud infrastructure services are vital inputs to many businesses and organisations across the UK economy and so are vital to economic growth in the UK.
- 1.14 We have sought views of a wide range of UK companies which use public cloud infrastructure services including major firms in banking, retail, energy, media and communications, transport and a range of other industries, all of whom told us how these services support their activities. Without public cloud infrastructure services many businesses providing digital services to consumers would not be able to function in the way they do today.
- 1.15 We have been told by these customers that cloud services bring many benefits compared to when they hosted IT services on-premises: cloud services are reliable, scalable, easy to use and maintain; providers support their customers and are innovative. These services may also be lower in cost for most use cases than on premises or private cloud.¹⁹

Scope and focus of the investigation

- 1.16 Public cloud infrastructure services provide access to computing resources on demand, via a network. Customers buy access to the computing resources as a service and typically do not own the underlying hardware and software.
- 1.17 In this market investigation we have considered the supply of public cloud infrastructure services in the UK. This refers to services that are open to all customers, with computing resources shared between them (public cloud computing) and which 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).²⁰

¹⁸ Ofcom's [Terms of reference](#), 05 October 2023.

¹⁹ See [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), page 25.

²⁰ Ofcom's Cloud services market study final report, [Terms of Reference](#), 5 October 2023 (Terms of Reference).

- 1.18 These public cloud infrastructure services are referred to throughout this report as cloud services. Software as a service (SaaS) does not form part of this definition.²¹
- 1.19 Cloud services are differentiated by the level of control the customer has over the management and maintenance of the computing resources. IaaS and PaaS together with SaaS form a vertical 'cloud stack', where each layer is notionally built on top of the previous one(s).²²
- (a) The IaaS layer provides access to raw computing resources (compute, storage and network) for processing workloads and storing data.²³ The hardware associated with these computing resources takes 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 and virtualisation.²⁴ The IaaS service model provides the customer with the highest level of control over the cloud stack, including over the operating system, applications and data.
 - (b) The PaaS layer provides access to a virtual environment for customers to develop, test, deploy and run applications. They include application development computing platforms and pre-built application components and tools which customers can then use to build and manage full applications. The customer has less control over the cloud stack compared to IaaS – customers still manage applications and data but not the PaaS computing platform (including its operating system) and the pre-built application components and tools.
 - (c) The SaaS layer comprises complete applications hosted in the cloud. Like PaaS, they can be offered by the cloud provider that owns the underlying raw compute resources or by an independent software vendor (ISV). The service provider(s) manages all hardware and software.
- 1.20 In addition to public cloud computing, there are two other cloud deployment models:
- (a) private cloud – a cloud deployment model in which computing resources (like the hardware) are used exclusively by one customer; and

²¹ Some services may not 'fit' neatly into these service models and the lines between each of IaaS, PaaS and SaaS may be blurred. However, we still consider them to be useful to inform our analysis in this market investigation. For an explanation of SaaS, please see paragraph 1.19 (c).

²² In practice, this vertical stack is not strictly applied. For example, SaaS may be built and deployed using IaaS only.

²³ A workload is a general term which may mean a customer's application, service, capability or other task or activity.

²⁴ Virtualisation is the process of using software to create an abstraction layer over servers that allows the hardware elements of a single server to be divided into multiple virtual servers, commonly called virtual machines. 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.

(b) hybrid cloud – a cloud deployment model in which public and private clouds are combined.

- 1.21 Cloud computing is distinct from traditional IT where assets (such as servers and network hardware) are located on-premises and managed by the end user.
- 1.22 Our market investigation has focused on public cloud infrastructure services, however other cloud deployment models and traditional IT have been considered where relevant.
- 1.23 On the basis of Ofcom’s findings and our guidelines on potential sources of harm we have focused our investigation on: analysing the structure of and competitive dynamics in cloud services (including the impact of AI); the barriers suppliers face when looking to enter or expand in cloud services; technical and commercial barriers to customers switching cloud provider or using multiple clouds; and the impact of Microsoft’s software licensing practices on competition.
- 1.24 The scope and focus of our investigation have been shaped by the evidence we have received including the submissions received in response to our issues statement, working papers, updated issues statement, PDR and from the main parties in formal hearings after the publication of our working papers and PDR.²⁵ Where relevant, we have taken these submissions into account throughout our assessment as appropriate. Where submissions raised issues that we consider go beyond the scope of our investigation we have noted this as appropriate.

Our approach to evidence gathering

- 1.25 During our investigation, we have requested information and cooperation from businesses operating in the relevant markets. This information-gathering served to develop our understanding of how these markets function and the effects of market features on competition.
- 1.26 We have gathered a broad range of evidence and consulted a large number of stakeholders. Stakeholders included a range of UK customers from both private and public sectors, all of the UK’s main cloud providers, intermediaries, ISVs, trade associations, academics and individuals.
- 1.27 Our approach to evidence gathering in this market investigation, including our procurement of primary market research, focused primarily on qualitative evidence instead of quantitative evidence due to the complexity of the relationship between cloud providers and customers and the purchasing journey for cloud services. We set out our position on quantitative survey evidence including the concerns about using quantitative surveys in this market due to lack of validity and vulnerability to

²⁵ The main parties to our inquiry were Microsoft, AWS and Google. See further detail on this engagement below

the quality and coverage of the customers spoken to in chapter 3.²⁶ When describing responses to our requests for information to, or calls with, customers in particular, we use the terms ‘a few’, ‘some’ and ‘many’ in a primarily qualitative way and we do not claim that these responses are statistically representative of the broader universe of UK cloud customers.

Information requests and customer calls

- 1.28 Throughout the investigation we held calls and issued formal and informal information requests on a wide range of topics to cloud providers, customers and other stakeholders.²⁷ We gathered further evidence across a range of topics following our PDR.

Consultations

- 1.29 We consulted on our emerging thinking in working papers and an updated issues statement and our provisional decisions in the PDR:
- (a) On 17 October 2023 we set out the areas on which we intended the investigation would focus in an issues statement.²⁸
 - (b) Between 23 May 2024 to 06 June 2024, we published six working papers setting out our emerging thinking and an updated issues statement.²⁹
 - (c) On 28 January 2025, we published our PDR setting out our provisional decisions on competition and remedies.³⁰
- 1.30 We invited all interested parties to make submissions in response to the above consultations and non-confidential versions of these submissions were published on our case page.³¹ We have carefully considered all submissions and evidence received.

Engagement with main parties

- 1.31 We considered AWS, Google and Microsoft to be the main parties to this investigation and we have ensured throughout the investigation that these parties

²⁶ As set out in appendix C.

²⁷ Given the heterogeneity of the market, information requests sent to customers and other external parties were not identical in terms of focus and specific questions asked.

²⁸ [Issues statement](#), 17 October 2023.

²⁹ [Updated issues statement](#), 06 June 2024; [Technical barriers working paper](#), 06 June 2024; [Licensing practices working paper](#), 06 June 2024; [Potential remedies working paper](#), 06 June 2024; [Competitive landscape working paper](#), 23 May 2024; [Egress fees working paper](#), 23 May 2024; [Committed spend agreements working paper](#), 23 May 2024.

³⁰ [Provisional Decision Report \(PDR\)](#), 28 January 2025.

³¹ [Cloud services market investigation - GOV.UK](#)

have had the opportunity to present all material facts and evidence, and that they have been able to make submissions in support of their respective positions.

1.32 We have engaged extensively with these main parties throughout our inquiry:

- (a) early in the investigation, we conducted site visits at the premises of AWS, Google and Microsoft where they provided presentations to us on the nature of their businesses and the market context in which they operate;
- (b) we held hearings with each party in July 2024 and again in March 2025 in which we explored a range of issues raised in the working papers and PDR respectively. Summaries of these hearings have been published on our case page;³²
- (c) we shared substantial parts of our analysis with the external legal and economic advisors of these parties in confidentiality rings. We also disclosed into confidentiality rings an unredacted version of the PDR, relevant underlying analysis as well as unredacted versions of certain submissions we had received; and
- (d) following the main party hearings in May 2025 we also held meetings with the main parties and provided them with an update on our evidence gathering and certain areas of focus since publication of the PDR, and we provided a short written document to these parties setting this update out.

Requests for disclosure of certain documents and information

1.33 On a number of occasions, both prior to and following the publication of the PDR, the main parties wrote to us requesting disclosure of certain underlying documents and evidence referred to in working papers, the PDR and submissions from the main parties.

1.34 The publication of the PDR, and the disclosure of relevant materials into confidentiality rings where appropriate, is the principal means by which we consulted parties likely to be substantially impacted by our decisions in this investigation.³³

1.35 The PDR set out a substantial body of evidence and our provisional assessments in considerable detail. We also provided a further update to the main parties on certain areas of our focus since publication of the PDR, including our additional

³² [Cloud services market investigation](#).

³³ Under section 169 EA02, where the CMA is proposing to make a relevant decision in a way which the CMA considers is likely to have a substantial impact on the interests of any person, the CMA must, so far as practicable, consult that person about what is proposed before making that decision. In consulting the person concerned, the CMA must, so far as practicable, give the reasons for the proposed decision.

evidence gathering.³⁴ Overall, we have provided sufficient gist of the main issues arising to enable the main parties to provide their views and evidence. We also considered further requests for disclosure from parties on a case-by-case basis throughout the investigation and disclosed materials where appropriate.

Jigsaw research

- 1.36 In addition to the evidence we gathered ourselves, we commissioned qualitative customer research from Jigsaw Research (Jigsaw).³⁵ This research has captured a different set of customers from those we contacted directly. The in-depth interviews allowed for an informative insight into customers' views on the issues investigated. Evidence from this research is set out in the Jigsaw report.³⁶ We draw on the evidence set out in the Jigsaw report where relevant throughout this report.

Approach to our assessment and the structure of this document

- 1.37 In our issues statement we set out four theories of harm based on both the structure of the market(s) we investigated and the conduct of relevant firms within these or related markets. These theories provided a useful framework for our evidence gathering and analysis. We set out our findings on each of these theories of harm separately³⁷ and we have considered them in the round.
- 1.38 The structure of this report reflects our assessment and is as follows:
- **Chapter 2** provides the industry background by introducing the main providers of cloud services and how they compete and by setting out the demand for cloud services by customers and how customers purchase these services.
 - **Chapter 3** describes the competitive landscape by setting out our approach to and findings on market definition, the structure of the cloud services markets, the level of concentration and the market outcomes we observe.
 - **Chapter 4** sets out our findings on the impact of AI on competition in cloud services.
 - **Chapter 5** provides our assessment of the barriers to entry and expansion that we have identified in cloud services.

³⁴ Update paper sent to AWS, Google and Microsoft on 16 May 2025.

³⁵ While our evidence gathering focused on large customers as they represent the majority of revenues for providers of cloud services, the research commissioned to Jigsaw Research complemented this evidence gathering by also engaging with smaller customers, eg start-ups.

³⁶ [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024).

³⁷ See chapters 6, 7 and 8.

- **Chapter 6** provides our assessment of the technical and commercial barriers (egress fees) to customers' switching between public clouds and their use of multiple clouds.
- **Chapter 7** provides our assessment of the impact of Microsoft's licensing practices on competition in the cloud services markets;
- **Chapter 8** considers the impact of committed spend agreements;
- **Chapter 9** brings together our findings on the features identified and whether individually or in combination they lead to an AEC in the UK; and
- **Chapter 10** sets out the action that should be taken to remedy, mitigate or prevent the AECs we have found.

Next steps

- 1.39 This document together with its appendices constitutes our final decision. It refers, where appropriate, to material published separately on the CMA website. The report, however, is self-contained and provides all material necessary for an understanding of our findings.

2. Industry background

- Cloud services are now a vital input to businesses and organisations across the UK economy, with £10.5 billion spent on them in 2024. They support many sectors' contribution to the UK's overall economic growth and it is therefore vital that competition works well in these markets for the benefit of these businesses and the wider UK economy.
- Cloud services revenues in the UK have grown significantly in recent years: over the period 2020-2024, providers' revenues in IaaS and PaaS have grown at an annual average rate of 29%. That growth is expected to continue although there is some indication that it could be at a slightly slower rate.
- Cloud services are used by a large number of customers of all sizes and across most industries in the UK. However, a small number of high-spend customers are responsible for a significant proportion of cloud providers' UK revenues.
- Customers buy cloud services through a variety of channels and consider a range of factors when purchasing cloud services depending on their characteristics and needs.
- Switching and multi-cloud (the use by a single customer of more than one cloud) require customers to take a series of steps and incur costs which may affect customers' ability and/or incentive to do so. Decisions on whether to switch and/or multi-cloud as well as the level of integration to adopt are driven by how customers weigh the different benefits and disadvantages.
- The specific benefits and disadvantages of switching and multi-cloud depend on customers' circumstances: their needs, preferences and use cases. The extent to which barriers to switching or multi-cloud prevent the adoption of the customers' preferred provider or types of multi-cloud is important to understand when considering whether there are any sources of harm to competition.

- 2.1 This chapter provides some background information on cloud services. It describes customers' demand for cloud services, the market position of the providers, how customers purchase cloud services, the parameters of competition, and switching and multi-cloud.
- 2.2 In a market investigation we need to understand how a market operates in practice as this allows us to apply the appropriate framework to our analysis. We collect and analyse evidence about the main characteristics of the relevant market(s) and

use this to inform our assessment of whether any features may be harming competition.³⁸

- 2.3 In order to understand the main characteristics of the reference market(s), this chapter sets out our assessment and conclusions on:
- (a) the nature of the customer base, such as the types of customers that are purchasing cloud services and trends in the usage of these services;
 - (b) the main providers, focusing on the vertically integrated suppliers of cloud services³⁹ (which we refer to as ‘providers’) including their business models and strategies;
 - (c) the customer journey, including how cloud services are purchased;
 - (d) customer preferences and the parameters of competition; and
 - (e) the steps necessary for a customer to switch provider, the different ways to use multiple providers, and the benefits and disadvantages associated with it.
- 2.4 More details, including a detailed discussion of the evidence underlying our findings, are included in appendix A: Demand for public cloud, and appendix B: Parameters of competition.

The demand for cloud services

Customers of cloud services

- 2.5 Evidence from cloud providers shows that there is a large number of customers of cloud services in the UK.⁴⁰ This includes customers which have transferred some or all their IT infrastructure from on-premises to the cloud – which has driven most of the sector’s growth in the past – and customers which have started new workloads directly on cloud.⁴¹
- 2.6 Customers are present in a range of different industries, including some with specialised use cases due to regulatory requirements (eg financial services) or procurement frameworks (as in the public sector).⁴² Our analysis of cloud

³⁸ See [CC3 \(Revised\)](#), [Guidelines for market investigations: Their role, procedures, assessment and remedies](#) (CC3 (Revised)) paragraph 97.

³⁹ That is, vertically integrated suppliers of cloud services that operate their own cloud infrastructure, ie they own the underlying raw computing resources.

⁴⁰ Responses to the CMA’s information requests [38].

⁴¹ Digital native customers are businesses that exist primarily or entirely online and use cloud technologies throughout their operation. Gartner (2019), [Cloud Shift Impacts All IT Markets](#), accessed on 22 October 2024.

⁴² See chapter 5.

providers' revenues shows that customers in financial services, IT or software services and retail are important categories.

- 2.7 Large companies are the major customers for cloud providers: evidence shows that a small number of high-spend customers account for a significant proportion of providers' UK revenues and a large number of low-spend customers are responsible for a small proportion of their revenues. In particular, for AWS, Microsoft and Google, the top [10-20]% of customers account for a very large majority of revenues and the top [0-5]% account for over half of revenues.⁴³

Growth in demand for cloud services

- 2.8 In order to understand better how cloud services may develop in the future, we have looked at trends in customer spending, cloud provider revenues and data centre capacity.
- (a) An industry report from Gartner indicated that worldwide customer spending for cloud services would grow significantly between 2023 and 2025.⁴⁴
- (b) A submission by one provider shows that it expects year-on-year revenue growth from its cloud services to continue but at a lower rate of growth.⁴⁵
- 2.9 We have estimated revenues from IaaS and PaaS in the UK from 2020 to 2024.⁴⁶ Over the period 2020-2024, the IaaS and PaaS markets have grown at a combined annual average rate of 29%.

Table 2.1: UK IaaS and PaaS revenues, 2020-2024 (£bn)

	2020	2021	2022	2023	£bn 2024	% Annual growth
IaaS	2.4	3.0	4.4	5.2	6.0	25.4
PaaS	1.3	1.8	2.8	3.6	4.4	35.9
IaaS and PaaS	3.7	4.8	7.2	8.8	10.5	29.4

Source: CMA analysis of data from IaaS and PaaS providers, IDC and Synergy. Numbers may not sum due to rounding.

- 2.10 The annual growth rates for the UK IaaS and PaaS markets from 2023 to 2024 were slightly lower than in the previous years. This, together with the evidence above, suggests that, while growth is expected to continue, it could be at a slightly slower rate.
- 2.11 We have looked at data on data centre capacity because it may indicate providers' expectations related to growth in customers' demand: if providers forecast

⁴³ Data refers to customers in the UK with a spend greater than \$1,000 in 2022. Responses to the CMA's information requests [38].

⁴⁴ [Gartner Forecasts Worldwide Public Cloud End-User Spending to Surpass \\$675 Billion in 2024](#), accessed on 28 August 2024. Note that Gartner's definition of public cloud services includes other segments beside IaaS and PaaS (eg SaaS, Cloud Business Process Services (BPaaS) and Cloud Desktop-as-a-Service (DaaS)).

⁴⁵ [38] response to the CMA's information request [38].

⁴⁶ CMA analysis of data from IDC, Synergy and various IaaS and PaaS suppliers. See section market structure and concentration in chapter 3 for more details on methodology.

increases in capacity, then this is consistent with an expectation among them that demand for cloud services will continue to increase.

- 2.12 Data from AWS, Microsoft, Google, IBM, Oracle, CoreWeave and other small providers⁴⁷ shows that data centre capacity has increased significantly since 2020 and is forecast to increase substantially in the near future.
- 2.13 As discussed in appendix K, public sector customers are likely to grow their use of cloud services with the majority of workloads still to be migrated and UK government policy explicitly promoting further migration of services to the cloud.⁴⁸

Providers' market positions

- 2.14 We have gathered evidence from customers, providers, industry reports and providers' internal documents on the position of the main providers of cloud services in the UK. Below and in appendix A, we set out the relevant information on each of these providers, as well as relevant information on their business strategies, including their pricing strategies and their customer base. We have also considered evidence on their business strategies from analyst reports, internal documents and customer views on each provider.⁴⁹ These show a sector subject to both challenges and opportunities driven by technological developments (eg in the area of AI) and evolving customer demand.
- 2.15 Overall, AWS and Microsoft are considered the leading providers of cloud services by customers and other providers. Analysts see Google as a smaller cloud provider than AWS and Microsoft, although they recognise that it is expanding across IaaS and PaaS. Large customers do not see Oracle, IBM or other providers as suitable alternatives to their main providers (generally AWS and Microsoft). Smaller providers are generally not seen as an effective alternative to large customers' main providers, but we note that they may be seen as suitable alternatives for certain workloads.

AWS

- 2.16 AWS, a subsidiary of Amazon, started providing cloud services in 2006.⁵⁰ AWS was the first provider of cloud services. It offers a wide variety of cloud services and has a global coverage.
- 2.17 The customers who participated in Jigsaw's research perceived AWS as 'the overall leader, in terms of the level of innovation and range and quality of services provided'. Customers also commonly saw AWS as a 'first mover' in the provision

⁴⁷ See appendix A.

⁴⁸ Note of meeting with [§].

⁴⁹ More details can be found in appendix A

⁵⁰ [About AWS](#), accessed on 10 February 2024.

of cloud services and while ‘there was a sense that other public cloud providers are catching up’, AWS still retained ‘a reputational advantage’.⁵¹

- 2.18 The customers who participated in Jigsaw’s research identified the following as AWS’ strengths: being an ‘innovator’, being ‘a requirement of doing business’, being ‘reliable’, its ‘ease of use and maintenance’ and ‘ease of set-up’, being ‘good for some uses/workloads’ and the level of ‘support and advice’. Customers also identified certain challenges to using AWS, including ‘greater risk of ‘lock in’ than competitors’, its ‘complex billing’ and ‘cost’, and the fact that ‘Amazon can be a competitor’ to some firms.⁵²
- 2.19 AWS identified the following strategic priorities as ‘the most important things for [its own] business to get delivered or done in 2024’: [REDACTED].⁵³
- 2.20 Other cloud providers described AWS as being strong across customer sizes, and particularly focused on startups and digital natives (ie customers without a pre-existing on-premises IT setup).⁵⁴

Microsoft

- 2.21 Microsoft started providing cloud services in 2008 through Windows Azure,⁵⁵ but made its cloud services offering more widely available in 2010.⁵⁶ It offers a large number of IaaS, PaaS and SaaS services, and has a global reach.
- 2.22 The customers who participated in Jigsaw’s research said that Microsoft was particularly strong for whole-enterprise solutions and saw Azure as a leading cloud platform, second only to AWS. Customers identified the following as Microsoft’s strengths: ‘familiarity and ease of set up’, being a ‘requirement of doing business’, its ‘good integration with other Microsoft products’, being ‘good for some uses/workloads’ and offering ‘good value for money via Enterprise Agreements’. They also noted that the key challenge to using Microsoft’s cloud was the ‘risk of ‘lock in’ from Enterprise Agreements and historical relationships’.⁵⁷
- 2.23 Overall, Microsoft’s internal documents show that the provision of AI-related services (including accelerated compute and access to FMs) is becoming a key priority for its cloud business and is allowing Microsoft to target digital natives to its cloud, a customer demographic it was not traditionally strong with.⁵⁸

⁵¹ [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), page 28.

⁵² [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), page 28.

⁵³ AWS’ response to CMA’s information request [REDACTED].

⁵⁴ Responses to the CMA’s information requests [REDACTED].

⁵⁵ [About Microsoft - Stories](#), accessed on 11 February 2024.

⁵⁶ [Windows Azure Platform Now Generally Available in 21 Countries](#), accessed on 11 February 2024.

⁵⁷ [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), page 28

⁵⁸ Microsoft’s responses to the CMA’s information requests [REDACTED].

- 2.24 Other cloud providers submitted that Microsoft has had the greatest success winning enterprise customers, with one suggesting that this is due to its licensing practices, and with another submitting that Microsoft's productivity software is critical for most enterprise IT customers and is software that these customers expect to be able to access.⁵⁹ We cover Microsoft's licensing practices in more detail in chapter 7.

Google

- 2.25 Google started providing cloud services in 2008 through Google App Engine, but made its offering in cloud services more widely available in 2011.⁶⁰ Since then, it expanded to offer a large number of IaaS, PaaS and SaaS services and has a global coverage.⁶¹
- 2.26 Google was the main provider for a number of customers we contacted. Some other customers identified Google as an effective alternative to their main provider, but Google was generally identified as not being a fully effective alternative to AWS and Azure. The most common reasons for customers not seeing Google as being a fully effective alternative were Google not having as advanced or as broad a range of functionality / features / services as the customers' main providers (AWS or Microsoft),⁶² or having no direct experience of Google's offering or a lack of internal skills to use Google's cloud services.⁶³
- 2.27 Following publication of our PDR, we asked the same large customers about Google's suitability as an alternative to its main public cloud provider and only one of these customers said that their view of Google had improved (citing the availability of Google's first-party AI models as a reason).⁶⁴
- 2.28 The customers who participated in Jigsaw's research thought of Google as a 'highly reputable' cloud provider, 'strong on analytics and AI'. Google had a smaller pool of users among the customers interviewed by Jigsaw and these were 'more likely to be tech-driven businesses such as start-ups, e-commerce, fintech or adtech companies'.⁶⁵
- 2.29 These customers identified the following benefits of using Google's cloud platform: 'familiarity' for those who use other Google services, access to 'BigQuery' (a Google PaaS analytics service), 'good integration with other platforms', being 'easy to use and manage', 'great support for start-ups', and being superior – in

⁵⁹ Responses to the CMA's information requests [redacted].

⁶⁰ [The History of Google Cloud Platform](#), accessed on 11 February 2024.

⁶¹ [Cloud Computing Services](#), accessed on 11 February 2024; [Global Locations - Regions & Zones](#), accessed on 11 February 2024.

⁶² Responses to the CMA's information requests [redacted].

⁶³ Responses to the CMA's information requests [redacted].

⁶⁴ [redacted] response to the CMA's information request [redacted].

⁶⁵ [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), page 27.

particular to AWS – in terms of its AI offering. Customers identified two key challenges to using GCP: its complex billing and the lack of configurability relative to competitors like AWS.⁶⁶

2.30 Google's internal documents show that Google's strategic priorities for its cloud business are:

- (a) Leading in the AI space: in particular, an internal document [REDACTED] for Google Cloud, Google stated that [REDACTED].⁶⁷
- (b) In another document [REDACTED], Google stated that [REDACTED].⁶⁸
- (c) winning startups, public sector, FTSE 100 & SI customers; and
- (d) expanding Google's partners network.⁶⁹

2.31 Microsoft described Google as a strong and growing competitor,⁷⁰ and AWS said that Google has achieved success in respect of its 'Big Data' analytics tools.⁷¹

IBM

2.32 IBM first started providing SaaS offering around 2008. It has subsequently expanded to provide a range of IaaS, PaaS and SaaS products.⁷² IBM offers a variety of products⁷³ and is active in 10 global regions.⁷⁴ IBM also has an offering, IBM Cloud Satellite, through which it provides certain cloud services on other public clouds and on-premises.⁷⁵ IBM is also active in hardware for on-premises solutions.⁷⁶

2.33 IBM was not the main provider for any of the customers we contacted. Amongst other customers, none identified IBM as an effective alternative to their main provider, with most identifying it as an ineffective or very ineffective alternative. The two main reasons given for these ratings for IBM (or for not giving a rating) were that IBM has more limited services or capabilities compared to larger providers,⁷⁷ and a lack of any experience or knowledge of IBM's offering on the part of the customer.⁷⁸

⁶⁶ [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), page 30.

⁶⁷ Google's response to the CMA's information request [REDACTED].

⁶⁸ Google's response to the CMA's information request [REDACTED].

⁶⁹ Google's responses to the CMA's information requests [REDACTED].

⁷⁰ [Microsoft's response to the CMA's PDR dated 28 January 2025](#), paragraphs 90 and 93; [AWS' response to the CMA's PDR dated 28 January 2025](#), paragraph 13.

⁷¹ [AWS' response to the CMA's PDR dated 28 January 2025](#), paragraph 5.

⁷² [What is IBM Cloud? Services Offered, Features & Pricing](#), accessed on 11 February 2024.

⁷³ [Cloud Products - IBM](#), accessed on 11 February 2024.

⁷⁴ [Locations for resource deployment - IBM](#), accessed on 11 February 2024.

⁷⁵ [IBM Cloud Satellite](#), accessed on 11 February 2024.

⁷⁶ [Enterprise Business Server Solutions - IBM](#), accessed on 11 February 2024.

⁷⁷ Responses to the CMA's information requests [REDACTED].

⁷⁸ Responses to the CMA's information requests [REDACTED].

- 2.34 The customers interviewed by Jigsaw saw IBM as a ‘secondary’ cloud. The types of customers using IBM were ‘more traditional [businesses], such as those involved in energy, health, regulatory support or insurance’. Customers identified the following benefits of using IBM’s cloud: ‘flexibility’, ‘easy financial management’ and ‘good governance’. Customers described IBM ‘as falling behind competitors’, and ‘some users had or were planning to switch away from using IBM’.⁷⁹
- 2.35 IBM identified several competitive factors. [REDACTED].⁸⁰

Oracle

- 2.36 Oracle entered cloud services with Oracle Cloud Infrastructure (OCI) in 2016 and has expanded to provide a wide range of services across IaaS, PaaS and SaaS.⁸¹ Oracle offers a smaller set of services than other providers and is active in 48 commercial and government regions.⁸²
- 2.37 The customers interviewed by Jigsaw saw Oracle as a ‘secondary’ cloud, ‘only relevant for certain workloads’. Like for IBM, the types of customers using Oracle’s cloud were ‘more traditional [businesses], such as those involved in energy, health, regulatory support or insurance’. A few customers ‘mentioned using Oracle for specific legacy systems that only run on Oracle cloud infrastructure or work more effectively if they do’.⁸³
- 2.38 Oracle also identified AI as an area of emerging focus.⁸⁴ On this point, another provider submitted that ‘similar to other providers, Oracle has worked aggressively to expand its range of AI and machine learning products and services with management continuing to highlight its technical differentiation for AI workloads’.⁸⁵
- 2.39 Cloud providers submitted that Oracle has fuelled its cloud business by leveraging its legacy incumbency in on-premises business, especially in data bases.⁸⁶

Other providers

- 2.40 There are a range of smaller providers offering IaaS and PaaS products, such as OVHcloud, Scaleway and others. The evidence we have seen from customers and analyst reports, however, indicates these are not considered strong alternatives to the leading providers.

⁷⁹ [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), page 31.

⁸⁰ IBM’s response to the CMA’s information request [REDACTED].

⁸¹ [Oracle Cloud Infrastructure Platform Overview](#), accessed on 11 February 2024.

⁸² [Public Cloud Regions and Data Centers - Oracle United Kingdom](#) accessed on 11 February 2024.

⁸³ [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), page 32.

⁸⁴ Oracle’s response to the CMA’s information request [REDACTED].

⁸⁵ [REDACTED] response to the CMA’s information request [REDACTED].

⁸⁶ Responses to the CMA’s information requests [REDACTED].

- 2.41 There are other cloud providers who may have a large presence globally, such as Alibaba, Huawei and Tencent, which have been identified as a competitor by at least one cloud provider.⁸⁷ However, these providers have not been identified as suitable alternatives to the main providers by UK customers we have contacted.⁸⁸
- 2.42 Another category of suppliers of cloud services is Independent Software Vendors (ISVs). ISVs are suppliers of cloud services, typically PaaS and/or SaaS, that do not usually own the underlying infrastructure.⁸⁹ As such, they tend to rely on cloud providers as an input to their own cloud services. That is, they may use cloud providers' IaaS as an input to develop and offer their own PaaS and/or SaaS. Examples of ISVs include VMware, MongoDB, Snowflake and Yugabyte.
- 2.43 ISVs can be considered, depending on the context, as customers, competitors, or partners (eg provider of complementary services) to cloud providers. ISVs and cloud providers may also have a relationship as ISVs may rely on cloud providers as a distributor of their services.
- 2.44 There are many ISVs providing PaaS in the UK and our share of supply analysis suggests that ISVs accounted for up to [30-40]% of UK PaaS revenue in 2024.⁹⁰
- 2.45 ISVs compete in specific product categories rather than across the entire range of PaaS products.

Our assessment of the main providers

- 2.46 Microsoft and AWS are considered the leading providers of cloud services by customers and other providers.
- 2.47 While some customers see Google as an effective alternative to their main providers, many are neutral. The most common reasons for customers not seeing Google as a fully effective alternative were Google not having as advanced or as broad a range of functionality / features / services as the customers' main providers (AWS or Microsoft), or having no direct experience of Google's offering or a lack of internal skills to use Google's cloud services.
- 2.48 Large customers do not see Oracle, IBM or other providers as effective alternatives to their main providers (generally AWS and Microsoft). Smaller providers are generally not seen as an effective alternative to large customers'

⁸⁷ Responses to the CMA's information requests [REDACTED]; [REDACTED] response to Ofcom's information request [REDACTED].

⁸⁸ For example, Alibaba received a very low rating by customers and was not identified by any customer as an effective alternative. Reasons for these ratings or not providing a rating included a lack of experience or knowledge of Alibaba's offering, that using Alibaba would be a supply chain risk or not appropriate (eg for data sovereignty or security reasons, that Alibaba has more limited services or capabilities compared to larger providers and that they are only considered for demand relating to China. Responses to the CMA's information requests [REDACTED].

⁸⁹ There are some exceptions to this – for example, Salesforce is an ISV, but operates its own infrastructure.

⁹⁰ CMA analysis of data from IDC, Synergy and various PaaS suppliers. See section market structure and concentration in chapter 3 for more details on methodology.

main providers, but we note that they may be seen as suitable alternatives for certain workloads.⁹¹

How customers purchase cloud services

- 2.49 The ways in which services are purchased can influence the nature of competition. In particular, where purchases are individually negotiated, the details of these negotiations can influence the competitive process as customers may vary in their needs and their ability to negotiate, and providers may tailor their negotiating strategies to particular customers.
- 2.50 There are a variety of purchasing channels through which customers can access cloud services. The use of each of these channels depends on the customers' characteristics and needs:
- (a) Smaller enterprise customers are more likely to purchase cloud services directly from providers through their online portals⁹² and/or their marketplaces.^{93 94} These customers generally pay for cloud services on a pay-as-you-go basis, paying the providers' listed prices.⁹⁵
 - (b) Some customers purchase cloud services through resellers who are a network of partners authorised to resell cloud services from providers.⁹⁶
 - (c) Large enterprise customers – for example, classified by one provider as those with an estimated spend of over £1 million per year⁹⁷ – generally procure cloud services through bilateral negotiations with providers.⁹⁸ A cloud provider said that this allows a range of customers, including those with higher annual contract values, to secure bespoke contracts tailored to their needs.⁹⁹
 - (d) A minority of customers carry out competitive tenders: these are particularly prevalent among public sector customers which often have specific requirements on how to procure IT services.¹⁰⁰ Public sector procurement is discussed in chapter 5 Barriers to entry and expansion. Tenders are less frequent in the private sector.

⁹¹ For more details on customer views of Oracle, IBM and other providers, see the section 'Main providers' in appendix A.

⁹² Online portals refer to the provider's website where customers can purchase cloud services at listed prices.

⁹³ Marketplaces are an online platform, where cloud providers and ISVs can offer services to customers, which run on the underlying infrastructure of the provider offering that marketplace. Both Google and Microsoft said customers can procure cloud services through their marketplaces. It is possible to buy services also through AWS' marketplace. [§] response to CMA's information request [§]; Responses to Ofcom's information requests [§].

⁹⁴ [§] response to CMA's information request [§].

⁹⁵ [§] response to CMA's information request [§].

⁹⁶ Responses to Ofcom's information requests [§].

⁹⁷ [§] response to CMA's information request [§].

⁹⁸ Responses to the CMA's information requests [§]; [§] response to Ofcom's information request [§].

⁹⁹ [§] response to CMA's information request [§].

¹⁰⁰ Responses to Ofcom's information requests [§].

- (e) Some providers offer incentive programmes to resellers to promote sales of certain services; for example, one provider offers payments for hitting certain sales thresholds,¹⁰¹ and another offers discounts for resellers purchasing its cloud services.¹⁰²

2.51 In particular, there are two main types of contracts that customers can enter into when purchasing cloud services: customer agreements and enterprise agreements. Some larger customers either engage in bilateral negotiations or tenders and are able to negotiate terms that depart from standard contracts.

- (a) Customer agreements are the standard contracts cloud providers offer through their online marketplaces. When customers have a customer agreement, they typically incorporate the cloud provider's listed terms and prices. Three providers said that the vast majority of customers are on this type of contract and do not negotiate additional terms.¹⁰³
- (b) Enterprise agreements, which are individually negotiated, are generally reserved for larger customers with higher spending.¹⁰⁴ One provider submitted that they contain commonly requested terms such as invoicing and regulatory compliance commitments.¹⁰⁵

2.52 We describe further evidence on how customers purchase cloud services and contract negotiations in appendix A.

Parameters of competition

2.53 We refer to the ways in which providers flex their offerings to meet customer preferences as the parameters of competition. Below, we summarise the evidence we have seen from both providers and customers on the relative importance of various parameters.¹⁰⁶

Evidence from providers

2.54 Cloud providers identified a variety of parameters of competition. They can be divided into price and non-price factors.

- (a) AWS, Microsoft, Google, IBM and OVHcloud identified price as a key parameter of competition for cloud services.¹⁰⁷ Providers compete on price

¹⁰¹ [redacted] response to Ofcom's information request [redacted].

¹⁰² [redacted] response to Ofcom's information request [redacted].

¹⁰³ Responses to the CMA's information requests [redacted]; [redacted] response to Ofcom's information request [redacted].

¹⁰⁴ [redacted] response to CMA's information request [redacted].

¹⁰⁵ [redacted] response to CMA's information request [redacted].

¹⁰⁶ In appendix B: parameters of competition, we provide more details on this evidence. In the appendix, we also present our assessment of parameters negotiated between cloud providers and large customers.

¹⁰⁷ Responses to Ofcom's information requests [redacted]; Responses to the CMA's information requests to [redacted].

through discounts,¹⁰⁸ including discounts under committed spend agreements,¹⁰⁹ cloud credits,¹¹⁰ and free tiers.¹¹¹

- (b) Aside from price, providers also identified a number of additional parameters of competition in the supply of cloud services. These included ease of migration,¹¹² security and data protection,¹¹³ innovation,¹¹⁴ availability of advanced cloud features,¹¹⁵ reliability¹¹⁶ of a platform,¹¹⁷ flexibility to deploy cloud services in combination with their traditional IT infrastructure,¹¹⁸ elasticity of its services,¹¹⁹ customer's loyalty,¹²⁰ brand trust,¹²¹ and technical support.¹²²

- 2.55 Microsoft and IBM said that different customer groups prioritise different factors when choosing a cloud provider.¹²³ According to these providers, these customer groups vary based on the industry they operate in, whether they are in a regulated sector or whether they are public sector customers.¹²⁴
- 2.56 Microsoft said that the increasingly heterogeneous nature of customers' needs has 'opened opportunities for existing and new cloud providers to differentiate themselves in different industry and workload verticals, without the need for hyperscaler scale'.¹²⁵
- 2.57 In response to our PDR, Microsoft and AWS commented on the relevance of AI as part of a cloud provider's offering. Microsoft submitted that rolling out of AI capabilities is a key parameter of competition between cloud providers.¹²⁶ AWS said that AI is driving a lot of competition, and has the potential to dramatically change cloud providers' market positions over the next one to five years.¹²⁷ We consider the impact of AI on cloud services in more detail elsewhere in this report.

¹⁰⁸ [redacted] response to CMA's information request [redacted].

¹⁰⁹ See chapter 8.

¹¹⁰ Responses to Ofcom's information requests [redacted]; [redacted] response to CMA's information request [redacted].

¹¹¹ Responses to Ofcom's information requests [redacted].

¹¹² Responses to Ofcom's information requests [redacted]. Responses to the CMA's information requests to [redacted].

¹¹³ Responses to the CMA's information requests [redacted]; [redacted] response to Ofcom's information request [redacted].

¹¹⁴ [redacted] response to Ofcom's information request [redacted]; [redacted] response to CMA's information request [redacted].

¹¹⁵ [redacted] response to CMA's information request [redacted].

¹¹⁶ Reliability is often captured by the uptime: a measure of the amount of time that a system or service is available and operational without any planned downtime.

¹¹⁷ Responses to the CMA's information requests [redacted]; [redacted] response to Ofcom's information request [redacted].

¹¹⁸ Responses to the CMA's information requests to [redacted].

¹¹⁹ [redacted] response to CMA's information request [redacted].

¹²⁰ [redacted] response to Ofcom's information request [redacted].

¹²¹ [redacted] response to CMA's information request [redacted].

¹²² Responses to the CMA's information requests [redacted].

¹²³ Responses to the CMA's information requests [redacted].

¹²⁴ Responses to the CMA's information requests [redacted].

¹²⁵ Microsoft's response to CMA's information request [redacted].

¹²⁶ [Microsoft's response to the CMA's PDR dated 28 January 2025](#), paragraph 79.

¹²⁷ [Summary of hearing with AWS, 12 March 2025](#), paragraph 10.

Evidence from customers

- 2.58 We initially asked large customers to rate the importance of a list of factors their organisation considers when choosing their main public cloud provider.
- (a) The following factors were identified as the most important by the large customers we heard from when choosing their main public cloud provider: service quality, price including discounts or cloud credits, data sovereignty requirements, range of cloud infrastructure services, number and location of data centres.¹²⁸
 - (b) These factors were followed in importance by other factors: cost and ability to use software licences, cloud-specific skills of employees, existing relationship with the cloud provider, ease of integration with existing technology.¹²⁹
 - (c) The following factors were identified as being generally of less importance, although they all were reported as of high importance by some customers: range of cloud infrastructure services offered by ISVs, ease of integration with other public clouds, AI capabilities.¹³⁰

Switching and multi-cloud

- 2.59 Switching between cloud providers involves a customer moving one or more workloads, or parts of workloads, from one provider's cloud to another. If a customer moves all its workloads from one cloud provider to another, we call this full switching. If it does so only for some of its workloads, we refer to it as partial switching.
- 2.60 A multi-cloud approach involves the placement by a customer of at least one workload on one provider's cloud and at least one workload on another provider's cloud.¹³¹
- 2.61 Both of these concepts are important as they can significantly influence the competitive dynamics in the supply of cloud services.
- 2.62 Customers' ability to switch or use multiple providers is key to the competitive process. In chapter 6 (barriers to switching and multi-cloud), we discuss the role of switching and multi-cloud in cloud markets. In the section on market outcomes in chapter 3 (competitive landscape), we assess the evidence on how common switching and multi-cloud is among customers. Here, we present:

¹²⁸ Responses to the CMA's information requests [X].

¹²⁹ Responses to the CMA's information requests [X].

¹³⁰ Responses to the CMA's information requests [X].

¹³¹ It is more rare to split a single workload across different providers.

- (a) how the switching process works for customers and the factors that customers account for when deciding to switch;
- (b) the different types of multi-cloud architecture that customers adopt; and
- (c) the benefits and disadvantages to customers of using multi-cloud.

Switching process

- 2.63 AWS, Google, IBM and Microsoft said customers must typically take a series of steps to switch between public clouds. The steps listed by these providers varied but included: assessing their existing environment, developing a migration/switching strategy, testing the integrity of their current IT environment, deploying workloads in the target public cloud, optimising the usage of the target cloud and validating their target cloud environment.¹³²
- 2.64 Google and AWS said the process of switching remains the same regardless of whether a customer is switching to or from another cloud provider, to or from another deployment model (eg, private cloud to public cloud), or to or from traditional on-premises IT.¹³³
- 2.65 Cloud providers said that there are certain financial and time costs associated with moving workloads between providers. These costs varied across submissions and included: customers' existing cloud set-up, licensing restrictions, complexity and size of the workloads being migrated, customers' priorities and chosen deployment method.¹³⁴
- 2.66 Evidence we have seen from customers we contacted indicates that, in addition to what providers told us, two other factors may influence the extent to which customers move workloads between different cloud providers:
- (a) First, customers may not move workloads if they consider that there is insufficient differentiation between providers¹³⁵ and in some cases customers have said that the effort involved outweighs the benefits.¹³⁶
 - (b) Second, customers may not move workloads if they consider a switch needs to occur at the right moment within the application lifecycle. In particular, some customers' responses show that they will not consider switching until their workload has been running on the cloud for a certain period of time.¹³⁷

¹³² Responses to the CMA's information requests [REDACTED].

¹³³ Responses to the CMA's information requests [REDACTED].

¹³⁴ Responses to the CMA's information requests [REDACTED].

¹³⁵ Responses to the CMA's information requests [REDACTED]; Notes of meetings with [REDACTED].

¹³⁶ Responses to the CMA's information requests [REDACTED]; Note of meeting with [REDACTED].

¹³⁷ [REDACTED] response to Ofcom's information request [REDACTED]; [REDACTED] response to the CMA's information request [REDACTED].

- 2.67 Some customers said that staff skillsets and expertise influenced their decisions to switch cloud providers. The evidence shows that where customers' in-house teams specialise in a certain cloud, this contributed to a reluctance to switch clouds as each provider had its own tools, frameworks, methodologies and best practices that required specific knowledge and training.¹³⁸
- 2.68 Customers who switched tended to do so for performance reasons, to improve cost-efficiency, or to reduce the complexity of their cloud architectures.¹³⁹ The Jigsaw report found that, in some cases, switching was a requirement for a customer following a merger or acquisition to rationalise providers.¹⁴⁰
- 2.69 Some customers pointed to the interaction between multi-cloud and switching. For example, one customer we contacted that had switched said that they had to operate two different public clouds and migrate systems between them, including maintaining periods of dual running.¹⁴¹ The Jigsaw report also found that the process of switching typically involved running two sets of services concurrently, then switching off the legacy system when ready.¹⁴²

Types of multi-cloud

- 2.70 In this section, we set out evidence from providers and customers on the different forms of multi-cloud use and degrees of integration that customers have between clouds.
- 2.71 AWS, Microsoft, Google and IBM identified three broad models of multi-cloud that customers may choose to adopt, as well as a set of benefits and disadvantages specific to each of them.¹⁴³
- (a) Cloud duplication, where customers mirror the entirety or part of their IT architecture on several clouds, enabling their workloads and applications to run equivalently on all of them.¹⁴⁴ Google said that this form of multi-clouding is relatively uncommon and is preferred by customers who require resilience over specific components of their cloud architecture.¹⁴⁵
 - (b) Integrated multi-cloud, where customers can mix and match cloud services from different public cloud providers¹⁴⁶ and there is a degree of integration between these services (eg, data is stored on one public cloud but analysed

¹³⁸ Responses to the CMA's information requests [REDACTED].

¹³⁹ Responses to the CMA's information requests [REDACTED]; Note of meeting with [REDACTED]; [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), paragraph 3.6.7.

¹⁴⁰ [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), paragraph 3.6.7.

¹⁴¹ Note of meeting with [REDACTED].

¹⁴² [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), paragraph 3.6.8.

¹⁴³ Responses to the CMA's information requests [REDACTED].

¹⁴⁴ Responses to the CMA's information requests [REDACTED].

¹⁴⁵ Google's response to CMA's information request [REDACTED].

¹⁴⁶ [REDACTED] response to the CMA's information request [REDACTED].

on a different one). Application components run on different clouds.¹⁴⁷

Google said that there is 'real customer appetite for integrated multi-cloud strategies', but certain practices, for example licensing restrictions, 'restrict or hinder many customers from adopting a truly integrated multi-cloud approach'.¹⁴⁸

- (c) Siloed multi-cloud, where customers use different cloud providers for different workloads with no or minimal integration between the different clouds.¹⁴⁹

Google said that this type of multi-clouding is the most prevalent amongst customers that have a multi-cloud architecture, as many of them are still at a relatively early stage of their cloud journey.¹⁵⁰

2.72 We have found this categorisation of the types of multi-cloud to be useful but too simple because the boundaries between types of multi-cloud, especially siloed and integrated, are blurred. There is a spectrum of how integrated these blurred architectures can be: they can be anywhere from partially integrated (eg a second cloud is used only for a specific service or application) which is closer to siloed multi-cloud, to highly integrated (eg integrating multiple applications and data hosted on different clouds).

2.73 There is a variety of types of multi-cloud being deployed by customers, each driven by specific use cases:

- (a) some customers told us that they integrate multiple public clouds for the purpose of simplifying their management and operation, by integrating their ancillary services, such as Identity Access Management (IAM).¹⁵¹
- (b) the Jigsaw report found that one of the most mentioned examples of integrated multi-cloud involved using Microsoft's IAM with services from another cloud provider.¹⁵²
- (c) this is also consistent with market research from Gartner which states that customers who wish to integrate multiple public clouds should aim to standardise policies and procedures and set up common tools for cost optimisation, security, IAM, monitoring and observability.¹⁵³

2.74 We consider the technical costs of using multiple cloud architectures in greater depth in the section on technical barriers.¹⁵⁴

¹⁴⁷ [redacted] response to the CMA's information request [redacted].

¹⁴⁸ Google's response to CMA's information request [redacted].

¹⁴⁹ [redacted] response to the CMA's information request [redacted].

¹⁵⁰ Google's response to CMA's information request [redacted].

¹⁵¹ For example, for integration of IAM services between public clouds. Responses to the CMA's information requests [redacted]; Note of meeting with [redacted]; [redacted] response to Ofcom's information request [redacted].

¹⁵² [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), paragraph 3.5.5(b).

¹⁵³ [redacted] response to the CMA's information request [redacted].

¹⁵⁴ See chapter 6.

- 2.75 Some customers integrate or network between applications on different public clouds, for example by building intermediary integration layers that connect cloud networks. We understand that these integrations tend to be between the storage services of multiple clouds, rather than directly between features of other services.¹⁵⁵
- 2.76 We also heard that some customers are, or have experimented with, integrating within applications and/or workloads across multiple public clouds. This includes using services with cross-cloud elements (such as querying the storage service on one cloud from a data warehouse service on another cloud), in addition to connecting only between the storage services.¹⁵⁶
- 2.77 Microsoft, Google and IBM told us that certain customer groups are more likely to adopt a multi-cloud strategy (or particular types of multi-cloud strategies) than others, namely:¹⁵⁷
- (a) digital native customers are better positioned to adopt multi-cloud strategies;¹⁵⁸
 - (b) Microsoft told us that larger enterprises are more likely to deploy cloud agnostic multi-cloud than smaller customers, who are more likely to adopt the multi-cloud org model (where different applications live in different clouds such that the organisation adopts multi-cloud but the applications do not).¹⁵⁹ Google said that large companies tend to have distinct departments with their own independent workloads, facilitating the adoption of a siloed multi-cloud approach;¹⁶⁰
 - (c) enterprises with complex regulatory compliance obligations and/or a requirement for enhanced resilience and stability (eg financial institutions) often opt for multi-cloud strategies;¹⁶¹ and
 - (d) merged entities may temporarily adopt a multi-cloud approach to combine each company's distinct cloud strategy. However, these enterprises typically choose one primary provider after the merger.¹⁶²

¹⁵⁵ Responses to the CMA's information requests [REDACTED]; Notes of meetings with [REDACTED].

¹⁵⁶ Due to differentiation in interpretation of the term 'workload', we have grouped responses that related to integrating within applications and workloads, as we understand the difference between the two approaches to be of no consequence for our analysis. Responses to the CMA's information requests [REDACTED]; Notes of meetings with [REDACTED].

¹⁵⁷ Responses to the CMA's information requests [REDACTED].

¹⁵⁸ [REDACTED]. [REDACTED] response to CMA's information request [REDACTED].

¹⁵⁹ Microsoft's response to CMA's information request [REDACTED].

¹⁶⁰ Google's response to CMA's information request [REDACTED].

¹⁶¹ Responses to the CMA's information requests [REDACTED].

¹⁶² [REDACTED] response to CMA's information request [REDACTED].

Benefits and disadvantages of using multi-cloud for customers

2.78 Using a multi-cloud architecture brings certain benefits to customers, but it also has some disadvantages. These vary depending on the type of multi-cloud architecture chosen and include:

- (a) Competitive tension: multi-clouding allows customers to leverage the threat of moving individual workloads to competing providers to increase their bargaining power with their cloud providers.¹⁶³
- (b) Resilience: multi-clouding minimises the risk of downtime – cloud duplication gives the highest degree of resiliency and downtime protection.¹⁶⁴ However, integrated multi-cloud also creates multiple points of failure across different clouds, increasing the risks customers are exposed to, including security ones.¹⁶⁵
- (c) Compliance: multi-clouding allows customers to comply with specific legal or regulatory requirements;¹⁶⁶
- (d) Geographic reach: multi-clouding can help customers achieve global coverage.¹⁶⁷
- (e) Flexibility: a more integrated approach gives customers the flexibility to use their preferred services from different public cloud providers.¹⁶⁸

2.79 The disadvantages include:

- (a) Higher costs: because of the greater complexity and increased staffing requirements, customers may incur higher costs when using multiple clouds.¹⁶⁹ However, a more integrated multi-cloud strategy could also minimise customers' costs, as they are able to select the most cost-effective solution for each individual workload.¹⁷⁰
- (b) Complexity: multi-clouding can increase management complexity.¹⁷¹ Relatedly, integrating workloads across multiple clouds can lead to challenges around data latency and data governance.¹⁷²

¹⁶³ [redacted] response to CMA's information request [redacted].

¹⁶⁴ [redacted] response to CMA's information request [redacted].

¹⁶⁵ Responses to the CMA's information requests [redacted].

¹⁶⁶ [redacted] response to CMA's information request [redacted].

¹⁶⁷ [redacted] response to CMA's information request [redacted].

¹⁶⁸ Responses to the CMA's information requests [redacted].

¹⁶⁹ [redacted] response to CMA's information request [redacted].

¹⁷⁰ [redacted] response to CMA's information request [redacted].

¹⁷¹ Responses to the CMA's information requests to [redacted].

¹⁷² [redacted] response to CMA's information request [redacted].

- 2.80 Customer views vary on this, reflecting their current level of integration, use cases or the appetite to use more integrated types of multi-cloud:
- (a) In relation to highly integrated multi-cloud, in a handful of cases customers said that they currently have, at most, a limited use case.¹⁷³ Reasons for this included that there are challenges with adopting this approach.¹⁷⁴ However, other responses from these customers suggest that they do have some form of integration.¹⁷⁵
 - (b) The responses from a few customers suggest that they may be more open to adopting an integrated multi-cloud approach in the future.¹⁷⁶ One customer said that while it currently has ‘no business reason to have full integration’, it is feasible that they may want to integrate in the future.¹⁷⁷ Another customer noted that, while its legacy applications migrated from on-premises are not suitable for integrated multi-cloud, the same is not necessarily true for its new cloud-native workloads.¹⁷⁸
 - (c) Customers had somewhat differing views on the disadvantages of multi-cloud. For example, one customer said that when utilising a multi-cloud strategy, consideration must be applied to risks associated with a business process traversing multiple cloud providers, to prevent the impact of a single provider failure from introducing increased operational resilience risk to the business process.¹⁷⁹ However, another customer said that it ran two related applications on two different clouds and this did not negatively affect the resiliency of either application.¹⁸⁰

Conclusions

- 2.81 Cloud services are a vital input to businesses and organisations across the UK economy, with £10.5 billion spent on them in 2024. Cloud services revenues in the UK have grown at an annual average rate of 29% over 2020-2024. That growth is expected to continue, although there is some indication that it could be at a slightly slower rate.
- 2.82 Cloud services are used by a large number of customers of all sizes and across most industries in the UK. A small number of high-spend customers are responsible for a significant proportion of cloud providers’ UK revenues.

¹⁷³ Responses to the CMA’s information requests to [redacted]; Note of meeting with [redacted].

¹⁷⁴ [redacted] response to the CMA’s information request [redacted]; Note of meeting with [redacted].

¹⁷⁵ Responses to the CMA’s information requests to [redacted]; Note of meeting with [redacted].

¹⁷⁶ Responses to the CMA’s information requests to [redacted].

¹⁷⁷ [redacted] response to the CMA’s information request [redacted].

¹⁷⁸ Note of meeting with [redacted].

¹⁷⁹ Note of meeting with [redacted].

¹⁸⁰ [redacted] response to the CMA’s information request [redacted].

- 2.83 Microsoft and AWS are considered the leading providers of cloud services by customers and other providers, followed by Google. Smaller providers are generally not seen as an effective alternative to large customers' main providers (generally AWS and Microsoft), but we note that they may be seen as suitable alternatives for certain workloads.
- 2.84 Customers buy cloud services through a variety of channels depending on their characteristics and needs. They consider a range of factors when purchasing cloud services, including service quality, price including discounts or cloud credits, data sovereignty requirements, range of cloud infrastructure services, number and location of data centres.
- 2.85 Switching and multi-cloud require customers to take a series of steps and incur costs which may affect their ability and/or incentive to do so. Decisions on whether to switch or multi-cloud, as well as the level of any multi-cloud integration to adopt, are driven by how customers weigh the different benefits and disadvantages. These depend on the customers' circumstances, needs, preferences and use cases.
- 2.86 We consider prevalence of switching and multi-cloud amongst customers in chapter 3 and the barriers to switching and multi-cloud in chapter 6.

3. Competitive landscape

- We have found a product market for IaaS based on standard compute infrastructure which is separate from IaaS based on accelerated compute. For the purposes of this report, when not otherwise specified, we use the term IaaS to refer to IaaS based on standard compute. We have considered PaaS in aggregate as a product market for the purposes of this market investigation. The geographic scope of both the IaaS and PaaS markets is the UK and EEA.
- Neither SaaS nor traditional 'on-premises' IT and private cloud form part of the same markets as IaaS and PaaS, because most customers do not see them as substitutes for those cloud services.
- The IaaS market is highly concentrated and the two largest cloud providers, Microsoft and AWS, each has a high share of supply: in 2024, they each had [30-40]% revenue shares in the UK+EEA region. The PaaS market is also concentrated but AWS and Microsoft have smaller shares of supply at [10-20]% and [20-30]% respectively and there is a larger number of small PaaS providers. Across IaaS and PaaS markets together in 2024, AWS and Microsoft's respective shares were [20-30]% and [30-40]% in the UK+EEA region.
- The third largest cloud provider, Google, has much lower shares of supply, at [5-10]% in IaaS and PaaS markets in 2024. There are also a few smaller cloud providers, including Oracle and IBM, whose share of supply is even smaller and which do not supply as wide a range of cloud services as Microsoft, AWS and Google.
- We have found that Microsoft and AWS have been generating sustained returns from their cloud services substantially above their cost of capital in cloud services for several years.
- While customers recognise that cloud services offer many benefits to them including in relation to quality and innovation, we consider that a more competitive market would have sustained better market outcomes. These could include more consistently competitive prices, as well as improvements in quality and innovation.
- Switching between cloud services providers is extremely rare. This is consistent with the existence of high barriers to switching and with many customers considering that switching costs outweigh the benefits of changing provider.
- The use of multi-cloud architectures is more common than switching in cloud services, particularly amongst larger customers. However, most smaller and medium-sized customers do not multi-cloud. The overall prevalence of multi-cloud, in conjunction with customer evidence suggesting that the costs of multi-cloud generally outweigh the benefits, indicates that some barriers exist and that, in particular, customers' ability to integrate workloads on more than one cloud is subject to barriers. Even where

customers have opted for a multi-cloud strategy, they generally concentrate a high proportion of their workloads with a 'primary' cloud provider.

- 3.1 This chapter provides an overview of the competitive landscape for cloud services. It includes our assessment of market definition, market structure and concentration, and market outcomes.

Market definition

- 3.2 Our guidelines for market investigations note that we look at the main characteristics of the market and the outcomes of the competitive process, along with the composition of the relevant market within which competition may be harmed.¹⁸¹ In this chapter, we set out our analysis of market definition, before considering market structure and outcomes.
- 3.3 Our guidelines state that defining the market enables the CMA to focus on the sources of any market power and provides a framework for its assessment of the effects on competition of features of a market.¹⁸²
- 3.4 Further, our guidelines state that market definition is a useful tool, but not an end in itself, and that identifying the relevant market involves an element of judgement. The boundaries of the market do not determine the outcome of our competitive assessment in any mechanistic way. The competitive assessment takes into account any relevant constraints from outside the market, segmentation within it, or other ways in which some constraints are more important than others.¹⁸³
- 3.5 There are normally two dimensions to the definition of a market:¹⁸⁴
- (a) a product dimension where the relevant product market comprises a set of substitute products; and
 - (b) a geographic dimension, where the relevant geographic market may be national (or wider), regional or local.
- 3.6 We may treat a group of markets together for the purposes of assessing competitive effects where a feature manifests itself in a similar way across several different markets and the CMA is able to reach a view about the effects of the feature on competition across the group of markets as a whole.¹⁸⁵

¹⁸¹ See [CC3 \(Revised\)](#), [Guidelines for market investigations: Their role, procedures, assessment and remedies](#) (CC3(Revised)), paragraph 94.

¹⁸² [CC3 \(Revised\)](#), paragraph 132.

¹⁸³ [CC3 \(Revised\)](#), paragraph 133.

¹⁸⁴ [CC3 \(Revised\)](#), paragraph 142 also identifies customer groups or temporal factors.

¹⁸⁵ [CC3 \(Revised\)](#), paragraph 152.

- 3.7 The hypothetical monopolist test (HMT) is a tool which can be used to identify effective substitutes and to check that the market is not defined too narrowly. The principle behind it rests on defining a market as a product, or collection of products, a sole supplier of which could hypothetically impose a small but significant non-transitory increase in price (referred to as the SSNIP test). The test can help to identify the constraints that would prevent a hypothetical monopolist from exercising market power. In practice it may often be used as a conceptual framework rather than quantitatively.¹⁸⁶
- 3.8 The market definition(s) used by the CMA need not always correspond with the market for the goods or services described in the Terms of Reference (the 'relevant market(s)').¹⁸⁷ The CMA may conclude that the market definition goes wider or narrower than those goods and services.¹⁸⁸
- 3.9 We consider the product and geographic market definition below.

Product market

- 3.10 The willingness and ability of customers to switch to other products is a driving force of competition. This means that, in forming views on market definition, we considered the degree of demand-side substitutability. We have, where relevant, included supply-side factors in defining the market. This may arise, for example, if firms supplying non-substitute products have the capabilities and assets to redirect production to goods and services that would be substitutes for those in the market.¹⁸⁹ In determining whether there is supply-side substitutability, we considered factors such as whether:¹⁹⁰
- (a) suppliers supply a range of different products in the same broad category, using the same set of assets and capabilities; or
 - (b) suppliers regularly introduce new products or reposition existing ones within the category.
- 3.11 The Terms of Reference (TOR) for this market investigation in relation to public cloud infrastructure services in the UK defined 'cloud infrastructure services' as comprising IaaS and PaaS.¹⁹¹ As noted above, the CMA may conclude that the market definition goes wider or narrower than those goods and services set out in the TOR. Therefore, we assessed whether demand-side substitutability and

¹⁸⁶ CC3 (Revised), paragraph 138.

¹⁸⁷ CC3 (Revised), paragraph 26.

¹⁸⁸ CC3 (Revised), paragraph 131.

¹⁸⁹ CC3 (Revised), paragraphs 134.

¹⁹⁰ CC3 (Revised), paragraph 134 and footnote 75.

¹⁹¹ Terms of Reference ([ofcom.org.uk](https://www.ofcom.org.uk)), page 2.

supply-side substitutability implied a narrower or wider market(s). For the purposes of our assessment, we considered whether:

- (a) IaaS should be segmented into separate, narrower markets by types of service;
- (b) IaaS based on standard compute is in a separate relevant market from IaaS based on accelerated compute;¹⁹²
- (c) there is a single market for IaaS and PaaS;
- (d) PaaS should be segmented into separate, narrower markets;
- (e) there is a single market for PaaS and SaaS; and
- (f) markets for IaaS and/or PaaS should be widened to include alternative IT models.

IaaS

Whether IaaS should be segmented into separate, narrower markets by types of service

- 3.12 IaaS includes three broad types of services that provide access to raw computing resources for processing workloads and storing data: compute, storage and networking.¹⁹³ This section considers whether IaaS should be segmented into separate, narrower markets, ie one for each of these types of services. We discuss IaaS based on accelerated compute separately.
- 3.13 We set out below the cloud providers' submissions in relation to market definition for IaaS before presenting our assessment. We did not receive customer views on the segmentation of IaaS.

Cloud provider views

- 3.14 Most cloud providers' relevant submissions related more generally to the similarities and differences between IaaS, PaaS and SaaS. We have set these out here as we consider these submissions relevant to the potential sub-segmentation between the IaaS elements.

¹⁹² Standard compute is that based on CPUs, whereas accelerated compute also relies on accelerator chips such as GPUs.

¹⁹³ In this context, by compute we mean services that provide customers access to virtual machines, but we note that compute can be used in other contexts to refer to other services which we consider to be PaaS.

- 3.15 Cloud providers submitted that the market is wider than IaaS, and therefore IaaS should not be subdivided:
- (a) Microsoft, AWS and Google said that there is not a single clear-cut IaaS/PaaS/SaaS categorisation, and that different people might disagree over how a particular solution should be classified.¹⁹⁴ These cloud providers said that they do not operate their businesses according to a strict IaaS/PaaS/SaaS segmentation.¹⁹⁵ AWS further submitted that customers do not tend to think about whether they should use 'IaaS' or 'PaaS' to solve their IT needs, and the solution for a particular IT need may involve one or more different services, such as compute, storage, and networking, working together in a specific way.¹⁹⁶
 - (b) AWS said that customers start by defining their objectives and needs and then look at a broad set of options that might help meet them, and do not typically evaluate products or services based on IaaS, PaaS or SaaS categories.¹⁹⁷ AWS said that all models can be used to deploy similar solutions and are substitutable with each other, and that comparisons between categories are insufficient to reflect customer behaviour and the competition facing any given product.¹⁹⁸ AWS submitted that customers care about solving for an IT need — whether on-premises, in the cloud or some combination thereof. Equally, they do not focus on industry labels related to cloud services, such as IaaS, PaaS and SaaS. Indeed, they often mix and match different IT services to fulfil the same need, including combinations of IaaS, PaaS and SaaS or one or more different services, such as compute, storage and networking, working together in a specific way.¹⁹⁹
 - (c) Google said the level of control that an enterprise has over its workloads can be scaled up and down far more flexibly than the IaaS/PaaS/SaaS segmentation suggests.²⁰⁰
 - (d) Oracle said that IaaS and PaaS should not be distinguished for the purposes of market definition.²⁰¹

¹⁹⁴ Responses to the CMA's information requests [🔗].

¹⁹⁵ Responses to the CMA's information requests [🔗].

¹⁹⁶ [AWS' response to the CMA's PDR dated 28 January 2025, paragraph 10.](#)

¹⁹⁷ AWS' response to the CMA's information request [🔗].

¹⁹⁸ AWS' response to the CMA's information request [🔗].

¹⁹⁹ [AWS response to the CMA's working papers and updated issues statement, 25 June 2024](#), paragraph 7; [AWS' response to the CMA's PDR dated 28 January 2025, paragraph 10](#); AWS' response to the CMA's information request [🔗].

²⁰⁰ Google's response to the CMA's information request [🔗].

²⁰¹ Oracle's response to the CMA's information request [🔗].

(e) Microsoft and AWS said suppliers are innovative and a supplier present at a particular layer may be a competitive constraint on other infrastructure layers.²⁰²

3.16 In response to our PDR, Cloudflare submitted that Network-as-a-Service (NaaS) should be included in the IaaS category. It explained that NaaS allows customers to operate their own networks without maintaining their own network infrastructure.²⁰³

Our assessment

3.17 In relation to the submissions above, whilst there is no widely accepted or clear distinction between IaaS, PaaS and SaaS, a degree of substitutability between the three does not mean that services at different ends of the spectrum are substitutable on the demand or supply side. We also note that even if customers do not delineate products by labels this does not necessarily mean that these activities should not be segmented into separate markets as different types of solutions may not be substitutable.

3.18 We consider that the different elements of IaaS, namely compute, storage and networking, serve different functions. We consider that these services, as defined, accommodate for 'borderline' services that do not necessarily fit neatly into any specific element of IaaS (eg NaaS – we consider networking a part of IaaS but acknowledge that some cloud providers offer solutions like NaaS which we consider closer to a PaaS offering). Each of the elements of IaaS are not demand-side substitutes.

3.19 On the supply side, we understand that all IaaS providers supply compute, storage and networking. We did not receive any evidence that suggests that firms compete to supply these elements under different conditions of competition.²⁰⁴ Accordingly, we consider there is a significant degree of supply-side substitution between each of these IaaS elements.

3.20 We have considered the extent to which there is variation in the competitive strength of cloud providers in the supply of standard compute, storage and networking in our wider assessment. However, strengths in any of these services has not been identified as a key parameter of competition either by cloud providers or customers.²⁰⁵ This suggests that variation in strengths in these areas does not significantly impact competition.

²⁰² AWS' response to the CMA's information request [§]; Microsoft's response to the CMA's information request [§].

²⁰³ [Cloudflare's response to the CMA's PDR](#) dated 28 January 2025, page 1-2.

²⁰⁴ CC3 (Revised), paragraph 134.

²⁰⁵ See appendix B.

- 3.21 The features that we have investigated are relevant across each of the IaaS services, and as such it is reasonable to consider IaaS in aggregate for the purposes of this market investigation.²⁰⁶
- 3.22 While compute, storage and network products are not demand-side substitutes, the same firms compete to supply these products under similar conditions of competition. In light also of the scope for the different features of IaaS to be considered in aggregate, we aggregate the supply of compute, storage and networking and analyse them as one market, reflecting that this does not affect our competitive assessment.

Whether IaaS based on standard compute is in a separate relevant market from IaaS based on accelerated compute

- 3.23 IaaS based on accelerated compute refers to infrastructure services that involve the use of powerful accelerator chips, and often high bandwidth networking, to provide accelerated computing services. Although IaaS based on accelerated compute has specific accelerate compute components, it may also include standard compute (eg CPUs), networking and storage elements that are required for it to operate.²⁰⁷
- 3.24 We examined demand for IaaS based on accelerated compute infrastructure in the context of the growing demand for IaaS services suitable for use in developing and deploying artificial intelligence (AI) models.²⁰⁸
- 3.25 As context to our assessment, accelerated compute infrastructure involves specialised hardware and accompanying software. The design of accelerator chips makes them more efficient to use for certain workloads than Central Processing Units (CPUs).²⁰⁹ Multiple accelerator chips can be networked together with high bandwidth cables and used in parallel to increase the efficiency of computations, whereas CPUs do not scale well for this type of operation.²¹⁰
- 3.26 There are two types of accelerator chips commonly used in accelerated compute infrastructure: GPUs and ASICs.
- (a) Graphical Processing Units (GPUs), which have many small, specialised processing cores that run in parallel to perform computations simultaneously. This makes them well-suited to computing the large number of mathematical

²⁰⁶ CC3 (Revised), paragraph 152.

²⁰⁷ Although cloud providers often specify whether an IaaS service is specifically compute, storage, or networking, all IaaS services will require some mix of standard compute, storage, and networking elements to function (in addition to other physical and digital assets necessary to deliver the service).

²⁰⁸ See eg [AI data center growth: Meeting the demand | McKinsey](#), accessed on 22 April 2025.

²⁰⁹ CMA's initial review of FMs, [Full report](#), page 13. The Central Processing Unit (CPU) is the primary component of a computer that acts as its 'control centre.' The CPU interprets, processes and executes instructions, most often from the hardware and software programs running on the device. [Glossary Central Processing Unit \(CPU\)](#).

²¹⁰ CMA's initial review of FMs, [Full report](#), page 13.

operations necessary in FM training and inference.²¹¹ The most popular GPUs in the market are Nvidia's.²¹²

- (b) Application-Specific Integrated Circuits (ASICs), which are hardwired for a specific application. ASICs developed for AI are typically hyper-specialised GPUs, with modifications that increase the efficiency of specific AI workloads. The AI accelerator chips developed by Google (TPUs), Amazon (Trainium and Inferentia) and Microsoft (Maia 100) are all ASICs.²¹³

- 3.27 Certain large-scale AI workloads, such as training a foundation model (FM) or large language model (LLM), involve large amounts of data and a high number of mathematical operations. Accelerated compute infrastructure is well suited to efficiently meet these requirements. Other types of workloads with similar requirements also use IaaS based on accelerated compute infrastructure, eg high-performance computing applications such as pharmaceutical discovery, weather forecasting and financial modelling.²¹⁴
- 3.28 We set out below the cloud providers' submissions, in relation to whether IaaS based on standard compute is part of a separate market from IaaS based on accelerated compute, before presenting our assessment. The evidence we have assessed relates to the requirements of AI workloads and AI developers, as it is AI development and deployment that has been driving demand for IaaS based on accelerated compute.

Cloud provider views

- 3.29 AWS submitted that for most 'AI and ML-related' workloads, factors typically impacting customer compute requirements include the size of the model, the type of training or inferencing to be performed, preferences and constraints on how quickly the workload should be performed, and cost, which can exist on a sliding scale.²¹⁵
- 3.30 However, AWS submitted that finding IaaS based on accelerated compute was distinct from IaaS based on standard compute was 'unduly narrow' because it fails to account for the fact that the deployment of FMs has created significant

²¹¹ Inference refers to the process of calling upon a FM, for example when using an FM-enabled product or service to generate content or perform analysis. FM deployment refers to the embedding of FMs in products and services to be used in this way.

²¹² [Analyzing NVIDIA's growth strategy: How the semiconductor leader is powering generative AI and the future of computing - CB Insights Research](#), accessed on 4 November 2024, [Nvidia: The chip maker that became an AI superpower - BBC News](#), accessed on 4 November 2024; [§] response to the CMA's information request [§].

²¹³ Responses to the CMA's information requests [§].

²¹⁴ [Cloud Compute Instances – Amazon EC2 Instance Types – AWS](#), accessed on 17 April 2025.

²¹⁵ AWS' response to the CMA's request for information [§].

opportunities for many types of IT providers to enter and expand which go beyond 'IaaS based on accelerated compute', notably because:²¹⁶

- (a) the meaning of accelerated compute is quickly expanding with new players and innovations;
- (b) there is overlap between the infrastructure required to supply IaaS based on accelerated compute and IaaS based on standard compute; and
- (c) there are different AI and machine learning use cases within each of accelerated and standard compute, particularly for inference and deployment of smaller or more efficient models.

3.31 AWS further submitted that hardware and AI models are rapidly being improved and sometimes replaced due to constant innovation, and it was nearly impossible to clearly delineate workload requirements for standard compute and accelerated compute because these exist on a 'multi-faceted and ever-changing spectrum of customer needs relative to hardware capabilities'.²¹⁷ AWS also submitted that [REDACTED] of IaaS based on accelerated compute is [REDACTED] such that the cost difference between IaaS based on accelerated and standard compute [REDACTED].²¹⁸

3.32 Regarding supply side substitutability, AWS said that many cloud providers are already active across IaaS based on standard compute and IaaS based on accelerated compute (including Microsoft, Google, IBM, Oracle, Alibaba, Tencent, CoreWeave, Vultr, Gensic Cloud, DigitalOcean, Scaleway, GCore, Civo, OVHcloud, and Acecloud) and that this is evidence of supply side substitutability.²¹⁹

3.33 Microsoft submitted that 'accelerated compute is generally more performant than standard compute, but less readily available and more expensive' and therefore customers may reallocate an accelerated compute workload to standard compute if it is possible without compromising performance or if performance is less important, or reallocate a standard compute workload to accelerated compute if it evolves to demand greater performance.²²⁰

3.34 Microsoft also submitted that it was wrong to segment IaaS based on accelerated and standard compute because:²²¹

- (a) Customers purchase compute and other cloud services based on workload requirements, mixing and matching resources to suit their workload –

²¹⁶ AWS' response to the CMA's PDR dated 28 January 2025, paragraph 16.

²¹⁷ AWS' response to the CMA's request for information [REDACTED].

²¹⁸ AWS' response to the CMA's request for information [REDACTED].

²¹⁹ AWS' response to the CMA's request for information [REDACTED].

²²⁰ Microsoft's response to the CMA's request for information [REDACTED].

²²¹ Microsoft's response to the CMA's PDR dated 28 January 2025, paragraph 17 and 75.

generally, customers do not buy 'AI compute' in isolation and instead integrate it with other PaaS and IaaS services;²²²

- (b) Demand side substitutability exists in both directions and customers often reallocate workloads between the two, depending on cost-efficiency, availability and performance considerations;²²³ and
- (c) Supply side substitutability exists due to cloud providers investments in common infrastructure.²²⁴

- 3.35 Regarding supply side substitutability, Microsoft submitted that common infrastructure investments include 'massive capex in data centres, networking, and storage, supporting both standard and accelerated compute' adding that 'cloud providers build scalable data centres with modular architectures, which means they can adjust GPU and CPU server capacity based on demand'.²²⁵ Microsoft said that [REDACTED].²²⁶
- 3.36 Microsoft further submitted that separating IaaS based on standard from accelerated compute 'does not align with the demand or supply of compute' because:²²⁷
- (a) Standard compute and accelerated compute are important for all workloads, including for narrow use cases of FM training and inferencing of large consumer models; and
 - (b) Customers deploy workloads that incorporate AI but require many different services that do not rely on accelerated compute.
- 3.37 Microsoft also submitted that AI and software development tools are optimised to facilitate the use of both CPUs and GPUs, and cloud providers offer solutions that allow workloads to switch between standard and accelerated compute, or add accelerated compute to standard compute instances, where necessary.²²⁸
- 3.38 Google said that, while at present standard compute chipsets and accelerated chipsets are sold and bought separately, which suggests there is separate demand for accelerated compute and standard compute, there is already evidence of convergence and it is too early to tell how the market will develop in the next

²²² Microsoft's response to the CMA's PDR dated 28 January 2025, paragraph 76 and 77(i).

²²³ Microsoft's response to the CMA's PDR dated 28 January 2025, paragraph 77(ii).

²²⁴ Microsoft's response to the CMA's PDR dated 28 January 2025, paragraph 76 and 77(iii).

²²⁵ Microsoft's response to the CMA's information request [REDACTED].

²²⁶ Microsoft's response to the CMA's information request [REDACTED].

²²⁷ Microsoft's submission to the CMA [REDACTED].

²²⁸ Microsoft's response to the CMA's information request [REDACTED].

few years.²²⁹ Separately, Google said that it is too early to know how accelerated and standard compute supply chains may evolve and interact going forward.²³⁰

- 3.39 Google submitted that there was ‘considerable uncertainty regarding the demand for and future pricing of accelerated compute’ and that ‘increased capabilities of some CPUs for AI inferencing may simultaneously increase supply of compute suitable for AI inferencing, which could further drive down prices’.²³¹
- 3.40 Oracle said that accelerated compute and standard compute are separate markets, but there is a grey area in the middle where the two are more substitutable with each other depending on customer needs.²³²

Our assessment

- 3.41 We considered the demand and supply-side substitutability between IaaS based on accelerated compute infrastructure and IaaS based on standard compute infrastructure and how they may change in the future.
- 3.42 First, we have considered the extent of demand-side substitution from IaaS based on standard compute to IaaS based on accelerated compute.
- 3.43 Accelerated compute infrastructure is technically differentiated from standard compute infrastructure: standard compute infrastructure relies on general-purpose CPU processors, whereas accelerated compute infrastructure additionally depends on accelerator chips that are typically more expensive. Accelerator chips can ‘accelerate’ the computational process leading to higher overall computing performance compared to using CPUs. Network cables are used for both IaaS based on accelerated and standard compute infrastructure. High bandwidth network equipment can be used for both standard and accelerated compute infrastructure but is more often worth the cost for many accelerated compute workloads.²³³
- 3.44 This technical differentiation leads IaaS based on standard compute and accelerated compute infrastructure to have largely different use cases. IaaS based on standard compute infrastructure is used for tasks such as running software, analysing data, managing network traffic and fetching data from memory.²³⁴ On the other hand, IaaS based on accelerated compute infrastructure is typically used

²²⁹ [Summary of hearing with Google, 17 March 2025](#), page 3, paragraph 20.

²³⁰ Google's submission to the CMA [§].

²³¹ Google's response to the CMA's information request [§].

²³² Note of meeting with Oracle [§].

²³³ AWS' response to the CMA's information request [§].

²³⁴ See AWS' description of its General Purpose Computing instances and their use cases ([Cloud Compute Instances – Amazon EC2 Instance Types – AWS](#), accessed on 16 May 2025) and Microsoft's description of Azure's General purpose virtual machines ([Virtual machine sizes overview - Azure Virtual Machines | Microsoft Learn](#), accessed on 16 May 2025).

for large scale AI workloads, carried out primarily by FM developers, and other high-performance computing tasks.²³⁵

- 3.45 Evidence indicated that accelerated compute infrastructure is significantly more expensive than standard compute infrastructure, but the former is more cost-effective, versus the latter, when used for high-performance computing tasks such as FM development and deployment.²³⁶ Cloud providers also offer larger discounts for IaaS based on accelerated compute infrastructure than they do for IaaS based on standard compute.²³⁷
- 3.46 We acknowledge submissions from cloud providers (presented above) showed that there is not a clear boundary between standard compute and accelerated compute use cases, particularly for the deployment of AI models. For workloads that customers could choose to deploy on either standard or accelerated compute, there are a range of factors that impact which type of compute may better meet their requirements, including cost and speed.²³⁸ We consider that demand for different types of compute being driven by different workload requirements is evidence of a lack of demand-side substitutability, rather than evidence of substitutability. We also consider that, whilst there is not a distinct boundary between standard compute and accelerated compute use cases – meaning they may be substitutable for some customers for some workloads – this does not mean the market should not be separated out if there is limited substitutability for most customers.
- 3.47 In relation to cloud providers' submissions on the substitutability of AI workloads, we note that cloud providers' IaaS revenue associated with AI almost exclusively derives from workloads using accelerated compute.²³⁹ This demonstrates that, although substitutability is technically possible, currently, most customers are choosing to deploy their AI workloads using IaaS based on accelerated compute and the extent of demand-side substitution to standard compute has been minimal.
- 3.48 There are a wide range of workloads that are currently deployed using IaaS based on standard compute.²⁴⁰ However, given the difference in price described above

²³⁵ See AWS' description of its Accelerated Computing instances and their use cases ([Cloud Compute Instances – Amazon EC2 Instance Types – AWS](#), accessed on 16 May 2025) and Microsoft's description of Azure's GPU accelerated virtual machines ([Virtual machine sizes overview - Azure Virtual Machines | Microsoft Learn](#), accessed on 16 May 2025).

²³⁶ Note of meeting with [REDACTED]; 'GPU vs. CPU Cost Analysis in Shared Hosting Environments' by Medium', accessed on 25 September 2024.

²³⁷ [REDACTED]. [REDACTED] response to the CMA's information request [REDACTED].

²³⁸ [REDACTED] response to the CMA's information request [REDACTED]; [REDACTED] response to the CMA's information request [REDACTED].

²³⁹ [REDACTED] submitted that accelerated compute made [REDACTED] of its 'AI infrastructure' revenue in FY24. [REDACTED] response to the CMA's information request [REDACTED]. [REDACTED] submitted revenue data for its generative AI focused IaaS offering, [REDACTED], which is [REDACTED]. [REDACTED] response to the CMA's information request [REDACTED]. [REDACTED] submitted that its 'AI IaaS' revenue was [REDACTED]. [REDACTED] response to the CMA's information request [REDACTED].

²⁴⁰ For example, AWS' website lists hundreds of IaaS based on standard compute offerings for a wide range of use cases including application servers, web servers, code repositories, ad serving, and analytics. [Cloud Compute Instances – Amazon EC2 Instance Types – AWS](#), accessed on 14 May 2025.

we would not expect most customers to substitute unless they require the specialised capabilities of accelerated compute for that workload.

- 3.49 Second, we have considered the extent of demand-side substitution from IaaS based on accelerated compute to standard compute.
- 3.50 We consider the technical differentiation between IaaS based on accelerated compute and IaaS based on standard compute described above to be relevant to the demand-side substitution from IaaS based on accelerated compute to standard compute too.
- 3.51 Evidence from a range of market participants demonstrated that there was limited demand-side substitution in this context, particularly in relation to the use of GPUs:
- (a) Some FM developers said that they predominantly rely on accelerator chips for the development of FMs.²⁴¹
 - (b) A chip designer said that most FMs are designed to work using GPUs.²⁴²
 - (c) While a cloud provider submitted that CPUs could be used for ‘AI/ML training, development and inference’²⁴³ and that [X] is specifically optimised to be deployable using ‘cost efficient infrastructure across [X]’, we note that this only refers to inference of a relatively small FM.²⁴⁴ This cloud provider also submitted that ‘most of [its] AI customers use GPUs’ and that to ‘compete hard for AI workloads’ it must be ‘continually improving upon [X]’ and ‘offering ample top-of-the-range GPUs’ – in this regard, there was no mention of CPUs.²⁴⁵
 - (d) Generally, internal documents received from one cloud provider discussed only investment in AI accelerator chips [X] when discussing its AI infrastructure strategy and investment.²⁴⁶
 - (e) Another cloud provider said that GPUs can be used for ‘demanding workloads requiring high performance computers’ and that ‘FM developers prefer GPUs for AI model training’. It added that GPUs can be used for other purposes (eg gaming infrastructure, video editing, fluid dynamics and non-large language model AI usage). This cloud provider submitted that, ‘in [its] view, the requirements of FM developers are substantially similar, if not identical, to high performance computing customers’.²⁴⁷

²⁴¹ Notes of meetings with [X].

²⁴² Note of meeting with [X].

²⁴³ [X] response to the CMA’s information request [X].

²⁴⁴ [X] response to the CMA’s information request [X].

²⁴⁵ [X] response to the CMA’s information request [X].

²⁴⁶ [X] response to the CMA’s information request [X].

²⁴⁷ [X] response to the CMA’s information request [X].

(f) Two other cloud providers said that accelerated compute is an essential input in the development and deployment of FMs:

(i) [X] submitted that ‘FM developers have specialised needs that differ from other public cloud customers’, the ‘most notable’ being ‘access to specialised chips, GPUs, that are in short supply’.²⁴⁸

(ii) [X] submitted that ‘both FM developers and FM providers need AI accelerators (GPUs) for their workloads. Depending on the scale of either the FM training (FM developer) or FM inferencing (FM provider), the required amount of GPU capacity may be significant’.²⁴⁹

3.52 Additionally, some evidence from FM developers, which are key customers of IaaS based on accelerated compute,²⁵⁰ showed they were unwilling or unable to substitute accelerated compute infrastructure with IaaS based on standard compute infrastructure. This was because of the differences in performance, efficiency and latency.²⁵¹ This is supported by the technical differentiation between IaaS based on accelerated compute and IaaS based on standard compute described above.

3.53 Third, we have assessed the supply-side substitutability of IaaS based on accelerated compute infrastructure and IaaS based on standard compute infrastructure.

3.54 Capacity shortages in the supply of accelerated compute infrastructure have prevailed despite excess capacity in standard compute infrastructure.²⁵² Enduring shortages suggest that it is not easy to shift from production of standard compute infrastructure into production of accelerated compute infrastructure, as otherwise we may have expected to have seen the productive assets used in standard compute infrastructure being shifted to help to meet demand for accelerated compute infrastructure.

3.55 We consider the fact that some cloud providers of standard compute also provide accelerated compute to be consistent with limited supply-side substitutability across asset types (ie cloud providers are incentivised to offer both standard and accelerated infrastructure to meet different customers’ needs). Some cloud providers supply both accelerated and standard compute, including some smaller cloud providers, such as Vultr.²⁵³ However, there is evidence of differences in the

²⁴⁸ [X] response to the CMA’s information request [X].

²⁴⁹ [X] response to the CMA’s information request [X].

²⁵⁰ Evidence assessed in chapter 4 indicates, for AWS, Microsoft and Google [X] the revenue associated with AI is generated from FM developers.

²⁵¹ Notes of meetings with [X]; [X] response to the CMA’s information request [X].

²⁵² WIRED (24 August 2023) ‘Nvidia Chip Shortages Leave AI Startups Scrambling for Computing Power’. Dell APJ chief says the industry won’t wait for Nvidia H100 - The Register, Nvidia sold half a million H100 AI GPUs in Q3 thanks to Meta, Facebook — lead times stretch up to 52 weeks: Report, accessed on 20 May. [X].

²⁵³ Vultr rides cloud GPU wave, widens ambitions, July 17 2023, accessed on 1 May 2025.

competitive conditions for IaaS based on accelerated and standard compute infrastructure. For example, although AWS, Microsoft and Google each supply both IaaS based on accelerated and standard compute infrastructure, their competitive positions differ between them – this is assessed in chapter 4.²⁵⁴ Furthermore, there are specialised cloud providers that are heavily focused on supplying IaaS based on accelerated compute infrastructure – eg [X] and [X]²⁵⁵ – and offer IaaS based on standard compute infrastructure only to a limited extent.²⁵⁶ We have assessed the competitive constraint these specialised providers impose as well as the barriers to entry and expansion they face in chapter 4.

3.56 As presented above, two cloud providers submitted that there was overlapping or common infrastructure between IaaS based on accelerated and standard compute, citing land, data centres, and network equipment.²⁵⁷ However, evidence received from one of these cloud providers and another, separate cloud provider showed that a large proportion of the costs associated with providing IaaS based on accelerated compute are not common costs:

- (a) Internal documents received from one cloud provider demonstrated that accelerated compute specific infrastructure (eg GPU-based servers) made up a large proportion of the total cost of operation (including data centre infrastructure and operation costs).²⁵⁸ We expect this experience to be common across suppliers of accelerated compute infrastructure, as the GPUs are common across their offerings.²⁵⁹ We note that the proportion spent on GPUs may fall over time if the price of GPUs falls relative to other costs.
- (b) A second cloud provider submitted that, of the total technical infrastructure investments associated with its public cloud division in 2025, approximately [X]% was expected to be allocated to ‘AI servers’ (powered by AI accelerator chips), with [X]% for ‘non-AI servers’ and the rest on data centres and networking that ‘generally support both AI and non-AI workloads’.²⁶⁰

²⁵⁴ Chapter 4 outlines evidence that Google is closer to Microsoft and AWS in terms of global accelerated compute capacity relative to its competitive position in the supply of IaaS based on standard compute.

²⁵⁵ Specialised providers of accelerated compute are discussed further in the later section setting out our assessment of the impact of AI on cloud services.

²⁵⁶ Responses to the CMA’s information request [X].

²⁵⁷ [AWS’ response to the CMA’s PDR dated 28 January 2025](#), paragraph 16; [Microsoft’s response to the CMA’s PDR dated 28 January 2025](#), paragraph 76 and 77(iii); AWS’ response to the CMA’s information request issued 28 March 2025, question 1, paragraph 1.3.

²⁵⁸ An internal [X] document set out that for an IaaS service based on [X]. [X] response to the CMA’s information request [X]. Another [X] internal document shows that for an IaaS service [X]. [X] response to the CMA’s information request [X].

²⁵⁹ As discussed in the chapter ‘Impact of AI on cloud services’, evidence from cloud providers shows that the majority of their accelerated compute capacity is powered by Nvidia’s GPUs.

²⁶⁰ [X] response to the CMA’s information [X].

- 3.57 Another cloud provider listed the specialised networking requirements of AI and high-performance computing as a factor causing its fixed asset costs to increase. The same provider also noted that additional power and space (and hence costs) were required to run and cool GPUs versus CPUs.²⁶¹
- 3.58 The above evidence shows that, while some infrastructure was common across IaaS based on accelerated compute and IaaS based on standard compute (eg land and some network equipment), the non-common infrastructure comprised a significant amount of the costs. This indicates that specific investments are needed from cloud providers to offer accelerated compute separate from those required to provide standard compute.
- 3.59 Some cloud providers' internal documents categorise investments separately from other cloud or data centre capex and they also track both their own and competitors existing and forecast metrics relevant to accelerated compute capacity (such as GPU numbers or capex) separately from IaaS based on standard compute.²⁶² For example:
- (a) In a [REDACTED].²⁶³ AWS said that it had invested a total of approximately [REDACTED].²⁶⁴
 - (b) In a finance review submitted to Microsoft's Board in July 2023, Microsoft's [REDACTED] between FY23 [REDACTED] and FY24 [REDACTED], whereas [REDACTED] in FY23 and [REDACTED] in FY24).²⁶⁵ Microsoft said that between July 2021 and June 2024 (including forecast spend), it expects to have spent in total [REDACTED].²⁶⁶
 - (c) In a Google [internal document], Google outlined how [REDACTED].²⁶⁷
- 3.60 The above evidence suggests that competing in the supply of IaaS based on accelerated compute infrastructure requires significant investments to either procure third-party accelerator chips or develop them in-house.
- 3.61 Finally, we considered how our assessment might be affected by future developments in the industry.
- 3.62 The technical capabilities of IaaS based on standard and accelerated compute are evolving, as are the demands of the workloads they support. In future, this may lead to changes (either more or less) in demand side substitution. However, the evidence above shows that there is currently a very limited degree of demand side

²⁶¹ Note of meeting with [REDACTED].

²⁶² For example, [REDACTED]. [REDACTED] response to the CMA's information request [REDACTED].

²⁶³ [REDACTED] response to the CMA's information request [REDACTED].

²⁶⁴ AWS' response to the CMA's information request [REDACTED].

²⁶⁵ Microsoft's response to the CMA's information request [REDACTED].

²⁶⁶ Microsoft's response to the CMA's information request [REDACTED]. [REDACTED].

²⁶⁷ Google's response to the CMA's information request [REDACTED].

substitution between IaaS based on accelerated compute infrastructure and IaaS based on standard compute infrastructure.

- 3.63 In future, overall customer demand for AI workload deployment could shift to IaaS based on standard compute if smaller and more efficient AI models (such as DeepSeek's R1) are favoured.²⁶⁸ Demand for AI deployment (inference) may also exceed that of AI training workloads, with demand for the former more easily being substituted for standard compute. On the other hand, some customer demand could shift from IaaS based on standard compute to accelerated compute in the future. Together, these factors could meaningfully advance the convergence of IaaS based on standard and accelerated compute. However, despite cloud providers' submissions (presented above) regarding this convergence, their internal documents demonstrated a clear focus on investment in accelerated compute, specifically GPUs and ASICs, for the purpose of meeting AI demand. This is discussed in chapter 4. This suggests that cloud providers are not currently anticipating or preparing for a significant degree of convergence between IaaS based on standard and accelerated compute.
- 3.64 Overall, the evidence discussed above shows that there is currently little demand and supply-side substitution between IaaS based on accelerated compute infrastructure and IaaS based on standard compute infrastructure. Therefore, our view is that IaaS based on accelerated compute is separate from the market for IaaS based on standard compute. However, this might change in future as competitive conditions and customer preferences evolve. We consider any interactions between IaaS based on accelerated compute and wider public cloud infrastructure services in chapter 4.

Whether there is a single market for IaaS and PaaS

- 3.65 PaaS provides 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. With PaaS, the customer has less control over the cloud stack compared to IaaS. Customers of PaaS manage applications and data but do not manage the PaaS computing platform (including its operating system) or the pre-built application components and tools.

²⁶⁸ Google's response to the CMA's information request [36]; [AWS' response to the CMA's PDR dated 28 January 2025](#), paragraph 16.

Cloud provider views

- 3.66 As described above, cloud providers generally submitted that IaaS and PaaS were substitutes and thus should be included in the same product market. We also received submissions from an academic.²⁶⁹

Customer views

- 3.67 We asked large customers about the mix of IaaS and PaaS they use and the main factors that influence their usage of these categories of public cloud services. Most customers we spoke to said they use a mix of IaaS and PaaS for their cloud workloads.²⁷⁰ In general, customers expressed their preference for either IaaS or PaaS, but some noted they may use the other in specific cases.²⁷¹
- 3.68 We asked large customers about the extent to which they regard IaaS and PaaS as substitutes for each other:
- (a) Some customers said that IaaS and PaaS are not substitutes or, even if IaaS and PaaS are technical substitutes, they would not switch between them because they have different characteristics.²⁷²
 - (i) Some of these customers said that IaaS offers more control and flexibility compared to PaaS. For example, one customer said that it frequently creates its own solutions based on IaaS because PaaS does not offer sufficient scale or Service Level Agreements, is higher cost and has an even higher level of lock-in.²⁷³
 - (ii) Similarly, another customer explained that PaaS is not a perfect alternative because it did not allow it to control its own services.²⁷⁴
 - (iii) Another customer said that, while it is possible to rearchitect and refactor an IaaS workload for PaaS, it uses PaaS sparingly because it is

²⁶⁹ Dr George R Barker, [Comment on The UK Competitive Market Authority's \(CMAs\) Cloud Services Market Investigation Provisional Decision Report](#); Dr George R Barker, [Comment on The UK Competitive Market Authority's \(CMAs\) Cloud Services Market Investigation Three Working Papers on The Supply of Public Cloud Infrastructure Services In the UK Covering The CMA's 1. Competitive Landscape Working Paper; 2. Egress Fee Working Paper; and 3. Committed Spend Agreements Working Paper](#). In response to our Competitive landscape working paper, Dr George R Barker submitted that the IaaS/PaaS/SaaS distinction used by the CMA cannot be tested by asking people questions about their beliefs or preferences in relation to the three CMA predetermined narrow products (Dr George R Barker, [Comment on The UK Competitive Market Authority's \(CMAs\) Cloud Services Market Investigation Three Working Papers on The Supply of Public Cloud Infrastructure Services In the UK Covering The CMA's 1. Competitive Landscape Working Paper; 2. Egress Fee Working Paper; and 3. Committed Spend Agreements Working Paper](#), page 35). We disagree with this submission and consider that we have followed our guidelines in defining the relevant product markets in this case. Dr Barker also submitted that the fact some customers shift from IaaS to PaaS is consistent with the market being competitive. We assess this point below where we consider customer evidence on the reasons why customers switch from IaaS to PaaS.

²⁷⁰ Responses to the CMA's information requests [REDACTED].

²⁷¹ See, for example, responses to the CMA's information requests [REDACTED].

²⁷² Responses to the CMA's information requests [REDACTED].

²⁷³ [REDACTED] response to the CMA's information request [REDACTED].

²⁷⁴ [REDACTED] response to the CMA's information request [REDACTED].

more constraining in terms of flexibility, propriety and specific ways of working as well as being more subject to lock-in.²⁷⁵

- (b) A few customers said that IaaS and PaaS have the potential to be substitutes for specific workloads.²⁷⁶
 - (i) For example, one customer said that substitutability is largely dependent on the service in context, but overall, there is minimal substitutability.²⁷⁷
 - (ii) Similarly, another customer stated that the capabilities of IaaS and PaaS are often not interchangeable and that it uses IaaS and PaaS to achieve different goals, but they can be substitutes in some cases.²⁷⁸
 - (iii) Another customer said that PaaS is a good substitute for IaaS when there is no customisation required at the operating level.²⁷⁹
- (c) A few customers said that IaaS and PaaS are easily substitutable for most workloads.²⁸⁰
- (d) Another customer said that it intends to switch from IaaS to PaaS, though this was due to a strategy shift to benefit from the lower management overhead and consumption-based cost model associated with PaaS.²⁸¹
- (e) One customer said that PaaS is first or second choice for most applications, including application refreshes for systems previously implemented as virtualised infrastructure. The customer explained it prefers the cost and operating advantages of managed services.²⁸²

Our assessment

- 3.69 Although cloud providers said that they do not necessarily segment their products along the IaaS/PaaS/SaaS categorisation, and there is a spectrum of products that do not necessarily fit neatly into these categories, this does not mean that the underlying products and services are in fact substitutes.
- 3.70 Customers for IaaS could be divided into different categories, including (a) ISVs who use IaaS as an input into the PaaS products they supply; and (b) other IaaS customers. We consider that the extent of demand-side substitutability between

²⁷⁵ [redacted] response to the CMA's information request [redacted].

²⁷⁶ Responses to the CMA's information requests [redacted].

²⁷⁷ [redacted] response to the CMA's information request [redacted].

²⁷⁸ [redacted] response to the CMA's information request [redacted].

²⁷⁹ [redacted] response to the CMA's information request [redacted].

²⁸⁰ Responses to the CMA's information requests [redacted].

²⁸¹ [redacted] response to the CMA's information request [redacted].

²⁸² [redacted] response to the CMA's information request [redacted].

IaaS infrastructure and PaaS services is limited for customers in both of these categories. This is because PaaS services rely on the infrastructure provided by IaaS providers which limits the competitive constraint PaaS can exert on IaaS (a change in price of IaaS would be passed on to PaaS).

- 3.71 Some customers said they had shifted from using IaaS to using PaaS. While this can indicate a degree of substitutability for some customers or specific workloads, it may also be the result of changes in customers' needs leading to different choices over time. If this were the case, we would categorise this shift as migration from one technology to the other, rather than substitutability.
- 3.72 In any case, we recognise that there is a spectrum of products and for some customers, for some workloads, IaaS and PaaS are substitutes. However, evidence from customers indicates that PaaS is not a good substitute for IaaS for most customers and workloads, and that most customers expressed that they are unwilling to substitute between the two, even if it may be technically possible to do so.
- 3.73 In light of the above, we consider there is not a sufficient degree of demand-side substitutability to warrant widening the IaaS market to include PaaS.
- 3.74 We also considered whether supply side substitution is sufficient to warrant widening the market to include PaaS. As set out above, to aggregate markets on the basis of supply-side substitution, we considered factors such as whether (i) suppliers supply a range of different products in the same broad category, using the same set of assets and capabilities; or (ii) suppliers regularly introduce new products or reposition existing ones within the category. While the three main cloud providers offer both IaaS and PaaS, it is unlikely that ISVs would be able to quickly or easily start providing IaaS, given the substantial investment in infrastructure that would be required (see chapter 5). Competitive conditions in IaaS and PaaS are different, for example see our assessment of shares of supply later in this chapter. Based on this, supply-side substitution is not sufficient to warrant widening the market to include PaaS.
- 3.75 As such, the market for IaaS should not be widened to include PaaS. We recognise that some PaaS services are more likely to pose a greater constraint on some IaaS services than other PaaS services, but we consider that the precise boundary of the market does not affect our analysis. In any case, we have also considered the market shares for IaaS and PaaS together to reflect this ambiguity and therefore provide context to the potential features that are relevant across both IaaS and PaaS.
- 3.76 In addition, we note that the evidence reviewed above is also relevant for the assessment of whether IaaS is substitutable for PaaS (ie responding to the question whether PaaS should be widened to include IaaS). The customer

evidence above indicates that there is limited demand-side substitutability. Competitive conditions in the supply of PaaS are significantly different from those in the supply of IaaS, as demonstrated by the large number of ISVs that do not supply IaaS services. In addition, IaaS providers rely on marketplaces to provide a full suite of PaaS products (supplied by ISVs) which complements (and sometimes is in competition with) their own PaaS offering. This suggests that while IaaS providers may own the key assets needed to produce and market some PaaS products, that is not true for the whole set of PaaS products.²⁸³

- 3.77 For these reasons, the PaaS market should not be widened to include IaaS, but note that the precise definition of the boundary of the market does not affect our competitive assessment.

PaaS

Whether PaaS should be segmented into separate, narrower markets

- 3.78 This section considers whether PaaS should be segmented into separate, narrower markets. We set out below the evidence we received in relation to market definition for PaaS before presenting our assessment. We did not receive customer evidence on the segmentation of PaaS.

Cloud provider submissions

- 3.79 As discussed at the beginning of the IaaS section above, cloud providers generally submitted that different elements of PaaS were substitutes and therefore should be in the same market.

Our assessment

- 3.80 PaaS consists of hundreds of individual products, many of which perform different functions; they are not precisely defined in the industry and each cloud provider uses different categories. Categories include, for example, databases, analytics, containers and machine learning. These categories are not mutually exclusive or collectively exhaustive and therefore are not straightforwardly subjected to demand- and supply-side substitution questions.
- 3.81 Although, based on demand-side and supply-side substitutability, PaaS could be subdivided into more granular segments,²⁸⁴ for the purposes of this market investigation it was not appropriate to do so. In this context, taking a rigid approach to applying demand-side substitution considerations to a large number of subcategory products which are both arbitrary and do not have clear boundaries

²⁸³ See 'Other providers' in chapter 2.

²⁸⁴ Including, for example, separating PaaS based on accelerate compute from PaaS based on standard compute.

would not have been helpful to the overall assessment of the wider supply of cloud services. Instead, we consider the features assessed in our competitive assessment manifest themselves in a similar way across all potential PaaS markets.²⁸⁵

- 3.82 Therefore, in line with our guidelines which state that we may treat a group of product markets together for the purposes of assessing competitive effects,²⁸⁶ we have considered PaaS services in aggregate for the purposes of this market investigation. Throughout the rest of the report, we refer to PaaS market as the aggregation of PaaS market segments as explained above.

Whether there is a single market for PaaS and SaaS²⁸⁷

Cloud provider submissions

- 3.83 As set out above, cloud providers generally submitted that IaaS, PaaS and SaaS were in the same market.
- 3.84 Microsoft said that there are grey areas between PaaS and SaaS services. In particular, it said that PaaS services can have some elements of SaaS solutions; but, while SaaS-like functions are delivered by PaaS services on its platform, Microsoft does not believe those products belong in the SaaS service model.²⁸⁸
- 3.85 Google said that the SaaS layer is even more fragmented than the IaaS or PaaS layers,²⁸⁹ and IBM said that SaaS is a very crowded field with many large to small companies, some of which also offer PaaS.²⁹⁰
- 3.86 The International Center for Law & Economics submitted that SaaS should not be considered as a meaningful layer of the market for cloud computing.²⁹¹

Customer views

- 3.87 We asked customers about the mix of PaaS and SaaS they use. Customers generally submitted that they use a mix of services.²⁹² Some customers suggested that they prefer a particular service model (ie IaaS, PaaS or SaaS).²⁹³ For example, one customer said that its principle is to operate as high up the stack as

²⁸⁵ CC3 (revised), paragraph 152.

²⁸⁶ CC3 (Revised), paragraph 152.

²⁸⁷ We have not received submission the substitutability of IaaS and SaaS, but given our conclusions on the market for PaaS and the even more marked different characteristics between IaaS and SaaS (relative to IaaS and PaaS) we consider IaaS to be in a separate market from SaaS.

²⁸⁸ Microsoft's response to the CMA's information request [REDACTED].

²⁸⁹ Google's response to the CMA's information request [REDACTED].

²⁹⁰ IBM's response to the CMA's information request [REDACTED].

²⁹¹ International Center for Law & Economics' response to the CMA's PDR, page 6.

²⁹² Responses to the CMA's information requests [REDACTED].

²⁹³ Responses to the CMA's information requests [REDACTED].

possible (ie use SaaS over PaaS and PaaS over IaaS) to minimise the amount of bespoke engineering work that it needs to perform and maintain.²⁹⁴

3.88 We asked customers about the extent to which they regard PaaS-based and SaaS-based solutions to be substitutes for each other. Customers generally said that they do not consider PaaS-based and SaaS-based solutions as substitutes.²⁹⁵ Most of these identified that PaaS and SaaS are used for different purposes, generally depending on the level of control the customer required over the application.²⁹⁶ For example:

- (a) One customer submitted that PaaS and SaaS serve distinct roles within its technology stack: it uses PaaS for developing and deploying customer applications due to the flexibility and control offer, whereas the customer uses SaaS for standard business applications due to the convenience and ease of use. The customer explained that the two solutions complement each other but address different needs.²⁹⁷
- (b) Similarly, another customer submitted that it uses PaaS where a more bespoke solution is required, but it uses SaaS services where the software requires less customisation and 'off-the-shelf' solutions are available.²⁹⁸
- (c) Another customer said that the two are not interchangeable at all, and explained that SaaS is the complete end to end delivery of an application, whereas PaaS is about the platform the applications run on.²⁹⁹

3.89 Some customers also submitted that the layers of cloud services require different skills and experience.³⁰⁰ For example, one customer said that PaaS is fundamentally not a substitute for SaaS and explained that each layer of cloud represents a need to have skilled personnel to manage it and develop against it.³⁰¹

3.90 Some customers indicated that SaaS may be a substitute for PaaS in certain circumstances.³⁰² For example:

- (a) One customer said that IaaS, PaaS and SaaS are all viable options for hosting models for its services, though noting that they all have architectural trade-offs and it predominantly chooses its workloads based on the skills it has available. 95% of this customer's workloads are IaaS.³⁰³

²⁹⁴ [redacted] response to the CMA's information request [redacted].

²⁹⁵ Responses to the CMA's information requests [redacted].

²⁹⁶ Responses to the CMA's information requests [redacted].

²⁹⁷ [redacted] to the CMA's information request [redacted].

²⁹⁸ [redacted] response to the CMA's information request [redacted].

²⁹⁹ [redacted] response to the CMA's information request [redacted].

³⁰⁰ Responses to the CMA's information requests [redacted].

³⁰¹ [redacted] response to the CMA's information request [redacted].

³⁰² Responses to the CMA's information requests to [redacted].

³⁰³ [redacted] response to the CMA's information request [redacted].

- (b) Another customer said that the substitutability of PaaS and SaaS needs to be considered on a use case by use case basis. It said that one of its principles is to adopt not adapt (ie use SaaS) where its strategy is to standardise and not differentiate via customisation.³⁰⁴

Our assessment

- 3.91 Customers use a combination of SaaS and PaaS (and IaaS) products. Customers almost universally submitted that SaaS is not a good substitute for PaaS, and many customers explained that the two serve different business requirements. Some customers prefer particular service models for their workloads based on factors including existing skills and knowledge, the need for customisation of applications and/or the costs of the different models.
- 3.92 In relation to supply-side substitution, based on the evidence we have set out above, competitive conditions for PaaS and SaaS are significantly different and there does not appear to be supply-side substitution between the two layers.
- 3.93 On this basis, we do not consider that PaaS should be widened to include SaaS.

Alternative IT models

- 3.94 This section considers the extent to which either (i) traditional (on-premises) IT (ie dedicated computing resources on-premises) or (ii) private cloud services (ie a cloud which is exclusive to one customer) are substitutes for IaaS and/or PaaS. Together, we refer to these as alternative IT models or environments.
- 3.95 In assessing the constraint from traditional on-premises IT and private cloud services, we sought to understand the extent to which customers find these to be suitable alternatives to cloud services. The extent that customers have moved from one environment to another for reasons other than changes in quality or price (for example due to changes in preferences or adjustments after trial and error) does not in itself enable us to draw strong inferences on the relevant market definition.
- 3.96 Customers may shift to a new technology and then be unwilling to shift back, meaning that the observation of ongoing 'switching' does not imply that customers would substitute back and forth in response to shifts in competitive offerings. Similarly, some customers may shift back to an old technology having tried a new technology and decided it was inappropriate for their requirements. Customers may also use new and old technologies as complements, deploying them for

³⁰⁴ [REDACTED] response to the CMA's information request [REDACTED].

different purposes, but requiring some movement of data, resources or spend between them.

- 3.97 We undertook this market investigation during a period of significant migration to public cloud, which was characterised by a process of customers learning which environment is most suitable for their requirements and moving their workloads accordingly. Therefore, we considered the evidence in light of the possibility that switching may arise from competitive diversion and/or migration trends and the associated learning process.
- 3.98 In our assessment of the evidence, we were mindful of the potential for asymmetric constraints and migration effects to create patterns of shifting demand that may not be related to ongoing competition. We also considered evidence on the extent to which traditional on-premises IT and private cloud environments represent a good substitute for customers of public cloud.

Cloud provider submissions

Traditional on-premises IT

- 3.99 AWS and Google said that from the customer's perspective, on-premises and public cloud are substitutable for most use cases.³⁰⁵ In relation to this, AWS submitted that customers view a wide range of IT services as substitutable, including on-premises, the cloud, or some combination. AWS further submitted that customers assess their IT needs on a workload-by-workload basis,³⁰⁶ and that customers typically look to solve a specific IT problem and then look at a broad set of options that might meet them. AWS submitted that customers rarely look simply to use the cloud as an end in itself.³⁰⁷
- 3.100 AWS also submitted that citing switching costs as a reason for customers not considering switching from cloud services to on-premises services is an irrelevant critique. It submitted that switching costs accrue regardless of whether a customer is switching from cloud to on-premises or between cloud providers, adding that, by the PDR's logic, every IT services provider (whether cloud or on-premises) would constitute a separate market for its services, which is not the case.³⁰⁸
- 3.101 We acknowledge that switching costs alone would not be sufficient to support a finding that on-premises IT is not in the same market as public cloud. However, below we considered wider factors based on a range of customer evidence. Finally, we note that the framing of the relevant question that we asked these

³⁰⁵ AWS' response to the CMA's information request [🔗]; Google's response to the CMA's information request [🔗].

³⁰⁶ [AWS' response to the CMA's working papers and updated issues statement, 25 June 2024](#) paragraph 7.

³⁰⁷ AWS' response to the CMA's information request [🔗]; [AWS' response to the CMA's PDR](#) dated 28 January 2025, paragraph 10.

³⁰⁸ AWS' submission to the CMA [🔗].

customers means that their answers are likely to refer to switching costs that are specific to switching between public cloud and non-public cloud IT environments.³⁰⁹

3.102 Some cloud providers also made submissions relating to whether, and to what extent, public cloud and traditional IT meet customers' requirements when deciding where to place different types of workloads. In this regard:

- (a) AWS submitted that many customers consider on-premises services to be more suitable than cloud services for a specific workload. This could be due to latency, customers' own preferences, the nature of customers' existing IT architecture, or portability issues imposed by Microsoft's licensing restrictions. AWS submitted that, in light of this, it has invested heavily in solutions that support on-premises infrastructure and AWS' services, including AWS Application Migration Services, AWS Database Migration Services, Kubernetes and direct connections to on-premises environments from AWS, point-to-point connections and site-to-site VPNs, among many other investments.³¹⁰
- (b) AWS also said that competitive pressure exerted by on-premises IT providers means customers do switch between cloud and on-premises solutions or use them alongside each other. AWS submitted that, as a result, customers often opt for a hybrid solution, using both on-premises and cloud services providers to fulfil the different aspects of their IT needs.³¹¹
- (c) Okta submitted that most organisations use hybrid cloud, and this is mostly because they have legacy on-premises systems that are difficult to move to a fully cloud-based system.³¹²
- (d) Google said that the key difference between the two solutions is flexibility and scalability, rather than technical differentiation. It also said that traditional IT, unlike cloud services, is costly and can be inefficient.³¹³
- (e) Microsoft said that traditional IT solutions suit customers with requirements for a high degree of control over their IT, or that have to comply with security or regulatory requirements.³¹⁴ It also said that traditional IT can be more resource-intensive to manage and maintain.³¹⁵ Microsoft said that, in comparison, public cloud meets a broad and diverse array of customer needs, including customers that do not want to make large investments in

³⁰⁹ The question explicitly identified non-public cloud IT environments and public cloud.

³¹⁰ [AWS response to the CMA's working papers and updated issues statement, 25 June 2024](#) paragraph 9.

³¹¹ [AWS response to the CMA's working papers and updated issues statement, 25 June 2024](#) paragraph 10.

³¹² Note of meeting with Okta, [REDACTED].

³¹³ Google's response to the CMA's information request [REDACTED].

³¹⁴ Microsoft's response to the CMA's information request [REDACTED].

³¹⁵ Microsoft's response to the CMA's information request [REDACTED].

their IT solutions or companies that are looking for access to the most cutting-edge computing capabilities.³¹⁶

- (f) Microsoft also submitted that cloud has advantages relative to other forms of IT models and there have been significant technological developments in the industry.³¹⁷

3.103 AWS submitted that traditional IT solutions are exerting competitive pressure on cloud services providers,³¹⁸ and customers switch regularly between cloud and traditional IT.³¹⁹ We set out the evidence AWS provided in this regard below:

- (a) AWS submitted several examples of customers that switched to hybrid solutions. For example, one customer moved its data out of the provider to a hybrid on-premises and cloud solution, and another customer moved from the cloud to a hybrid solution by building its own network of servers to 'switch seamlessly between cloud providers and its own servers'.³²⁰
- (b) AWS told us about a news article which identified a customer that said it boosted its profits from leaving the cloud.³²¹ The news article also referred to a survey by Citrix that found that 94% of large US organisations it surveyed had worked on repatriating data or workloads from the cloud in the last three years.³²² AWS also submitted that the Flexera 2025 State of the Cloud report finds that 21% of cloud-based workloads and data have been repatriated from cloud providers back to on-premises.³²³
- (c) AWS submitted that [REDACTED] customers choose from a variety of options including on-premises data centres, off-site data centres dedicated to a specific user (ie private cloud), hosted installations, co-located data centres, managed service providers, software providers, and cloud offerings. In particular, AWS refers to some internal documents that assess the market and name several competitors across the wider IT industry.³²⁴
- (d) AWS submitted survey evidence on switching. We set out this evidence and our assessment of it below.
- (e) Finally, AWS submitted that it is implausible that a substantial share (between 24% and 36%, based on a series of external quantitative surveys³²⁵) of cloud customers who have migrated or plan to migrate

³¹⁶ Microsoft's response to the CMA's information request [REDACTED].

³¹⁷ Microsoft's submission to the CMA [REDACTED].

³¹⁸ [AWS' Economic response to the CMA's Competitive landscape working paper](#), paragraph 5.

³¹⁹ AWS' response to the CMA's information request [REDACTED].

³²⁰ AWS' response to the CMA's information request [REDACTED].

³²¹ Transcript of hearing with AWS [REDACTED].

³²² [Are rainy days ahead for cloud computing? - BBC News](#).

³²³ AWS' submission to the CMA [REDACTED].

³²⁴ AWS' submission to the CMA [REDACTED].

³²⁵ [See Economic response to the CMA's Competitive Landscape Working Paper](#), paragraphs 24-26.

workloads from cloud to on-premises have all or mostly done so for non-competitive reasons.³²⁶ It added that examples of reasons for switching – such as latency constraints, poor client application architecture, and suboptimal performance (and price) – are all reflective of relative quality or price dimensions between cloud and on-premises solutions. AWS submitted that these cases suggest that customers initially migrated workloads from on-premises to the cloud but ultimately chose to move them back to on-premises after assessing and preferring its relative quality and price parameters. AWS said that this is competition at play and exactly the type of switching that is relevant for a competitive assessment.³²⁷

- 3.104 AWS, Google and Oracle said that there are very few technical differences between on-premises and cloud services.³²⁸ Oracle added that the difference is in method of delivery, not the underlying technology.³²⁹ We do not consider these arguments to be probative of the constraint imposed on public cloud by on-premises IT because 'the method of delivery' is a key characteristic of cloud technology which influences the customer choice of cloud over on-premises. This is supported by evidence described in appendix A where cloud providers noted, among others, flexibility, elasticity, pay-as-you-go and the ability to quickly increase resources as determining factors in the choice of a cloud provider.
- 3.105 AWS submitted that on-premises IT providers are continuing to invest in order to retain existing customers and attract customers away from cloud services providers.³³⁰ In addition, AWS cited research that said that Dell, VMware and HPE are pivoting towards cloud-like 'as a service' operating models, which AWS said shows strong supply side substitutability.
- 3.106 The International Center for Law & Economics submitted that cloud computing is one component of IT. It added that investments in cloud computing represent a small portion of global IT spending, with estimates ranging from 7 to 12%.³³¹

Private cloud

- 3.107 AWS said that, from a customer's perspective, IT services are substitutable for most use cases, regardless of their delivery method.³³²

³²⁶ AWS' submission to the CMA [REDACTED].

³²⁷ AWS' submission to the CMA [REDACTED].

³²⁸ AWS' response to the CMA's information request [REDACTED]; Google's response to the CMA's information request [REDACTED]; Oracle's response to the CMA's information request [REDACTED].

³²⁹ Oracle's response to the CMA's information request [REDACTED].

³³⁰ [AWS' Economic response to the CMA's Competitive landscape working paper](#) paragraph 21 and 22.

³³¹ [International Center for Law & Economics' response to the PDR](#), page 4.

³³² AWS' response to the CMA's information request [REDACTED].

- 3.108 Similarly, Google said that customers multisource and switch between different deployment models and technologies, which highlights that these hosting methods are largely interchangeable from a technical perspective.³³³
- 3.109 Google, Microsoft, Oracle and IBM said that private cloud offers more control, security and reduced latency, while public cloud has greater scalability, flexibility and may be cheaper.³³⁴
- 3.110 Google said that even for customers with latency or security requirements there remains strong competition for and substitutability between the different deployment models and technologies.³³⁵ It provided an example where a customer chose instead to invest in a private cloud platform.³³⁶
- 3.111 AWS provided an example of a customer switching from its public cloud to private cloud, though the customer ultimately brought workloads back to public cloud to benefit from the continuous investment and innovation in hardware and other services.³³⁷
- 3.112 [X] submitted that it competes for customers both with other private cloud suppliers and with the large public cloud providers. It submitted that it is technically possible for large customers to switch providers from public to private cloud. [X] added that, for small and medium enterprise (SME) customers, competition may occur between private cloud providers and AWS/Azure.³³⁸

Our assessment of cloud provider submissions

- 3.113 With respect to AWS' submission that the use of cloud is not 'an end in itself', we note that a product does not need to be consumed as 'an end in itself' for it to be in a separate market. In particular, even if customers are not seeking to use public cloud as an end to itself, they may find that their requirements are best met by public cloud. Similarly, they may find that their needs are not met sufficiently well by private cloud. To the extent that customers do not view alternative IT as a good substitute (because alternative IT does not meet those requirements well), we could still find that public cloud is a separate market.
- 3.114 Customer evidence in the round shows that on-premises and private cloud are generally not good substitutes for public cloud. We note that some cloud providers' submissions, for example AWS' submission that customers consider on-premises services more suitable than cloud services for specific workloads, or Microsoft's

³³³ Google's response to the CMA's information request [X].

³³⁴ Google's response to the CMA's information request [X]; Microsoft's response to the CMA's information request [X]; Oracle's response to the CMA's information request [X]; IBM's response to the CMA's information request [X].

³³⁵ Google's response to the CMA's information request [X].

³³⁶ Google's response to the CMA's information request [X].

³³⁷ AWS' response to the CMA's information request [X].

³³⁸ [X] submission to the CMA [X].

submission that traditional IT solutions suit customers with specific requirements, are consistent with on-premises IT being suitable only in certain circumstances and are therefore not inconsistent with a finding of separate markets.

- 3.115 With respect to AWS' submissions that customers use hybrid solutions across public cloud and one of either private cloud or traditional IT, we note that using two different products alongside each other is ambiguous as to whether it is indicative of substitutability given that parallel use would be consistent with two products being complementary. The use of hybrid services would be consistent with firms using public cloud services for certain use cases and on-premises IT or private cloud for other use cases. It is therefore informative to consider evidence on substitutability, including any information on use cases or benefits of different types of IT environment in different scenarios, between it and public cloud.
- 3.116 With respect to anecdotal examples of customers switching, we note that this shows only that such switching occurs. We acknowledge that it is technically possible to switch workloads from public to private clouds (or to traditional IT). Therefore, we acknowledge that on-premises may exert some competitive constraint on cloud services for some customers, and some customers may switch between the two IT models for competitive reasons, but we have considered this alongside other evidence to assess the significance of this constraint. We add that, even if AWS' submission that 24-36% of cloud customers have migrated or plan to migrate workloads from cloud to on-premises solutions was correct,³³⁹ this means 64-76% have never migrated or planned to migrate any workloads, which is also relevant.
- 3.117 We have seen some AWS internal documents that reference competitors including on-premises data centres and hosted-installation providers, cloud providers and other IT services providers. One document considered only other public cloud providers when discussing competitors to its storage, compute and database offering, but also mentioned other on-premises IT providers in a more general discussion.³⁴⁰ Another internal document from AWS compares its profits growth to that of [X] and [X].³⁴¹ We have also seen other internal documents from AWS and Google that focus on other cloud providers only.³⁴² The documents we have reviewed suggest public cloud providers look at a variety of IT providers, including providers of alternative IT models, but with a strong focus on public cloud providers. This is consistent with on-premises and private cloud providers

³³⁹ We place limited evidential weight on the surveys that AWS submitted on switching workloads to traditional IT. We set out our reasoning for this in 'Evidence from surveys' below.

³⁴⁰ AWS' response to the CMA's information request [X]. AWS submitted that some of the competitors listed also provide on premises and hybrid solutions.

³⁴¹ AWS' response to the CMA's information request [X].

³⁴² See chapter 7.

providing a weaker constraint than public cloud providers. This is also consistent with customer evidence assessed below.

- 3.118 The cloud providers' submissions include several assertions that cloud and alternative IT models are substitutable. We have assessed these statements alongside other evidence on substitutability which includes views from various stakeholders, responses to requests for information sent to cloud customers, and qualitative market research that we commissioned.
- 3.119 With respect to the ICLE's submission that cloud computing is just one component of IT, we note we have considered the extent to which traditional IT is a good alternative to public cloud and have found that most customers are unable or unwilling to switch from public cloud to traditional IT for most workloads. As such, the view that cloud computing is simply a component of IT is not a relevant consideration.

Customer views

- 3.120 We asked large customers about the extent to which they consider IT environments other than the public cloud (namely private cloud and traditional IT) a substitute for public cloud computing and about the main factors that influenced their choice between the two. Most customers we spoke to answered by referring to non-public cloud without specifying traditional IT or private cloud, but where possible we have differentiated their answers between the two.
- 3.121 Most customers said that non-public cloud IT environments were not a good substitute for public cloud and they would not switch from public cloud to non-public cloud IT environments,³⁴³ or that doing so would involve significant resources.³⁴⁴ Some of these customers said that public cloud had different characteristics, such as advanced functionality, cost, elasticity, scalability, flexibility and resiliency.³⁴⁵
- (a) One customer said that public cloud offers 'unparalleled' capability to provision and use infrastructure components at scale, in multiple geographies, in a timely manner and with the latest functionality and security features. This customer said that, in comparison, it uses non-public cloud environments where public cloud cannot meet its requirements, such as for latency or operational risk.³⁴⁶

³⁴³ As noted above, the framing of the question that these customers answered means that their answers are likely to refer to switching costs that are specific to switching between public cloud and non-public cloud IT environments.

³⁴⁴ Responses to the CMA's information requests [X].

³⁴⁵ Responses to the CMA's information requests [X].

³⁴⁶ [X] response to the CMA's information request [X].

- (b) One customer said that it would only move workloads to the public cloud where technical benefits, such as scalability, elasticity and resiliency can be achieved. The customer said that it is guided by the specifics of the requirements for the workload to determine the best solution for the organisation.³⁴⁷
- (c) Another customer said that, while it uses on-premises IT for legacy applications, it would take significant effort to construct a stack with the full security and observability capabilities as virtual machines, and even then would still miss many elements of a public cloud value proposition.³⁴⁸
- (d) Similarly, another customer said that the public cloud is materially more cost-effective, gives it flexibility and scale, allows it to rapidly respond to changes in demand and provides a level of innovative services that non-public clouds do not.³⁴⁹
- (e) Some customers also said that they have adopted a cloud first policy that specifically targets moving their workloads to cloud services. For example, one customer's technology strategy aims for 'SaaS first, cloud second, on prem only where necessary'.³⁵⁰ This suggests that some customers do not see cloud and on-premises IT as substitutable.

3.122 Some customers said that non-public cloud IT environments were only a substitute for public cloud for certain workloads or had significant qualifications about how good of a substitute it was. For example:³⁵¹

- (a) One customer said non-public clouds can substitute in a limited setting, such as hosting virtual machines and layering other components on top, such as SQL Server to mimic SQL PaaS;³⁵²
- (b) Similarly, another customer said that substitutability depends on the application stack development;³⁵³
- (c) One customer said that applications that use purely basic IaaS services can be run on non-public cloud environments, but even in this scenario there are complexities due to networking and security configurations;³⁵⁴ and
- (d) Another customer said that, while currently its applications could be run in non-public cloud, it is feasible that it will evolve to an extent that those

³⁴⁷ [REDACTED] response to the CMA's information request [REDACTED].

³⁴⁸ [REDACTED] response to the CMA's information request [REDACTED].

³⁴⁹ [REDACTED] response to the CMA's information request [REDACTED].

³⁵⁰ [REDACTED] response to the CMA's information request [REDACTED].

³⁵¹ Responses to the CMA's information requests [REDACTED].

³⁵² [REDACTED] response to the CMA's information request [REDACTED].

³⁵³ [REDACTED] response to the CMA's information request [REDACTED].

³⁵⁴ [REDACTED] response to the CMA's information request [REDACTED].

applications may not be suitable for on-premise hosting. This customer also said that non-public cloud environments work best for certain types of workloads, such as those with low latency requirements or certain types of data.³⁵⁵

3.123 Two customers said that non-public cloud IT environments are a good substitute for public cloud.³⁵⁶

(a) One customer said that most of its workloads can be relatively easily migrated to non-public cloud environments.³⁵⁷

(b) Another customer said that most environments are a viable substitute.³⁵⁸

3.124 We also asked customers whether they had ever switched any workloads hosted on a public cloud to non-public cloud IT environments instead, such as private cloud or traditional IT. We also asked what workloads they switched and the reasons why. This question can be expected to have identified both (i) switching due to competitive diversion and (ii) switching because customers' requirements were not met by public cloud. Where possible, we have set out customers' reasons for switching, in order to inform whether non-public cloud IT environments pose a strong constraint on the remaining workloads.³⁵⁹ We set out the evidence below:

(a) Most customers said that they had not switched to on-premises,³⁶⁰ although two of these were exploring switching.³⁶¹

(b) All of the customers that did switch only did so for certain workloads, and some only did so because the characteristics of public cloud did not meet the requirements for those workloads, rather than as a response to a price rise or degradation of quality.³⁶² For example:

(i) One customer switched because of application latency constraints and poor client application architecture,³⁶³ and

³⁵⁵ [REDACTED] response to the CMA's information request [REDACTED].

³⁵⁶ Responses to the CMA's information requests [REDACTED].

³⁵⁷ [REDACTED] response to the CMA's information request [REDACTED].

³⁵⁸ [REDACTED] response to the CMA's information request [REDACTED].

³⁵⁹ This question asked for the reasons customers switched, but not for the reasons why customers did not switch. As such, the question was inherently weighted towards switching behaviours and the answers should be considered in light of this tendency to provide greater detail in relation to switching.

³⁶⁰ Responses to the CMA's information requests [REDACTED].

³⁶¹ Responses to the CMA's information requests to [REDACTED].

³⁶² Responses to the CMA's information requests [REDACTED].

³⁶³ [REDACTED] response to the CMA's information request [REDACTED].

(ii) Another customer was unable to get the right level of performance and price in public cloud, and so switched back to traditional data centre services.³⁶⁴

(c) One customer said that it had switched to private cloud and another had switched to a non-public cloud environment, but it was unclear whether this was to private cloud or traditional IT. The former customer said that it had moved some workloads to its private cloud, but it was extremely complicated to move workloads from one location to another, and it did not wish to do so frequently.³⁶⁵ No other customers we spoke to identified moving to the private cloud.

3.125 A handful of customers submitted evidence on the substitutability of private cloud for public cloud. These customers said that the public cloud can technically be replicated on private cloud.³⁶⁶

(a) Of these customers, most qualified their answer, saying that there would be cost and/or effort involved in switching.³⁶⁷

(b) A few of these said that their operations are subject to significant variation in loads over time, some of which may be better suited to a public cloud environment.³⁶⁸

(c) Further, one customer said that, while it could do everything it needed on a private cloud, it would take a huge amount of time, effort and re-engineering and would be incredibly disruptive.³⁶⁹

(d) Another customer said that private cloud potentially results in the loss of shared development, but private and public cloud use fundamentally the same technology.³⁷⁰

3.126 A supplier of professional services also commented on customer needs.³⁷¹ It submitted that there was a growing competitive constraint from on-premises and private cloud. It also said that private cloud was not an exact replica of public cloud, but it was a viable environment for anything which is high volume, of relatively low value and typically not at 'the bleeding edge' of innovation. It added that public cloud was superior for the best scalability and innovative services and

³⁶⁴ [REDACTED] response to the CMA's information request [REDACTED].

³⁶⁵ [REDACTED] response to the CMA's information request [REDACTED].

³⁶⁶ Responses to the CMA's information requests [REDACTED].

³⁶⁷ Responses to the CMA's information requests [REDACTED].

³⁶⁸ Responses to the CMA's information requests [REDACTED].

³⁶⁹ [REDACTED] response to the CMA's information requests [REDACTED].

³⁷⁰ [REDACTED] response to the CMA's information requests [REDACTED].

³⁷¹ Suppliers of professional services assist customers with their cloud strategy (see, for example, [Cloud Offerings - KPMG UK](#); [Cloud Consulting Services & Solutions | Accenture](#)), which may include negotiating with cloud providers on customers behalf.

that customers might choose public cloud for other reasons such as technical strategy or commitments with cloud providers.³⁷²

Jigsaw report

- 3.127 The Jigsaw report said that the customers spoken to generally perceived that the benefits of public cloud matched their requirements better than traditional IT or private cloud. For example, customers spoken to said that public cloud was reliable and came with less risk to continuity than on-premises or private cloud.³⁷³ These customers added that using public cloud was a requirement of doing business and in many cases, they could not operate at all or as effectively without use of public cloud.³⁷⁴
- 3.128 The Jigsaw report also noted that some customers said that for some use cases the high costs of public cloud had been a significant driver to set up their own private data centre. The Jigsaw report said that high-cost use cases include machine learning via GPUs (we have considered GPUs in the section on accelerated compute above) and data transfer (we have considered the costs of data transfer in the section on egress fees in the Barriers to switching and multi-cloud chapter).³⁷⁵

Evidence from surveys

- 3.129 AWS submitted evidence from several surveys on customers switching from public cloud to traditional IT. AWS said that these surveys showed that there is a high diversion ratio, and estimated rates of cloud repatriation show that it was incorrect to exclude on-prem from the competitive assessment.³⁷⁶ We set out some details from these surveys below:
- (a) One survey asked customers if they had migrated applications/workloads from a hyperscale public cloud to some other venue in the past 12 months. The survey included AWS, Azure, GCP, IBM and Alibaba as examples of hyperscale public clouds. AWS said that the results showed that, of the organisations that switched workstreams away from 'hyperscalers', 65% switched to their own on-prem services.³⁷⁷
 - (b) Another survey asked the customers spoken to that had ever switched cloud infrastructure providers if they switched to another cloud services provider or an on-premises solution. AWS submitted that the results showed that 30% of the customers spoken to switched to on-prem services (when excluding

³⁷² Note of meeting with [REDACTED].

³⁷³ [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), page 25.

³⁷⁴ [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), page 25.

³⁷⁵ [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), page 26.

³⁷⁶ [AWS' economic response to the Competitive landscape working paper](#), paragraph 5.

³⁷⁷ [AWS' economic response to the Competitive landscape working paper](#), paragraph 13.

private cloud customers).³⁷⁸ 29% of the 26% that had ever switched tells us that, overall, 7.5% of IaaS/PaaS users that responded submitted that they switched to an on-premises solution.³⁷⁹

- (c) AWS said that in addition to diversion ratios, it is also informative to consider the percentage of cloud customers that have moved any workloads back to on-prem. AWS submitted that multiple other surveys show that the share of cloud customers that have moved workloads back to on-prem or plan to do so is significant, between 24% and 36%.³⁸⁰ AWS submitted that Barclays' survey of CIOs also found that 83% of enterprise CIOs that responded to the survey plan to repatriate at least some IT workloads from the cloud to private cloud or on-premises in 2024.

3.130 We set out our observations on this evidence below.

3.131 We place limited evidential weight on the surveys that AWS submitted on switching workloads to traditional IT. There are significant weaknesses to using quantitative surveys in this market for our purposes due to ambiguity in terminology and vulnerability to the representativeness of the customers spoken to.³⁸¹

3.132 Further, we note that even if we did not have any reservations about the probative value of the research, the surveys do not necessarily indicate that traditional IT should be included in the market. In particular:

- (a) The statistics calculated are more akin to a 'switching ratio' than a diversion ratio because they capture all customers switching workloads, irrespective of changes in price or quality. Switching may arise from competitive diversion and/or migration trends and the associated learning process. In particular, customers may move workloads to the public cloud and, in doing so, learn that their requirements are not well met by public cloud and consequently move these workloads back to alternative IT models. For survey evidence to be informative on the constraint from traditional IT, it is necessary to focus on whether and/or to what extent customers would use traditional IT instead of public cloud (not after or in addition), in particular in response to an increase in the price of public cloud by 5% to 10%. Without sufficient clarity on the

³⁷⁸ AWS analysis of [REDACTED].

³⁷⁹ CMA analysis of [REDACTED].

³⁸⁰ [AWS' economic response to the Competitive landscape working paper](#), paragraph 24.

³⁸¹ As set out in appendix C. In fact, these weaknesses are recognised by the research that AWS submitted: The research report submitted by AWS recognises that the repatriation terminology is 'imprecise' because it implies a permanent and absolute transition away from the public cloud, when a shift may just be temporary or involve a limited number of workloads. We consider that such ambiguity makes it difficult to interpret the results in a meaningful way. Further, the research report submitted by AWS indicated that the results are sensitive to the survey used, stating that analysing the results of 'a different survey gives us a different story'.

[Cloud repatriation: What it is, what it isn't, and why it's not going away](#), pages 1 and 3.

motivation for customers' switching, survey evidence on a switching ratio is ambiguous.

- (i) As such, we have considered the reasons for switching and the extent to which customers view traditional IT as a good alternative to public cloud (set out in 'Customer views' above).³⁸²
- (b) The switching ratios overstate the strength of traditional IT. For example, the research identified above gave customers that had ever switched one of their cloud infrastructure providers a binary choice between whether they switched to another cloud services provider or an on-premises solution. We note that:
 - (i) It is not clear how many workloads switched to either environment, nor what the value of those workloads were. To highlight the issues with this, we can imagine a customer that switched one workload to traditional IT and another customer that switched all its workloads traditional IT. These customers would be equally weighted in this survey. While AWS submitted that it is informative to consider the percentage of cloud customers that have moved any workloads, we consider that traditional IT is likely a significantly more important constraint in the latter scenario compared to the former.
 - (ii) The binary choice means that customers who switched different workloads to different destinations would have been classified only as switching to traditional IT or switching to another cloud provider, when they should have been classified as switching to both. We therefore cannot observe the switching ratio based on responses to this question.
 - (i) The use of the term 'ever' creates uncertainty over the frequency and recency of any switching reported by customers, when both factors are relevant in considering the extent of the competitive constraint from traditional IT. In particular, we are more likely to consider traditional IT as a relevant constraint if a customer switches frequently and regularly than a customer with an isolated instance of switching years before.
- (c) We also note that some of the survey evidence suggests that there is a low overall switching ratio to traditional IT. For example, just 7.5% of IaaS/PaaS users surveyed in the research set out above had ever switched to on-

³⁸² While we consider it is appropriate to place limited weight on this survey evidence, we note that the research itself also considers the reasons for switching. In particular, the research asked the customers spoken to about the top factor behind moving workloads away from hyperscalers' (defined below) public clouds. Some customers' answers seem to relate to the characteristics of the cloud not matching with the requirements for that workload, including application lifecycle considerations (ie different IT environments for test/dev and production), regulatory/governance requirements and data locality or sovereignty. This would suggest that the switching prevalence estimated is unlikely to exclusively represent competitive diversion.

[Cloud repatriation: What it is, what it isn't, and why it's not going away](#), pages 3 and 4.

premises. We do not consider that this is consistent with a situation where traditional IT poses a significant competitive constraint on cloud services.

Our assessment of alternative IT models

- 3.133 The evidence shows that most customers do not consider alternative IT models to be a good substitute for public cloud, whether IaaS or PaaS solutions. Customers indicated that their requirements are not well met by alternative IT models and identified distinct characteristics of public cloud, including advanced functionality, cost, elasticity, scalability, flexibility and resiliency.
- 3.134 While some customers identified that alternative IT models are a good substitute for public cloud, these answers are mostly caveated, qualified or only for certain workloads or situations. Where customers had switched to alternative IT models, they generally indicated that they had done so because the workloads were not well suited for public cloud in the first place.
- 3.135 The evidence we have seen indicates that some public cloud workloads could technically be hosted on private cloud. However, evidence from customers shows that non-public cloud environments are generally poor substitutes for public cloud, and that there would be significant costs and time associated with switching to private cloud, with only two customers indicating that they had previously switched.
- 3.136 While some responses to our request for information implied that on-premises solutions do constrain public cloud solutions, customers generally responded that non-public cloud IT environments are not a good substitute for public cloud, with customers saying they would not switch from public cloud to non-public cloud IT environments, or that doing so would involve significant resources. The Jigsaw report also found that customers of public cloud consider that their requirements are best matched by public cloud, and the characteristics of public cloud are not matched well by alternative IT models. The report found that for some customers, using public cloud was a requirement of doing business and in many cases, they could not operate at all or as effectively without use of public cloud.
- 3.137 In this context, we did not find AWS' customer anecdotes of switching to be persuasive and there was no way to verify that the examples were representative. It is appropriate to attach weight to the customer evidence we have gathered using random selection and for which we have gathered information on the context and background to the decisions being made.
- 3.138 In addition, AWS' examples themselves were not necessarily inconsistent with the reasoning set out above. For example, the news article referenced by AWS quoted that '[t]here is a repatriation going on of things that should never have been in the

cloud or that won't work in the cloud'.³⁸³ This indicates that there is a learning process associated with migration to the public cloud (as we set out above). However, this form of switching reflects that their requirements for the workload were not satisfied by public cloud, rather than directly addressing whether the alternative IT models are substitutable.

- 3.139 While we acknowledge that, as submitted by AWS and discussed above, some IT providers are turning to 'cloud-like' services, we consider customer evidence to be more relevant to understand the competitive constraint from alternative IT models to public cloud. This evidence is discussed in the paragraphs above.
- 3.140 AWS also submitted several surveys, which it said show traditional on-premises IT should be included in the market. However, we consider it appropriate to attach only limited weight to quantitative surveys in this market, as we set out above. Further, even if we were to look past the issues with quantitative survey evidence in this market, we do not think the surveys are persuasive. The switching statistics produced cannot distinguish between competitive diversion and switching due to a learning process, the switching ratios overstate the strength of traditional IT, and the surveys suggest that some customers switch for reasons that are not inconsistent with separate markets. The switching ratios found by the surveys also have characteristics that risk overstating the strength of constraint from traditional on-premises IT including not considering the proportion of revenue that switched. In any case, in some instances the switching ratios found by these surveys are low.
- 3.141 The submission that customers use traditional on-premises IT and private cloud in hybrid solutions with public cloud provides little evidence on substitutability, and may instead be indicative of complementarity. Indeed, all customers that said they would not substitute their public cloud workloads to non-public cloud environments also had hybrid environments.
- 3.142 Based on the evidence set out above, we do not consider traditional on-premises IT a close substitute for public cloud for a high proportion of customers. Most customers would likely not view on-premises IT or private cloud as a suitable alternative for much of their cloud usage. Customers place workloads on the public cloud because it best meets their requirements for those workloads. We recognise that alternative IT models pose some level of constraint on public cloud and some customers may be able to switch their workloads to such environments. However, such switching may not be sufficient to constrain a hypothetical monopolist from raising prices by a small but significant amount. As such, we considered:

³⁸³ This is a quote from the chief commercial officer at a company that helps companies to migrate from the cloud to its colocation data centres. [Are rainy days ahead for cloud computing? - BBC News](#)

- (a) Each of the markets for IaaS and PaaS are separate from traditional on-premises IT; and
- (b) Each of the markets for IaaS and PaaS are separate from private cloud.

Conclusions on the product market definition

3.143 Our conclusions on the product market definition are that:

- (a) the product market for IaaS based on standard compute is separate from IaaS based on accelerated compute infrastructure. When not otherwise specified, we use the term IaaS to refer to IaaS based on standard compute infrastructure. We have specified where we refer to IaaS based on accelerated compute infrastructure;
- (b) there is a market for IaaS services in aggregate, ie we have not segmented IaaS (based on standard compute) into narrower markets;
- (c) the market for IaaS should not be widened to include PaaS;
- (d) we considered PaaS services in aggregate for the purposes of this market investigation;
- (e) the market for PaaS should not be widened to include SaaS;
- (f) IaaS and PaaS are separate from alternative IT models, including traditional on-premises IT and private cloud; and
- (g) we refer to the IaaS and PaaS markets together as the ‘cloud services markets’.

3.144 In summary, for the purposes of this market investigation, we defined the following markets:

- (a) IaaS based on standard compute;
- (b) PaaS.

3.145 Our guidance states that market definition is a ‘useful tool, but not an end in itself, and identifying the relevant market involves an element of judgement. The boundaries of the market do not determine the outcome of our competitive assessment of a market in any mechanistic way’.³⁸⁴

³⁸⁴ CC3 (Revised), paragraph 133.

- 3.146 Our competitive assessment is robust to variations of the precise product market definition used as submitted by parties, meaning that our findings would not change based on the precise boundaries of the relevant product markets.
- 3.147 This is because we have considered competitive constraints coming from inside and outside of the relevant product markets, to the extent relevant, when conducting the competitive assessment. If the market were wider, the additional constraints that would then fall within the market would be relatively weak compared to those we have assessed as in-market constraints. If the market were narrower, the constraints that would then fall outside the market would be relatively strong out-of-market constraints (albeit weaker than the in-market constraints). As a result, any difference in how we define the relevant product markets would not affect our decision.

Geographic market

- 3.148 Our guidelines state that geographic markets can be defined based on the location of either suppliers or customers,³⁸⁵ by considering the degree of substitutability, ie the extent to which suppliers can switch their areas of supply and the extent to which customers in one area may be served in another area.³⁸⁶
- 3.149 The Terms of Reference in this case state the supply of public cloud infrastructure services in the UK.³⁸⁷ As noted at the start of this chapter, the market definition(s) used by the CMA need not always correspond with the relevant market(s) described in the Terms of Reference; specifically, the CMA may conclude that the market definition goes wider or narrower than those goods and services.³⁸⁸

Geographic market for IaaS and PaaS

- 3.150 In this section we considered whether the market for public cloud infrastructure services is national or whether it should be expanded to Europe-wide (ie UK and European Economic Area (EEA)) or global. Evidence is common across IaaS and PaaS, so for the purposes of the geographic market definition, we have considered them together.

Cloud provider views

- 3.151 Cloud providers said that it is not necessarily a requirement to have physical infrastructure based in the UK to compete effectively.³⁸⁹ Microsoft added that a cloud provider can rely on physical infrastructure outside but nearby the UK and

³⁸⁵ CC3 (Revised), paragraph 145.

³⁸⁶ CC3 (Revised), paragraph 147.

³⁸⁷ Terms of Reference, page 2.

³⁸⁸ CC3 (Revised), paragraph 26 and 131.

³⁸⁹ Responses to the CMA's information requests [38].

still compete effectively³⁹⁰ and AWS said that many of its UK customers use infrastructure located outside of the UK.³⁹¹ This is consistent with the capacity shares relative to the revenue shares, which are both set out later in this chapter.

3.152 However, some submissions pointed to the importance of having infrastructure in the UK:

- (a) Oracle said that, while it cannot say for definite, having physical infrastructure in the UK is important because the UK no longer being part of the EU means data sovereignty is even more important.³⁹²
- (b) Microsoft and IBM said that UK-based infrastructure may be important for certain customers such as those with specialised security needs.³⁹³
- (c) AWS said that, while infrastructure does not need to be located near customers, cloud providers may choose to do this to decrease latency or to address customer preferences for data location.³⁹⁴

3.153 AWS and Google said that significant UK-specific investments are not necessary to competitively provide cloud infrastructure services to UK customers.³⁹⁵

However, Microsoft, Oracle and IBM said that a cloud provider would need to make some UK-specific investments:

- (a) Microsoft said that costs would include legal, sales, marketing and administrative roles and that a cloud provider may choose to buy or lease data centres, servers and network equipment;³⁹⁶
- (b) Oracle said that costs would include sales, marketing (though noting marketing could rely on corporate resources) and access to data centre infrastructure;³⁹⁷
- (c) IBM said that there is a multitude of UK-specific investments required to compete effectively, and that it has invested in [REDACTED].³⁹⁸

3.154 While three cloud providers said that, where appropriate, they tailor certain aspects on narrower geographical basis, most cloud providers set their pricing, advertising and marketing strategies globally.³⁹⁹

³⁹⁰ Microsoft's response to the CMA's information request [REDACTED].

³⁹¹ AWS' response to the CMA's information request [REDACTED].

³⁹² Oracle's response to the CMA's information request [REDACTED].

³⁹³ Microsoft's response to the CMA's information request [REDACTED]; IBM's response to the CMA's information request [REDACTED].

³⁹⁴ AWS' response to the CMA's information request [REDACTED].

³⁹⁵ AWS' response to the CMA's information request [REDACTED]; Google's response to the CMA's information request [REDACTED].

³⁹⁶ Microsoft's response to the CMA's information request [REDACTED].

³⁹⁷ Oracle's response to the CMA's information request [REDACTED].

³⁹⁸ IBM's response to the CMA's information request [REDACTED].

³⁹⁹ Responses to the CMA's information requests [REDACTED].

- 3.155 An academic submitted that several studies indicate that the market is in fact global, or at least wider than the region combining UK with Europe.⁴⁰⁰

Evidence from customers

- 3.156 We asked large customers to rate the importance of a list of factors their organisation considers when choosing their main public cloud. In answering this question, customers rated the importance of (a) the number and location of data centres and (b) data sovereignty, alongside other factors. They were asked to rate these factors from one to five, with one being not important at all, and five being very important.
- 3.157 Number and location of data centres and data sovereignty were identified as some of most important factors by customers we spoke to when choosing their main public cloud provider.
- 3.158 In relation to the number and location of data centres, the majority of customers considered this to be important or very important:
- (a) A handful of customers said that it was important that cloud providers were able to cover their operating regions, both in the UK and more broadly.⁴⁰¹
 - (b) A few said that a European presence was important,⁴⁰² and one customer said that its preference is to host personal data in the UK or a country deemed to be 'adequate' by the UK Government, eg EU countries.⁴⁰³
 - (c) A few said that they require data centres in the UK.⁴⁰⁴ One said that this was due to GDPR implications.⁴⁰⁵
- 3.159 In relation to data sovereignty, the majority of customers considered this to be important or very important.
- (a) Most of these customers identified the need for regulatory compliance, data protection and/or security.⁴⁰⁶

⁴⁰⁰ Mr Parisi is former senior fellow of the GW Competition & Innovation Lab at The George Washington University. He cited studies by the Dutch and French Competition Authorities. We note however that neither of these studies defined a global market for public cloud services. Mr Parisi also submitted that the interplay between cloud providers and consumers in the international arena shows that most cloud providers are active globally and supply their services across the globe. R. Parisi, [The Cloud Services Markets' Competitive Landscape: A contribution to the Competition and Markets Authority](#), pages 3, 7 and 9; [Opinion 23-A-08 of 29 June 2023 on competition in the cloud sector](#), page 108.

⁴⁰¹ Responses to the CMA's information requests [REDACTED].

⁴⁰² Responses to the CMA's information requests [REDACTED].

⁴⁰³ [REDACTED] response to the CMA's information request [REDACTED].

⁴⁰⁴ Responses to the CMA's information requests [REDACTED].

⁴⁰⁵ [REDACTED] response to the CMA's information request [REDACTED].

⁴⁰⁶ Responses to the CMA's information requests [REDACTED].

- (b) One customer specified that having a cloud provider that operates in multiple regions of the world means that it can meet data sovereignty requirements within its regions of business operation.⁴⁰⁷

- 3.160 As set out above, we asked large customers to respond on the suitability of a list of public cloud providers as alternatives to their main public cloud provider based on their perception or any direct experience. Customers could also identify and respond on public clouds that we did not list.
- 3.161 Some customers identified Alibaba as a public cloud provider. Alibaba is heavily concentrated in China,⁴⁰⁸ and it was not identified as an effective alternative to UK customers' to the three main cloud providers by customers we spoke to. In particular, Alibaba received a low average rating and was not identified by any customers as a suitable alternative.
- 3.162 Reasons that customers gave for their response included a lack of experience or knowledge of Alibaba's offering,⁴⁰⁹ that using Alibaba would be a supply chain risk or not appropriate (eg for data sovereignty or security reasons),⁴¹⁰ that Alibaba has more limited services or capabilities compared to the three main cloud providers,⁴¹¹ and that they are only considered for demand relating to China.⁴¹²

Our assessment

- 3.163 Overall, the evidence suggests that the markets for IaaS and PaaS are wider than the UK, but not as wide as global for the following reasons.
- 3.164 First, it suggests that the markets are wider than the UK because:
- (a) customers can theoretically choose data centres globally and do choose data centres outside of the UK;
 - (b) some customers identified that having data centres across their operating regions was an important factor when selecting a public cloud provider;
 - (c) the three main cloud providers to UK customers (AWS, Microsoft and Google) are active globally and set their strategies globally; and
 - (d) most cloud providers said that UK infrastructure was not necessary to compete effectively for UK customers.

⁴⁰⁷ [REDACTED] response to the CMA's information request [REDACTED].

⁴⁰⁸ [Alibaba Cloud Global Locations - Deploy Around the World Including Mainland China](#).

⁴⁰⁹ Responses to the CMA's information requests [REDACTED].

⁴¹⁰ Responses to the CMA's information requests [REDACTED].

⁴¹¹ Responses to the CMA's information requests [REDACTED].

⁴¹² Responses to the CMA's information requests [REDACTED].

3.165 Second, some of the evidence suggests that it is not as wide as global because:

- (a) the majority of customers considers the location of data centres and data sovereignty important or very important factors in choosing their main public cloud, with certain customers requiring UK, or sometimes EEA, data centres for regulatory or security purposes;
- (b) cloud providers acknowledge the importance of having physical infrastructure in or near the UK to meet certain security needs or decrease latency; and
- (c) supply chain risks and/or security reasons mean that cloud providers that focus on regions outside of the UK/EEA (eg Alibaba) are generally not considered by customers.

Conclusions on geographic scope

3.166 On the basis of this evidence, we have found that the geographic scope of both the IaaS and PaaS markets is the UK and EEA.

Market structure and concentration

3.167 In this section we considered the structure of the cloud services markets based on our shares of supply analysis. We have calculated shares of supply using various metrics to give an overall picture of the market structures and an indication of how those structures are likely to evolve over time.

Framework for our assessment

3.168 The calculation of market shares of the suppliers of the reference products provides useful background data for the assessment of the levels of firms' market power and may be relevant to our assessment of the theories of harm.⁴¹³ For example:

- (a) High market shares can be a sign that a firm faces weak constraints from rivals. If its market share has been stable or even increasing over time, especially in the face of demand or supply shocks, this could indicate that the firm has market power and that market outcomes are worse than they could be.⁴¹⁴
- (b) Certain practices are more likely to be of concern when a firm has market power. Our guidance sets out how market power may be considered in relation to certain vertical relationships. For example, if a vertically integrated firm has significant market power in an upstream market, it may have an

⁴¹³ CC3 (Revised), paragraph 100 and 101.

⁴¹⁴ CC3 (Revised), paragraph 187.

incentive to refuse access to the input or to raise its prices, and consequently increase the costs of competing downstream firms.⁴¹⁵

- (c) More generally if market concentration is increasing, this may be consistent with certain practices restricting, preventing or distorting competition.

- 3.169 We have considered shares of supply alongside other indicators, such as high barriers to entry and expansion, high profitability and high barriers to switching.⁴¹⁶
- 3.170 We have calculated shares of supply based on three different metrics: (i) shares by revenue; (ii) shares by capacity; and (iii) shares based on flows of new business. Below, we set out our methodology and results for each of these measures of shares of supply in turn.
- 3.171 We present shares of supply for IaaS and PaaS separately, which aligns with our findings on the relevant product markets (see above).
- 3.172 We also present shares of supply for IaaS and PaaS combined (cloud services). This is because categorising cloud services between IaaS and PaaS is not always simple: some services will be on the boundary between two categories and categorisation required some judgement.⁴¹⁷ The combined IaaS and PaaS shares are not sensitive to this allocation.
- 3.173 We attach greater weight to the shares in IaaS and in PaaS separately, since these reflect the relevant product markets. However, we do not rely on the precise point estimates of shares, but rather we take the shares of supply as a good indicator of the relative size of cloud providers.
- 3.174 Shares of supply presented in this chapter include IaaS based on accelerated compute which could not be distinguished across cloud providers' data, but generally we understand that the contribution of accelerated compute to shares of supply is small.
- 3.175 We have gathered data on a UK+EEA, UK and global basis. We have considered where available shares which align with our findings on the geographic market for the cloud services markets, which is UK+EEA. However, UK only data is also informative of cloud providers' success in competing for UK customers and therefore useful in assessing their strength in the markets concerned. Where we present shares based on UK+EEA data below, we present some UK only shares in appendix D.

⁴¹⁵ CC3 (Revised), paragraph 268-270.

⁴¹⁶ We consider barriers to entry in a separate chapter. We consider profitability in Market outcomes below. We consider barriers to switching as part of Prevalence of switching and multi-cloud in Market outcomes below.

⁴¹⁷ See detail in appendix D. We also understand that it may be the case that some providers allocate revenues differently across internal functions, in the instance where a provider's own compute is used to provide a PaaS service. Combined IaaS with PaaS shares also accommodate this factor.

3.176 As we set out above, IaaS is an input into PaaS. This means that any IaaS sales lost to non-vertically integrated PaaS providers are also IaaS sales by (in most cases) one of the three main cloud providers. PaaS services therefore can be thought of as a means for cloud infrastructure providers to ‘distribute’ their infrastructure capacity. The risk of losing sales to PaaS-only providers may provide some competitive constraints on vertically integrated PaaS providers, but PaaS-only providers may be limited in their ability to negotiate the margin they pay to the three main cloud providers for IaaS. IaaS is a core input into PaaS providers’ services, and therefore PaaS-only providers may also be limited in their ability to compete independently of IaaS providers on price to win customers from them.

Shares of supply by revenue

3.177 Shares by revenue are typically the most direct measure of the distribution of customer demand in a market as they take account of differences in the prices and quality of firms’ offerings.⁴¹⁸

Our approach

3.178 Our analysis has focused on how the relative position of cloud providers has evolved as the market has grown since 2020 and how their positions may be expected to change in the near future.

3.179 We have calculated shares of supply by revenue using the annual UK and EEA revenues⁴¹⁹ from AWS, Microsoft, Google, Oracle, IBM, UpCloud, Civo, [X], Hyve, Wasabi, Scaleway, CoreWeave and Centerprise.⁴²⁰ We categorised individual services for some cloud providers using the methodology set out in appendix D: Market structure and concentration methodology. Some of these cloud providers supplied data on a different basis to that which we requested. Further information about the way in which data was supplied is also detailed in appendix D. We do not consider that any limitations in the way cloud providers have supplied data materially affect the results of our analysis set out below, nor that our results would change significantly if some services were categorised differently.

3.180 For the total revenue across all cloud providers (ie the denominator for our shares) we also used data from IDC and Synergy. We placed equal weight on the IDC and Synergy data sources. Some of the firms present in both the IDC and Synergy

⁴¹⁸ CC3 (Revised), Annex A, paragraph 2.

⁴¹⁹ For the purposes of this analysis, we defined UK/EEA revenues as revenues generated from UK/EEA customers and we defined UK/EEA customers as cloud services customers that are operating or trading in the UK/EEA. We defined annual revenues as revenues generated within a calendar year.

⁴²⁰ In the following analysis we refer to UpCloud, Civo, [X], Hyve, Wasabi, Scaleway, CoreWeave and Centerprise as small UK IaaS providers.

data sets were not IaaS and PaaS suppliers in line with the definitions of IaaS and PaaS set out earlier in this chapter, and therefore we excluded these from our calculation of the market size. We used first party revenue submitted by cloud providers when available, as we considered this would be the best estimate of suppliers' UK and EEA revenue. We used data from IDC and Synergy to supplement this revenue, combining the two data sets using different methodologies for IaaS and for PaaS to provide an estimate of the rest of the market. These methodologies are set out in appendix D.⁴²¹

3.181 We consider that shares of supply by revenue for IaaS may overstate the strength of smaller cloud providers (termed as 'other' below) for two reasons:

- (a) The IDC data set includes a relatively sizeable aggregated 'other' category and we do not have visibility of the revenue breakdown by cloud provider within that category. However, the revenue associated with this 'other' category is included as part of the market size as reported. Based on the firms included within the broader data set, this aggregated 'other' category is likely to include firms providing a broad range of services including some that would not be captured by our definition of IaaS.⁴²²
- (b) A few firms provided a single IaaS service such as compute, storage or networking and do not provide all three services. For the purposes of our IaaS market share calculations, we have included all cloud providers in the calculation even if they provide only one of compute, storage or networking. However, these cloud providers are likely to be less close competitors to the three main cloud providers than cloud providers that offer a comprehensive offering consisting of all three services.

UK+EEA shares of supply for IaaS

3.182 UK+EEA shares of supply for IaaS for 2020 – 2024 are presented in the table below. These shares include revenues from IaaS based on accelerated compute which could not be distinguished across all Parties' data, but generally we understand the contribution of accelerated compute to the total revenues to be small.

Table 3.1: UK+EEA shares of supply for IaaS, 2020 – 2024

%

⁴²¹ Since the PDR we received updated historical market size figures from IDC and Synergy for 2020-2024, and an additional provider response. Historical market shares of supply for UK IaaS, PaaS and IaaS + PaaS have been re-calculated using the most up to date data.

⁴²² As outlined in appendix D, we excluded a number of providers through desktop research which we considered did not provide services which aligned with our definition of IaaS as set out above in this chapter from both Synergy and IDC data sets across the years 2020-2024. Therefore, we consider that the 'other' category may also contain revenue from providers which do not align with our definition of IaaS.

	2020	2021	2022	2023	2024
Microsoft	[30-40]%	[30-40]%	[30-40]%	[30-40]%	[30-40]%
AWS	[30-40]%	[30-40]%	[30-40]%	[30-40]%	[30-40]%
Google	[5-10]%	[5-10]%	[5-10]%	[5-10]%	[5-10]%
IBM	[0-5]%	[0-5]%	[0-5]%	[0-5]%	[0-5]%
Oracle	[0-5]%	[0-5]%	[0-5]%	[0-5]%	[0-5]%
Other	[20-30]%	[20-30]%	[10-20]%	[10-20]%	[10-20]%

Source: CMA analysis of first-party revenue data and IDC and Synergy data. Shares may not sum to 100 due to rounding. 'Other' category includes the other firms included in the IDC and Synergy data sets (paragraph 3.180).

3.183 Our analysis shows that:

- (a) Microsoft is the largest provider of IaaS. Its share increased from [30-40]% in 2020 to [30-40]% in 2024;
- (b) AWS is the second largest provider of IaaS. Its share was [30-40]% in 2020, and decreased to [30-40]% in 2024;
- (c) Google is the third largest IaaS provider and has a much lower share of IaaS revenues than AWS and Microsoft: it held a [5-10]% share in 2020 and this has increased to [5-10]% in 2024.
- (d) For IBM and Oracle, shares have remained in the [0-5]% range from 2020 to 2024.⁴²³
- (e) [X], the largest of the other cloud providers for which we have collected first party revenue, had a declining share of supply in UK+EEA IaaS between 2021 and 2024, and held [0-5]% share in 2024. The combined share of all other IaaS suppliers has decreased from [20-30]% in 2020 to [10-20]% in 2024.

3.184 The evidence shows that the IaaS market in the UK+EEA is highly concentrated, with persistently high shares held by two firms.⁴²⁴ Our analysis shows that concentration in IaaS is stable in the sense that the combined share of AWS and Microsoft has remained persistently high and stable from 2020-2024 at around [60-80]%. Over this period, AWS' share has decreased slightly while Microsoft's share has increased.

3.185 The next largest cloud provider is Google, but in 2024 Google remained less than a third of the size of AWS and Microsoft. Google's share has remained stable from 2022-2024.

⁴²³ CMA analysis of first-party revenue data and IDC and Synergy data. Responses to the CMA's information requests issued under section 174 to [X].

⁴²⁴ CC3 (Revised), paragraph 187.

UK+EEA shares of supply for PaaS

3.186 UK+EEA shares of supply for PaaS for 2020 – 2024 are presented in the table below.

Table 3.2: UK+EEA shares of supply for PaaS, 2020 – 2024

	2020	2021	2022	2023	2024
Microsoft	[20-30]%	[20-30]%	[20-30]%	[20-30]%	[20-30]%
AWS	[20-30]%	[20-30]%	[20-30]%	[20-30]%	[10-20]%
Google	[5-10]%	[5-10]%	[5-10]%	[5-10]%	[5-10]%
IBM	[0-5]%	[0-5]%	[0-5]%	[0-5]%	[0-5]%
Oracle	[5-10]%	[5-10]%	[0-5]%	[0-5]%	[0-5]%
Other	[40-50]%	[40-50]%	[40-50]%	[40-50]%	[40-50]%

Source: CMA analysis of first-party revenue data and IDC and Synergy data. Shares may not sum to 100 due to rounding. 'Other' category includes the other firms included in the IDC and Synergy data sets (see paragraph 3.180).

3.187 Our analysis also shows that:

- (a) Microsoft is the largest provider of PaaS and its share has increased slightly from [20-30]% in 2020 to [20-30]% in 2024;
- (b) AWS is the second largest provider of PaaS in this region. Its share has remained relatively stable from 2020-2024, with a share of [20-30]% in 2020 to [10-20]% in 2024.
- (c) Google's share in PaaS has grown slightly from 2020-2024, with a share of [5-10]% in 2020 to [5-10]% in 2024.
- (d) Oracle's share in PaaS has decreased from [5-10]% in 2020 to [0-5]% in 2024.
- (e) IBM's share has remained in the 0-5% range from 2020 to 2024.⁴²⁵
- (f) The combined share of the other smaller cloud providers and ISVs has remained stable at [40-50]% in 2020 and [40-50]% in 2024. This is much higher than for other smaller providers' share in IaaS. Shares in PaaS do not capture the distribution of infrastructure (IaaS) which ISVs use.⁴²⁶

3.188 The PaaS market is less concentrated than the IaaS market and AWS and Microsoft account for a smaller share of PaaS than IaaS at [30-50]% combined. Concentration in PaaS is stable with the combined share of AWS and Microsoft

⁴²⁵ CMA analysis of first-party revenue data and IDC and Synergy data. Responses to the CMA's information requests [X].

⁴²⁶ As we set out above, IaaS is an input into PaaS and we do not know the distribution of the infrastructure which ISVs (which do not self-supply IaaS and use IaaS as an input into their PaaS services) use.

remaining between [30-50]% and [30-50]% between 2020 and 2024. As we note above, IaaS is a core input into PaaS providers' services, and therefore PaaS-only providers may also be limited in their ability to compete independently of IaaS providers on price to win customers from them.

- 3.189 Google has a much smaller share in PaaS than AWS or Microsoft. The combined share of all other smaller cloud providers and ISVs in PaaS has remained stable over time from 2020-2024.

UK+EEA shares of supply for IaaS and PaaS combined

- 3.190 UK+EEA shares of supply for IaaS and PaaS combined for 2020 – 2024 are presented in the table below.⁴²⁷

Table 3.3: UK+EEA shares of supply for IaaS and PaaS combined, 2020 – 2024

	2020	2021	2022	2023	2024
Microsoft	[20-30]%	[30-40]%	[30-40]%	[30-40]%	[30-40]%
AWS	[20-30]%	[20-30]%	[20-30]%	[20-30]%	[20-30]%
Google	[5-10]%	[5-10]%	[5-10]%	[5-10]%	[5-10]%
IBM	[0-5]%	[0-5]%	[0-5]%	[0-5]%	[0-5]%
Oracle	[0-5]%	[0-5]%	[0-5]%	[0-5]%	[0-5]%
Other	[30-40]%	[20-30]%	[20-30]%	[20-30]%	[20-30]%

Source: CMA analysis of first-party revenue data and IDC and Synergy data. Shares may not sum to 100 due to rounding. 'Other' category includes the other firms included in the IDC and Synergy data sets (see paragraph 3.180).

- 3.191 Our analysis shows that for IaaS and PaaS combined:
- (a) Microsoft is the largest provider of IaaS and PaaS combined and its share has increased from [20-30]% in 2020 to [30-40]% in 2024;
 - (b) AWS is the second largest provider of IaaS and PaaS combined and its share has remained broadly stable, albeit declining slightly from 2020-2024: its share was [20-30]% in 2020, decreasing to [20-30]% in 2024;
 - (c) Google is the third largest provider of IaaS and PaaS combined and has a lower share than AWS and Microsoft. Google has increased its share from [5-10]% in 2020, to [5-10]% in 2024; and

⁴²⁷ Similar to shares of supply for IaaS, these shares include revenues from IaaS based on accelerated compute which could not be distinguished across all IaaS and PaaS provider data, but we understand the contribution of accelerated compute to the total revenues to be small.

- (d) for each of IBM and Oracle, shares have remained in the 0-5% range from 2020 to 2024.⁴²⁸

- 3.192 This analysis also shows that concentration in IaaS and PaaS combined is largely stable over time, with the combined share of AWS and Microsoft remaining between [50-70]% and [60-80]% over the last four years. Google remains much smaller than AWS or Microsoft.
- 3.193 The combined share of the other smaller cloud providers and ISVs has remained largely stable, albeit with a slight decrease remaining between [30-40]% and [20-30] % between 2020 and 2024.

UK shares of supply

- 3.194 Below we present a short summary of the results from the UK shares of supply by revenue. As outlined above, we have considered UK+EEA measures where available as this aligns with our geographic market definitions; however, UK shares can still be informative to build an overall picture of the cloud services markets.
- 3.195 These shares are largely consistent with our analysis of shares by revenue based on the UK+EEA region; however, in this region, AWS not Microsoft is the largest supplier of cloud services. Overall, with respect to this geographic region the cloud services markets are slightly more concentrated than at the UK+EEA level. Our analysis shows that:
- (a) For IaaS, UK shares of supply from 2020-2024 show that IaaS remains highly concentrated. Our analysis shows that Microsoft and AWS are the largest cloud providers by revenue with a combined share of [70-90]% in 2024 and that concentration has remained relatively stable over the period.
 - (b) For PaaS, UK shares of supply from 2020-2024 show that PaaS is concentrated but less so than IaaS. Our analysis shows that AWS and Microsoft are the largest single cloud providers by revenue and that shares of supply have remained relatively stable over the period. AWS' and Microsoft's combined share totalled [40-60]% in 2024.
 - (c) For IaaS and PaaS combined, UK shares of supply show that AWS and Microsoft are the largest providers of cloud services in the UK. Our analysis shows that shares of supply have remained relatively stable over the period 2020-2024. AWS' and Microsoft's combined share was [60-80]% in 2024.

⁴²⁸ CMA analysis of first-party revenue data and IDC and Synergy data. Responses to the CMA's information requests [X].

Provider and stakeholder views

- 3.196 In response to our Competitive landscape working paper, Google submitted that the trends in Microsoft's growth over time are significant, and that the CMA's analysis indicates that Microsoft is displaying particularly fast growth.⁴²⁹ Google said Microsoft's market share is growing rapidly and is 'closing the gap on AWS' with the difference in shares falling between 2019 and 2022.⁴³⁰ In response to the PDR, Google submitted that it agrees with the CMA's provisional findings on the structure of, and the cloud providers respective positions in, the markets for cloud infrastructure services.⁴³¹
- 3.197 AWS made two submissions with respect to the market definition used to calculate shares of supply:
- (a) In response to our working paper, AWS said the CMA's emerging views on the level of concentration were flawed because they exclude the constraint from on-premises solutions and exclude SaaS.⁴³²
 - (b) Additionally in response to the PDR, AWS submitted that analysis of market shares considering IaaS and PaaS in isolation, as distinct from each other and from rival IT providers (including on premises and private cloud) is divorced from how competition works in the cloud infrastructure services space.⁴³³
- 3.198 We conducted the shares of supply analysis in line with our market definition considering the share of supply for each of IaaS and PaaS.⁴³⁴ As some products or services may not fit neatly into the definitions of IaaS or PaaS, we have also presented shares of supply based on IaaS with PaaS, which are not sensitive to the classification of products and services between IaaS and PaaS.
- 3.199 In response to the Competitive Landscape working paper, AWS cited data showing that AWS' share in IaaS, PaaS and IaaS with PaaS has decreased from 2019 to 2023.⁴³⁵ As described above, our analysis also illustrates that AWS' share of supply for IaaS, PaaS and IaaS and PaaS combined all declined slightly from 2022 to 2023; however, its share remains persistently high. We place more weight on our own analysis of shares of supply (for IaaS, for PaaS, and for IaaS and PaaS combined) because we have used first party revenue data where available

⁴²⁹ [Google Cloud's response to the CMA's Competitive landscape working paper dated 23 May 2024](#), paragraph 8.

⁴³⁰ [Google Cloud's response to the CMA's Competitive landscape working paper dated 23 May 2024](#), paragraph 8a.

⁴³¹ [Google's response to the CMA's PDR dated 28 January 2025](#), paragraph 7.

⁴³² [AWS' economic response to the CMA's Competitive landscape working paper dated 02 August 2024](#), paragraph 28.

AWS has additionally made this argument in [§].

⁴³³ [AWS' response to the CMA's PDR dated 28 January 2025](#), paragraph 10.

⁴³⁴ See Market definition section above.

⁴³⁵ [AWS' economic response to the CMA's Competitive landscape working paper dated 02 August 2024](#), paragraph 32.

and used two third-party sources of data, from IDC and Synergy, which we refined based on the definitions of IaaS and PaaS we set out earlier in this chapter.⁴³⁶

- 3.200 AWS also submitted an analysis of UK concentration using the Herfindahl–Hirschman index (HHI).⁴³⁷ We have also calculated UK HHIs from 2020-2023 based on first party, IDC and Synergy data.
- 3.201 Overall, we note the limitations of HHI as a measure of concentration, in this instance in part due to the aggregation of some firms within the third-party data which is used as an input in both our and AWS calculations, which can inflate the HHI. Therefore, as we present above, we consider shares of supply to be our primary indicator of market concentration.
- 3.202 Our guidance states that we are likely to regard any market with an HHI in excess of 2,000 as highly concentrated, and any market with an HHI in excess of 1,000 as concentrated.⁴³⁸
- 3.203 AWS said its analysis showed concentration had decreased in segments considered by our working paper ie, IaaS, PaaS and IaaS and PaaS combined, between 2019 and 2023.⁴³⁹ We note that the HHI for cloud services provided by AWS indicates it is likely to be at least concentrated, and that the decline shown in the HHI calculated by AWS was only for 2023 and we do not interpret this as a trend, particularly in light of five previous years of overall stability.⁴⁴⁰
- 3.204 We have also calculated UK HHIs based on first party, IDC and Synergy data, which give the following results:
- (a) IaaS: HHI has been greater than 3,000 for 2020 to 2023;
 - (b) PaaS: HHI has been greater than 1,900 for 2020 to 2023; and
 - (c) IaaS and PaaS combined: HHI has been greater than 2,000 for 2020 to 2023.⁴⁴¹
- 3.205 Appendix D describes in more detail some reasons for the differences in our and AWS' HHI estimates. We consider our estimate of HHI for IaaS to have the fewest limitations. However, given the caveats associated with the data used to calculate

⁴³⁶ See appendix D.

⁴³⁷ The HHI potentially reflects both the number of firms in the industry and their relative size. It is defined as the sum of the squares of all the market shares in the market, and thus gives proportionately greater weight to the larger market shares. [CC3 \(Revised\)](#), Annex A, paragraph 7.

⁴³⁸ [CC3 \(Revised\)](#), Annex A, paragraph 7.

⁴³⁹ [AWS' economic response to the CMA's Competitive landscape working paper dated 02 August 2024](#), paragraph 31.

⁴⁴⁰ AWS' submission to the CMA [§].

⁴⁴¹ Due to the aggregation of the 'other' segments in the IDC and Synergy data used in this analysis we note that these HHIs are an overestimate. CMA analysis of first-party revenue data and IDC and Synergy data. Refer appendix D.

HHI, particularly in PaaS, we place limited weight on both our own and AWS' calculations.

- 3.206 In response to the PDR, AWS submitted that over 30% per year overall growth in cloud services means that even if a particular cloud provider's relative share were to decrease or remain static, there is still an upside for them and their investors to invest in a growing industry. It also submitted that this means that there are significant opportunities for new cloud providers to enter and existing cloud providers to expand.⁴⁴² AWS has also submitted that the CMA should take a much broader and dynamic view when assessing the market structure, taking into account long-term trends, investments, entry, and the growth rates of entrants.⁴⁴³
- 3.207 As we outline above, there is growth in cloud services. However, we consider that even in growing IaaS and PaaS markets, the barriers to entry and expansion that we identify in chapter 5, such as high sunk cost investments, economies of scale and the benefits of having product portfolios across IaaS and PaaS, will persist.
- 3.208 Our analysis of shares of supply illustrates that smaller cloud providers have not been increasing their share significantly over recent years in IaaS, despite market growth. Oracle, although gaining share in IaaS in the UK+EEA, has a declining share in PaaS, and it remains small relative to AWS, Microsoft and even Google. Google itself remains small in comparison to AWS and Microsoft. We have also considered below whether smaller cloud providers are winning an outsize proportion of revenue growth, which could provide a signal if smaller cloud providers were beginning to expand. However, as set out below, smaller firms are winning only a moderate share of overall revenue growth.
- 3.209 AWS submitted that since 2006, its shares of supply have steadily dropped as competitors continue to enter and grow in this space. AWS submitted that, by the PDR's analysis, AWS is growing at a slower pace than the reported percentage growth rate of the 'cloud services sector', and that its shares decreased between 2020 and 2023. It submitted that companies like Google and Oracle were growing more rapidly, in percentage terms, than AWS and have gained ground, with an increase of shares over the same period. AWS submitted that this contradicts the CMA's finding that the market structure is and will remain static.⁴⁴⁴
- 3.210 In addition, Microsoft submitted that recent financial announcements from Google since the publication of the PDR show an inconsistency in placing Google in a second tier. It submitted that the market shares cited by the PDR provide an

⁴⁴² [AWS' response to the CMA's PDR dated 28 January 2025](#), paragraph 12.

⁴⁴³ AWS' submission to the CMA [REDACTED].

⁴⁴⁴ [AWS' response to the CMA's PDR dated 28 January 2025](#), paragraph 13.

incomplete picture of Google's competitive position, which it submitted was strong.⁴⁴⁵

- 3.211 In this section we present a range of measures of share of supply to assess cloud providers' competitive position from 2020-2024 in the UK, UK+EEA and globally. Our analysis illustrates that the shares of Microsoft and AWS have remained relatively stable over recent years. Combined with this, other cloud providers including Google and Oracle are winning either small shares each year, or shares show no overall increasing trend, which suggests these cloud providers are struggling to win share and constrain larger rivals. We have not seen any clear and stable trends in forecast capacity shares and shares of flow that would indicate that current shares of supply are likely to change significantly. However, we do not consider shares of supply in isolation. As we outline above, we have considered shares of supply alongside other indicators of competition, such as closeness of competition between cloud providers, high barriers to entry and expansion, high profitability and high barriers to switching.⁴⁴⁶
- 3.212 In response to the Competitive landscape working paper, two academics said that they considered the cloud services market to be competitive:
- (a) One said on a European basis the cloud services market is very deconcentrated and provided market shares from Statista.⁴⁴⁷ These suggested lower shares for Microsoft and AWS in a European market for cloud services compared to the results of our analysis. For the reasons mentioned above, we consider our analysis as more reliable than analysis from Statista.
 - (b) Another said the CMA's evidence is consistent with the hypothesis that the computer storage and processing power market is competitive, including evidence that the market shares of the second and third firms have increased over three years.⁴⁴⁸
- 3.213 In response to the PDR, one of the same academics submitted, without providing supporting evidence, that:

⁴⁴⁵ [Microsoft's response to the CMA's PDR dated 28 January 2025](#), paragraphs 84-87.

⁴⁴⁶ We consider barriers to entry in a separate chapter. We consider profitability in Market outcomes below. We consider barriers to switching as part of Prevalence of switching and multi-cloud in Market outcomes below. See also chapter 9.

⁴⁴⁷ [X] submission to the CMA [X].

⁴⁴⁸ Dr George R Barker, [Comment on The UK Competitive Market Authority's \(CMAs\) Cloud Services Market Investigation Three Working Papers on The Supply of Public Cloud Infrastructure Services In the UK Covering The CMA's 1. Competitive Landscape Working Paper; 2. Egress Fee Working Paper; and 3. Committed Spend Agreements Working Paper](#), page 50. We note that Dr Barker is a member of the Oxford Cross Disciplinary Machine Learning Research Cluster (OXML), which is supported by Microsoft. (see page 1).

- (a) high market shares may also reflect an efficient market with high economies of scale and scope in production and consumption, and high market shares should be treated as a strong indicator of an efficient market;
- (b) various pieces of evidence in the PDR cannot be used to refute the legal presumption that the market is competitive; and
- (c) constant and changing shares are consistent with competition, and various pieces of evidence in the PDR were consistent with the market being competitive.⁴⁴⁹

3.214 We disagree that high market shares should be treated as a strong indicator of an efficient market and note that in general a highly concentrated market might be an indicator that one or more firms hold unilateral market power.⁴⁵⁰ In assessing whether competition in a market is working effectively, we have looked at a range of evidence including shares of supply in the relevant markets.

3.215 We also note that an industry report submitted in response to our request for information states that cloud services are highly consolidated, with just two primary vendors.⁴⁵¹

Shares of supply by capacity

3.216 Shares by capacity look at the market structure from the supply side. These show the relative strength of each cloud provider in terms of their production capability and can be an appropriate measure of market shares for homogenous products where the capacity is efficient enough to make expansion profitable in response to a price rise or to reduced output by competitors.⁴⁵² Absent barriers to competition or switching, firms with greater capacity have a greater ability and incentive to compete for business and thereby exert a competitive constraint on rivals. Shares by capacity may also reflect individual cloud providers' expectations relating to their growth. Where barriers to entry are high, firms with lower shares by capacity may be more likely to become constrained or have a lesser ability or incentive to compete for new demand than firms with greater capacity.

3.217 Shares by capacity may also provide a useful sense-check for estimated shares of supply of the service that relies on the same capacity. In cloud services, a company's share of total data centre capacity may serve as a proxy for its share of revenue from the sale of infrastructure directly to customers as IaaS, indirectly as part of PaaS solutions, or to ISVs.

⁴⁴⁹ [Dr George R Barker's response to the CMA's PDR dated 28 January 2025](#), pages 58-62.

⁴⁵⁰ [CC3 \(Revised\)](#), paragraph 187.

⁴⁵¹ [\[redacted\]](#) response to the CMA's information request [\[redacted\]](#).

⁴⁵² [CC3 \(Revised\)](#), Annex A, paragraph 2.

- 3.218 However, we note that cloud providers may vary in the extent to which their capacity is utilised. A less-effective competitor may have a larger share of capacity but have more spare capacity and a smaller revenue share. For this reason, we consider that capacity shares are most informative when considered alongside shares of revenue.
- 3.219 We present shares by capacity for the calendar years 2020-2026. The shares for 2020-2024 are based on realised capacity while the 2025-2026 shares are based on the cloud providers' own forecasts (including any planned or ongoing expansion projects). Estimates of shares of data centre capacity should be treated with caution as any significant announcements on new capacity since these estimates were collected may have an impact on those shares.

Our approach

- 3.220 We calculated the shares of supply by capacity using data from AWS, Microsoft, Google, IBM, Oracle, CoreWeave and some smaller IaaS providers that serve UK customers⁴⁵³ on their data centre capacity in megawatts (MW) within UK+EEA,⁴⁵⁴ globally, and in the UK. We consider that UK+EEA shares give the best reflection of the market structure with respect to UK customers.⁴⁵⁵
- 3.221 These shares do not include the capacity of other small IaaS providers. As such, each cloud provider's share is likely a small overestimate across all cloud providers and should be interpreted as an indicator of relative share between the cloud providers included in the calculation.
- 3.222 As noted above, the unit we have used for data centre capacity is megawatts (MW). Our analysis in chapter 5 indicates that [REDACTED].⁴⁵⁶
- 3.223 We note that these shares include revenues from IaaS based on accelerated compute which could not be distinguished across all Parties' data. Generally, we understand the contribution of accelerated compute to data centre capacity to be small.

UK + EEA shares of supply by data centre capacity

- 3.224 UK+EEA shares of supply by data centre capacity for 2020–2026 are presented in the table below.

⁴⁵³ Centreprise, [REDACTED], Scaleway and Wasabi.

⁴⁵⁴ For the purpose of this analysis, we have allocated providers' capacities to UK+EEA if they were classified by the provider as relating to Europe. Notes about the data provided by the providers are set out in appendix D.

⁴⁵⁵ See our discussion of the geographic market for IaaS and PaaS in the section Market Definition earlier in this chapter.

⁴⁵⁶ See our discussion of cloud providers' operating efficiencies in chapter 5, in the section 'Economies of scale'.

Table 3.4: UK + EEA shares of supply by data centre capacity, 2020 – 2026

	%						
	2020	2021	2022	2023	2024	2025F	2026F
Microsoft	[40-50]%	[40-50]%	[40-50]%	[40-50]%	[40-50]%	[>50]%	[>50]%
AWS	[40-50]%	[30-40]%	[30-40]%	[30-40]%	[20-30]%	[>50]%	[>50]%
Google	[10-20]%	[10-20]%	[10-20]%	[10-20]%	[10-20]%	[>50]%	[>50]%
Oracle	[0-5]%	[0-5]%	[0-5]%	[0-5]%	[0-5]%	[>50]%	[>50]%
CoreWeave	[0-5]%	[0-5]%	[0-5]%	[0-5]%	[0-5]%	[>50]%	[>50]%
IBM	[0-5]%	[0-5]%	[0-5]%	[0-5]%	[0-5]%	[>50]%	[>50]%
Other	[0-5]%	[0-5]%	[0-5]%	[0-5]%	[0-5]%	[>50]%	[>50]%

Source: CMA analysis of data from AWS, Microsoft, Google, IBM, Oracle, CoreWeave and some smaller IaaS providers. Shares may not sum to 100 due to rounding. Data from 2025 – 2026 are forecasts. 'Other' category includes Centerprise, [redacted], Scaleway and Wasabi.

3.225 Our analysis shows that:

- (a) Microsoft has become and is forecast to remain the largest cloud provider by capacity in the UK+EEA region: Microsoft's share was [40-50]% in 2020 and has increased to [40-50]% in 2024. It is forecast to increase its capacity, to [redacted] by 2026.
- (b) AWS is the second largest cloud provider by capacity in the UK+EEA region: AWS' share was [40-50]% in 2020 and has decreased to [20-30]% in 2024. It is forecast to decrease further to [redacted] by 2026.
- (c) Google is the third largest cloud provider by capacity and is significantly smaller than Microsoft and AWS: Google's share was [10-20]% in 2020 and has increased to [10-20]% in 2024. It is forecast to remain around [redacted] in 2026.
- (d) Smaller cloud providers do not have a similar scale. Oracle, CoreWeave and IBM's share of capacity has been [0-5]% each.
- (e) Microsoft and AWS' combined UK+EEA shares of supply by capacity was [redacted]% in 2024, and are forecast to increase to [redacted]% by 2026.⁴⁵⁷

3.226 UK+EEA shares of supply by capacity are consistent with other evidence relating to market structure in cloud services. By this measure, AWS and Microsoft are the largest cloud providers and Microsoft's share has been increasing relative to AWS' share. In contrast to revenue shares, Microsoft has been the largest cloud provider in the UK+EEA based on capacity since 2021. Smaller cloud providers have a small share of capacity, which is not forecast to change meaningfully.

⁴⁵⁷ CMA analysis of data centre capacity in megawatts submitted cloud providers. Responses to the CMA's information requests [redacted]. Responses to the CMA's information request [redacted].

Global shares of supply by data centre capacity

3.227 Global shares of supply by data centre capacity for 2020 – 2026 are presented in the table below.

Table 3.5: Global shares of supply by data centre capacity, 2020 – 2026

	2020	2021	2022	2023	2024	2025F	2026F
AWS	[40-50]%	[40-50]%	[40-50]%	[40-50]%	[30-40]%	[30-40]%	[30-40]%
Microsoft	[30-40]%	[30-40]%	[30-40]%	[30-40]%	[30-40]%	[30-40]%	[30-40]%
Google	[10-20]%	[10-20]%	[10-20]%	[10-20]%	[10-20]%	[10-20]%	[10-20]%
Oracle	[0-5]%	[0-5]%	[0-5]%	[5-10]%	[5-10]%	[5-10]%	[5-10]%
CoreWeave	[0-5]%	[0-5]%	[0-5]%	[0-5]%	[0-5]%	[0-5]%	[0-5]%
IBM	[0-5]%	[0-5]%	[0-5]%	[0-5]%	[0-5]%	[0-5]%	[0-5]%
Other	[0-5]%	[0-5]%	[0-5]%	[0-5]%	[0-5]%	[0-5]%	[0-5]%

Source: CMA analysis of data from AWS, Microsoft, Google, IBM, Oracle, CoreWeave and some smaller IaaS providers. Shares may not sum to 100 due to rounding. Data from 2025 – 2026 are forecasts. 'Other' category includes Centerprise, Scaleway, and Wasabi.

3.228 Our analysis shows that:

- AWS is the largest global cloud provider in terms of capacity, although its share has declined: AWS' share was [40-50]% in 2020 and has decreased to [30-40]% in 2024. It is forecast to increase to [30-40]% by 2026.
- Microsoft is the second largest in terms of capacity with a relatively stable share: Microsoft's share was [30-40]% in 2020 and has decreased slightly to [30-40]% in 2024. It is forecast to remain at [30-40]% by 2026.
- Google is the third largest in terms of capacity: Google's share was [10-20]% in 2020 and increased to [10-20]% in 2024. It is forecast to decrease slightly to [10-20]% by 2026.
- Oracle's share went from [0-5]% in 2020 to [5-10]% in 2024.
- Other smaller cloud providers do not have a similar scale. CoreWeave's share was [0-5]% in 2024, and IBM's share was [0-5]% in 2024.
- Microsoft and AWS' combined global shares of supply by capacity was [60-80]% in 2024, and it is forecast to increase to [60-80]% by 2026.

3.229 Our analysis of global shares of supply by capacity indicates that AWS and Microsoft have the largest capacity, followed by Google and then by Oracle, CoreWeave and IBM. Based on our analysis of capacity forecasts to 2026, the

positions of AWS, Microsoft and Google are not forecast to change in the next few years.⁴⁵⁸

UK shares of supply by data centre capacity

- 3.230 Shares of supply by capacity show that Microsoft is by far the largest-capacity cloud provider in the UK, followed by AWS.⁴⁵⁹
- 3.231 We have found that the geographic scope of the markets for both IaaS and PaaS is the UK+EEA, and that most cloud providers said that UK infrastructure was not necessary to compete effectively for UK customers.⁴⁶⁰ Therefore, we place less weight on these shares.

Cloud provider views

- 3.232 In response to our Competitive landscape working paper, AWS provided data on its share of capacity based on three different third party sources for data centre power capacities as the denominator, and said its share of supply is below [10-20]% both globally and in the UK.⁴⁶¹
- 3.233 The denominators that AWS used for the global figures were sourced from third party research firms. We place more weight on our analysis of shares of data centre capacity because they are based on first-party data rather than third-party research, noting the caveat as described above that this means our figures are slight overestimates. As described above, we place less weight on UK shares by capacity.

Shares of supply by flows of new business

- 3.234 In markets where customers face switching costs, shares of supply based on the 'stock' of customers (eg shares by 'installed base' or share of revenues, where revenues tend to be 'recurring' and have relatively low churn) may not reflect recent changes in the relative competitive position of suppliers.⁴⁶² In such cases, it can be useful to consider evidence on shares of supply on a 'flow' basis (eg shares of new customers or new revenues).

⁴⁵⁸ CMA analysis of data centre capacity in megawatts submitted by cloud providers. Responses to the CMA's information requests [§]. Responses to the CMA's information request [§].

⁴⁵⁹ Based on the inconsistency between these shares and the shares of supply by revenue, as well as other evidence set out in Market definition, we do not consider that these shares give an accurate representation of the structure of the market. We consider this limitation does not apply to UK shares of supply by revenue as this metric considers UK customers, which may source cloud services from outside the UK. CMA analysis of data centre capacity in megawatts submitted by the cloud providers.

⁴⁶⁰ See our assessment of the geographic market for IaaS and PaaS, in the 'Market Definition' section above.

⁴⁶¹ AWS' submission to the CMA [§].

⁴⁶² We have found that full switching between providers is extremely rare in cloud services markets. See the section 'Prevalence of switching and multi-cloud' for more detail.

- 3.235 In this section we present the following shares based on the flow of new business:
- (a) Shares by overall revenue growth. This involves calculating the year-on-year revenue growth of each cloud provider, and expressing that growth as a proportion of the total revenue growth of all of the cloud providers combined (calculated for both IaaS, and IaaS and PaaS combined).
 - (b) Shares by revenue from newly acquired customers. This involves calculating the revenues that each cloud provider earned from customers acquired in each year as a proportion of the total revenues from new customers earned by all of the cloud providers combined.
 - (c) Shares by new customers acquired. This involves calculating the number of customers that each cloud provider acquired as a proportion of the total customers acquired by all of the cloud providers in each year.
- 3.236 We consider that the analysis gives a useful indication of the relative importance of revenue growth from existing and new customers and which cloud providers are gaining new business at a faster pace. However, the granularity of the data we have gathered for shares of revenue from new customers and shares of new customers means we cannot distinguish between the following:
- (a) If the new customers that a cloud provider acquires are: (i) customers completely new to the cloud (representing competition for customers); (ii) customers that are only new to that provider and placing a new workload (representing competition for new workloads); or (iii) customers that are only new to that provider and switching an existing workload (representing competition for existing workloads).
 - (b) If changes in a cloud provider's revenue from existing customers is caused by: (i) some existing customers decreasing/increasing their spend on existing workloads without switching (eg cost optimisation, business expansion); (ii) some existing customers switching existing workloads to or from another cloud provider (representing competition for existing workloads); or (iii) some existing customers placing new workloads with that provider.
- 3.237 Based on the limitations of these measures outlined above, we consider shares by overall revenue growth the most informative of cloud providers' success in competing for new business and consider shares by new customers acquired to be the least informative.

Our approach

- 3.238 We calculated the various shares by business flows using data from AWS, Microsoft, Google, Oracle, IBM and some smaller IaaS providers that serve UK

customers,⁴⁶³ on the number of new UK customers they acquired in each of the calendar years 2020-2023, the revenues generated from those new customers, and their overall annual UK revenues. We also used IaaS and PaaS data procured from IDC and Synergy (as described above). We have the necessary data to present shares by new business flows for the years 2021 to 2024.

3.239 There are some caveats to the analysis based on flows of newly acquired customers:

- (a) The definition of a new customer is binary and is also based on an arbitrary threshold. A new customer was generally defined as a customer that spent more than \$100 for the first time in a year (in that cloud provider's revenue data). Therefore, a customer that spends \$200 and then increases to \$10 million in a year would be classified as an 'existing' customer in our analysis when it is more akin to a new customer, given its relatively limited previous engagement with cloud infrastructure services.⁴⁶⁴ We note that, due to the definition of a newly acquired customer, these measures may be particularly volatile; therefore, as we outline above, we considered these alongside other measures to build a complete picture of the cloud services markets.
- (b) In relation to the new customer data, each cloud provider submitted data on a slightly different basis.⁴⁶⁵
- (c) The analysis of the number of new customers does not account for the size of customers, so all are given the same weighting. This could produce a large share for a cloud provider winning a large number of very small customers.
- (d) We only have this data for UK customers, and this does not align with our geographic market definition of UK+EEA.⁴⁶⁶

3.240 We first set out shares by year-on-year revenue growth (for both IaaS, and IaaS and PaaS combined). We then consider shares based on newly acquired customers and revenue from newly acquired customers.

Shares of year-on-year IaaS revenue growth

3.241 For each firm we calculated its annual growth in IaaS revenues, then divided this by the sum of annual revenue growth across all cloud providers to give the

⁴⁶³ Centerprise, Civo, [REDACTED], Hyve, Wasabi and UpCloud. Shares of supply by newly acquired customers and by revenue from newly acquired customers are calculated for Microsoft, AWS, Google, Oracle, IBM and these selected smaller IaaS providers only and do not include all cloud providers active in the UK.

⁴⁶⁴ We tested the sensitivity of this assumption by requesting new customer data from Microsoft, AWS and Google based on different spend thresholds (\$500, \$1000 and \$10,000). See appendix D for further explanation.

⁴⁶⁵ See appendix D. We do not consider this to have a material effect on the analysis.

⁴⁶⁶ We have considered where available shares which align with our findings on the geographic market for the cloud services markets, which is UK+EEA. However, UK only data is also informative of providers' success in competing for UK customers and therefore useful in assessing their strength in the markets concerned, as UK customers may source cloud services from outside the UK.

share.⁴⁶⁷ We set out shares based on UK year-on-year IaaS revenue growth in appendix D, with a summary of the results below.

UK+EEA shares of year-on-year IaaS revenue growth

3.242 UK+EEA shares of year-on-year IaaS revenue growth for 2021 to 2024 are presented in the table below.

Table 3.6: UK+EEA shares of year-on-year IaaS revenue growth, 2021 – 2024

	%			
	2021	2022	2023	2024
Microsoft	[40-50]%	[40-50]%	[40-50]%	[30-40]%
AWS	[30-40]%	[30-40]%	[20-30]%	[30-40]%
Google	[5-10]%	[10-20]%	[10-20]%	[5-10]%
IBM	[><]%	[><]%	[0-5]%	[><]%
Oracle	[0-5]%	[0-5]%	[0-5]%	[0-5]%
Other	[5-10]%	[10-20]%	[10-20]%	[10-20]%

Source: CMA analysis of data from AWS, Microsoft, Google, IBM, Oracle and some smaller IaaS providers. Shares may not sum to 100 due to rounding. 'Other' category includes the other firms included in the IDC and Synergy data sets (see paragraph 3.180).

3.243 Our analysis shows that:

- (a) Microsoft has consistently held the largest share of revenue growth: it won [40-50]% of overall new revenues in 2021 and [30-40]% of overall new revenues in 2024;
- (b) AWS' share of revenue growth has been relatively stable between 2021 and 2024: it won [30-40]% and [30-40]% of overall new revenues in 2021 and 2024, respectively;
- (c) Google's share of overall revenue growth has increased from 2020 to 2023 from [5-10]% in 2021 to [10-20]% in 2023, decreasing to [5-10]% in 2024; and
- (d) For IBM and Oracle, shares of growth have remained below [0-5]%. Oracle's share of overall revenue growth has been growing over this period.⁴⁶⁸

3.244 Shares of year-on-year revenue growth in IaaS show that AWS and Microsoft have the largest shares of overall revenue growth. Google has the third largest share of overall revenue growth, consistent with its position in the market.

⁴⁶⁷ Similar to the shares of supply for IaaS, the underlying data include revenues from IaaS based on accelerated compute which could not be distinguished across all IaaS provider data, but generally we understand the contribution of accelerated compute to the total revenues to be small.

⁴⁶⁸ CMA analysis of first-party revenue data and IDC and Synergy data. Responses to the CMA's information requests [X].

UK shares of year-on-year IaaS revenue growth

- 3.245 Shares of year-on-year revenue growth in IaaS in the UK show that AWS and Microsoft have consistently held the largest shares of overall revenue growth. AWS' and Microsoft's individual shares of overall year-on-year IaaS revenue growth have been above [30-40]% from 2021-2024. Google has the third largest share of overall revenue growth, consistent with its position in the market. Its share has ranged from [5-10]% to [10-20]% over the period.

Shares of year-on-year IaaS and PaaS revenue growth

UK + EEA shares of year-on-year IaaS and PaaS revenue growth

- 3.246 For each firm we calculated its annual growth in IaaS plus PaaS revenues, then divided this by the sum of annual revenue growth across all providers to give the share.
- 3.247 UK+EEA shares of year-on-year IaaS and PaaS revenue growth for 2021–2024 are presented in the table below.

Table 3.7: UK+EEA shares of year-on-year IaaS and PaaS revenue growth, 2021 – 2024

	%			
	2021	2022	2023	2024
Microsoft	[30-40]%	[30-40]%	[30-40]%	[30-40]%
AWS	[20-30]%	[20-30]%	[20-30]%	[20-30]%
Google	[5-10]%	[10-20]%	[5-10]%	[5-10]%
IBM	[><]%	[><]%	[0-5]%	[><] %
Oracle	[0-5]%	[0-5]%	[0-5]%	[0-5]%
Other	[20-30]%	[20-30]%	[20-30]%	[30-40]%

Source: CMA analysis of data from AWS, Microsoft, Google, IBM, Oracle and some smaller IaaS providers. Shares may not sum to 100 due to rounding. 'Other' category includes the other firms included in the IDC and Synergy data sets (see paragraph 3.180).

- 3.248 Our analysis shows that:
- (a) Microsoft's share of overall revenue growth has fallen from [30-40]% in 2021 to [30-40]% in 2024;
 - (b) AWS' share of overall revenue growth fell over 2021 to 2024: it won [20-30]% and [20-30]% of overall new revenues in 2021 and 2024, respectively;
 - (c) Google's share of overall growth has decreased from 2021 to 2024: it won [10-20]% of overall new revenues in 2021 and [5-10]% of overall new revenues in 2024;

- (d) For IBM and Oracle, shares of overall revenue growth have remained below [0-5]%; Oracle's share of overall revenue growth has been from 2021-2024; and
- (e) Share of revenue growth from all other providers increased from [20-30]% in 2021 to [30-40]% in 2024. This is across an estimate of other providers based on more than 300 other providers.⁴⁶⁹

3.249 Shares of year-on-year revenue growth in IaaS and PaaS combined show that AWS and Microsoft have consistently held the largest shares of overall revenue growth from 2021-2024. Google has the third largest share of overall revenue growth, consistent with its position in the market.

UK shares of year-on-year IaaS and PaaS revenue growth

- 3.250 Shares of year-on-year revenue growth in IaaS and PaaS combined in the UK show that Microsoft and AWS have consistently held the largest shares of overall revenue growth. Microsoft's and AWS' individual shares of overall year-on-year IaaS revenue growth have ranged from [20-30]% to [30-40]% from 2021-2024.
- 3.251 Google has the third largest share of overall revenue growth, consistent with its position in the market. Its share has ranged from [5-10]% to [10-20]% over the period. Other providers have a larger share of revenue growth in IaaS and PaaS combined than the IaaS market, and this has been growing over the period 2021-2024. The share of other providers has ranged from [10-20]% to [20-30]%.

Shares of supply by revenue from newly acquired UK cloud services customers

- 3.252 For each cloud provider, we requested the public cloud services revenues from new UK customers in each year, then divided this by the sum of annual revenue growth across public cloud services revenues from all providers' new UK customers to give the share. For the purposes of this analysis, we have defined a new UK customer as one that spends at least \$100 in a year for the first time.⁴⁷⁰
- 3.253 When interpreting shares of supply by revenue from new customers, we note that these measures may be particularly volatile and sensitive to one-off events which can create temporary spikes. As such, we consider the general trend of these series over four years to be more informative than the shares in any particular year.

⁴⁶⁹ CMA analysis of first-party revenue data and IDC and Synergy data. Responses to the CMA's information requests issued under section 174 to [X].

⁴⁷⁰ See appendix D for details.

3.254 UK shares of supply by revenue from newly acquired customers for 2021-2024 are presented in the table below.

Table 3.8: UK shares of supply by revenue from newly acquired customers, 2021-2024

	%			
	2021	2022	2023	2024
Microsoft	[40-50]%	[60-70]%	[50-60]%	[30-40]%
Google	[5-10]%	[5-10]%	[10-20]%	[5-10]%
AWS	[10-20]%	[10-20]%	[10-20]%	[0-5]%
Oracle	[0-5]%	[10-20]%	[5-10]%	[30-40]%
CoreWeave	[0-5]%	[0-5]%	[0-5]%	[5-10]%
IBM	[10-20]%	[0-5]%	[0-5]%	[0-5]%
Other	[5-10]%	[5-10]%	[5-10]%	[0-5]%

Source: CMA analysis of data from AWS, Microsoft, Google, IBM, Oracle and some smaller IaaS providers. Shares may not sum to 100 due to rounding. 'Other' category includes Centerprise, Civo, [REDACTED], Hyve, Scaleway, Wasabi and UpCloud.

3.255 Our analysis shows that:

- (a) Microsoft's share of revenue from new customers increased, from [40-50]% in 2021 to [50-60]% in 2023, then fell to [30-40]% in 2024;
- (b) AWS' share of revenue from new customers decreased from [10-20]% in 2021 to [10-20]% in 2023, then fell further to [0-5]% in 2024;
- (c) Google's share of revenue from new customers has increased slightly from [5-10]% in 2021 to [10-20]% in 2023, then fell to [5-10]% in 2024;
- (d) Oracle's shares ranged between [0-5]% and [10-20]% in 2021 to 2023, then rose significantly to [30-40]% in 2024. CoreWeave's shares were [0-5]% in 2021 to 2023, then rose to [5-10]% in 2024. IBM's shares ranged between [0-5]% and [10-20]% in 2021 to 2023, and was [0-5]% in 2024.⁴⁷¹

3.256 Shares of supply by revenue from new customers show that Microsoft is winning significantly more completely new business than other providers. Relative to its overall position in the market, AWS' share of revenue from new customers is low.

3.257 In 2024, the shares of Microsoft, AWS and Google each fell, while the shares of Oracle and CoreWeave grew. We do not consider this to be indicative of a substantial weakening in the position of the main cloud providers. As noted above, the general trend of these series over four years is more informative than the shares in any particular year. In particular, we note the following:

⁴⁷¹ CMA analysis of revenue and customer acquisition data from Microsoft, AWS, Google, Oracle, CoreWeave and IBM. Responses to the CMA's information requests [REDACTED].

- (a) Oracle generated a substantially larger share of revenue from new customers in 2024. This was the result of a one-off event which will only increase its share of revenue from new customers to that extent in 2024.⁴⁷²
- (b) CoreWeave also generated a larger share of revenue from new customers in 2024. We are aware that this revenue is associated with the acquisition of [REDACTED], namely [REDACTED].⁴⁷³

3.258 We considered that our findings on the shares of supply by revenue from new customers may be sensitive to the specific definition of a new customer used, so we conducted a sensitivity analysis. We gathered data from Microsoft, AWS and Google in 2023 and 2024, and additionally from Oracle in 2024, on revenues from new customers as the spend threshold for defining a new customer varied. In addition to the \$100 threshold presented above, we collected data for thresholds of \$500, \$1,000, and \$10,000.

3.259 Our sensitivity analysis showed that these findings were robust to changes in the definition of a new customer.⁴⁷⁴ In 2023 and 2024, Microsoft continued to have a higher share of revenue from new customers than Google or AWS, at each threshold for defining a new customer (\$100, \$500, \$1,000, and \$10,000), though Microsoft's share was smaller in the higher threshold scenarios.

Shares of supply by newly acquired UK cloud services customers

3.260 We have also calculated shares of supply by the number of newly acquired UK customers.⁴⁷⁵ This analysis shows that Microsoft is winning UK customers at a significantly higher rate than AWS and Google.

3.261 However, we view this measure as less informative than shares based on revenue from newly acquired customers, as it does not consider the relative size of each new customer. We also note that, as with shares by revenue from new UK customers, these measures may be particularly volatile and sensitive to one-off events which can create temporary spikes.

Cloud provider views

3.262 In response to our Competitive landscape working paper, Microsoft provided an analysis of cloud provider year-on-year revenue growth, based on quarterly

⁴⁷² Oracle submitted that the reason for the large increase of revenue in 2024 from new customers is associated with the acquisition of Cerner customer contracts in FY24. [REDACTED].

⁴⁷³ CoreWeave submitted that [REDACTED] the revenue generated by CoreWeave from the UK in 2024 relates to revenue generated from [REDACTED] which became a UK customer of CoreWeave in 2024. CoreWeave's response to the CMA's information requests [REDACTED].

⁴⁷⁴ Refer to appendix D.

⁴⁷⁵ Refer to appendix D.

financial data.⁴⁷⁶ Microsoft said this showed Oracle has experienced double-digit revenue growth since Q4 2021, and IBM had double-digit growth between Q1 2020 and Q3 2022 (after which the revenue data was no longer available).

- 3.263 We note this is different to the results presented above (Shares of supply by year-on-year UK IaaS revenue growth and Shares of supply by year-on-year UK IaaS and PaaS revenue growth) which reports each provider's share of overall year-on-year revenue growth, rather than each provider's year-on-year growth rate. We put more weight on the shares of supply by year-on-year revenue growth we have calculated because this is based on first party revenue data.
- 3.264 Microsoft submitted that instead of focusing on a single share of supply metric, a better indicator of future relative market position would be UK cloud providers' share of global capex. On this basis, it submitted that the 2022 shares are: AWS at 41% (\$52.7 billion), Microsoft at 28% (\$35.2 billion), Google at 25% (\$32.3 billion), Oracle at 5% (\$6.6 billion) and IBM with 1% (\$1.2 billion).⁴⁷⁷
- 3.265 We consider the weight we can place on this analysis is limited for the following reasons: this data relates to a global market definition rather than UK+EEA; capex in any one year is not representative of a cloud provider's total asset base (a cloud provider who has been active in the market for longer will have a larger total asset base); and capex may be related to investment in graphics processing units for a cloud provider's AI business rather than only data centres for the supply of IaaS. In addition, cloud providers may not have the same impact on competition from the same capex.
- 3.266 In response to the PDR, Microsoft submitted that that the PDR shares on IaaS revenue growth and IaaS/PaaS revenue growth show that GCP is growing while Microsoft's share is falling, and AWS' shares remain stable. Microsoft submitted that that this indicates that Google is growing faster than AWS and Microsoft and that it will be a larger competitor going forward.⁴⁷⁸
- 3.267 Our shares of supply analysis illustrates recent growth of Google over the period 2020-2024. However, Google's share in IaaS and in PaaS in the UK and UK+EEA remains small. By revenue, its share is less than [10-20]%, and its share of overall revenue growth does not suggest that this will increase significantly (ie its share of overall revenue growth is not consistently and significantly larger than its shares of supply by revenue). In fact, considering shares of supply by revenue, especially from 2022-2024 in the UK+EEA, Google's share has been stable.

⁴⁷⁶ [Microsoft's response to the Competitive landscape, Committed spend agreements and Egress fees](#) working papers dated 23 May 2024, paragraph 20.

⁴⁷⁷ [Microsoft's response to the Competitive landscape, Committed spend agreements and Egress fees](#) working papers dated 23 May 2024, paragraphs 34-37.

⁴⁷⁸ [Microsoft's response to the CMA's PDR dated 28 January 2025](#), paragraph 81.

- 3.268 Microsoft has held a consistently high share of revenue growth, with its share of overall revenue growth in the UK and UK+EEA largely greater than 30% across 2021-2024.
- 3.269 As supported by our analysis of historical market shares by revenue, capacity and shares of flow, smaller shares of other cloud providers suggest they are struggling to win share and constrain larger rivals. Our analysis illustrates that the shares of Microsoft and AWS have remained relatively stable over recent years. Combined with this, other cloud providers including Google and Oracle are winning either small shares each year, or shares show no overall increasing trend, which suggests these providers are struggling to win share and constrain larger rivals.
- 3.270 Microsoft submitted that shares of supply by revenue flows as presented in the PDR likely understate Google's growth over the period as its lower profitability in the past is due in part to its own internal accounting issues.⁴⁷⁹
- 3.271 We do not consider Google's internal accounting system to have an impact on shares of overall revenue growth or share of revenues/number of new UK customers, as these changes impacted its profitability based on public accounts, not its submission to us.
- 3.272 In response to the PDR, Microsoft submitted that shares of supply likely overstate Microsoft's growth, as its shares of new customers is inconsistent with the fact that new customers who joined Microsoft between January and December 2023 only generated [REDACTED] of Azure revenues in their first year of spend.⁴⁸⁰
- 3.273 As we outline above, we consider analysis of new customers alongside other measures of share of supply to provide a useful indication of the relative importance of revenue growth from new and existing customers and which cloud providers are gaining new business at a faster rate. It may be the case that new customers are relatively small (as outlined above); therefore, a high share of number of new customers may not translate into a high share of revenue from new customers, and we have interpreted the shares in this light.

Conclusions on market shares and concentration

- 3.274 We have found that the IaaS market is highly concentrated and the PaaS market is concentrated. Forward-looking metrics suggest that these levels of concentration are likely to endure.
- 3.275 Microsoft is now the largest provider in both IaaS and PaaS markets in the UK+EEA and its share is increasing. AWS is the second largest provider in both

⁴⁷⁹ [Microsoft's response to the CMA's PDR dated 28 January 2025](#), paragraph 82.

⁴⁸⁰ Microsoft's submission to the CMA [REDACTED].

markets. Google's share in each is much lower than either Microsoft or AWS and while it has grown over the period 2020-2024, it has been more stable in recent years. Other providers are much smaller, particularly in IaaS.

- 3.276 We place most weight on shares which map to the product and geographic markets we have defined: these are separate product markets for IaaS and PaaS in the UK+EEA, rather than the shares which consider IaaS and PaaS combined. If we changed the boundaries we have defined for IaaS and PaaS slightly, this is unlikely to change our overall findings of concentration because we have found that the markets for IaaS and PaaS are highly concentrated and concentrated respectively.
- 3.277 Although the IaaS market is growing overall and this may provide opportunities for smaller providers to gain share as there is still new business to be won, the relatively small shares of other providers and lack of growth in these shares over the period we have considered, indicates that these providers are likely to remain considerably smaller and are unlikely to provide a strong constraint on their rivals.
- 3.278 The PaaS market remains concentrated but much less so than IaaS, with a longer tail of smaller competitors. As outlined above, most PaaS providers in the market are not vertically integrated with IaaS. Therefore, any sale they make may also be an IaaS sale for one of Microsoft, AWS or Google, as IaaS is an input to PaaS. This weakens the constraint which PaaS-only providers exert on vertically integrated firms supplying both IaaS and PaaS. In addition, concentration in IaaS (as an input to other cloud services) can be a cause for concern on its own.
- 3.279 Forward-looking metrics suggest that these market positions are likely to endure:
- (a) Shares of supply by capacity show that AWS and Microsoft are the largest providers and Microsoft has been the largest provider in the UK+EEA based on capacity since 2021.
 - (b) Shares of year-on-year revenue growth in (i) IaaS and (ii) IaaS and PaaS combined both show that Microsoft and AWS have the largest single provider shares of overall revenue growth; and shares of supply by revenue from new customers show that Microsoft is winning significantly more new customers than AWS and Google.
 - (c) While Google's market share has increased over the period, it remains significantly smaller than that of Microsoft or AWS, whose market shares remain high and stable.

- (d) We did not see any clear and stable trends in the measures of share by flow that would indicate that current shares of supply are likely to change significantly.⁴⁸¹

Market outcomes

Introduction

- 3.280 Our guidance sets out that, in a market investigation, the CMA will normally consider outcomes of the competitive process such as profitability, prices, levels of innovation, product range and quality. This is because outcomes of the competitive process can provide evidence about the functioning of a market.⁴⁸² Evaluating these outcomes helps the CMA determine whether there is an adverse effect on competition and, if so, the extent to which customers may be harmed by it, that is the degree and nature of ‘customer detriment’.⁴⁸³
- 3.281 In this chapter, we set out our analysis in relation to the following market outcomes in the supply of public cloud infrastructure services: profitability, prices, quality and innovation, multi-cloud and switching.

How we use evidence on market outcomes

- 3.282 Profitability and prices are among the more observable and measurable outcomes of the competitive process, but should be treated as indicators that may be useful within the context of our overall assessment of the market rather than as features of the market for the purpose of the AEC test.⁴⁸⁴ Other less quantifiable factors, such as quality and innovation, are also important to customers.⁴⁸⁵
- 3.283 Microsoft and AWS submitted that evidence of market outcomes show that the cloud market is working well, and that the CMA should provide clear benchmarks against which to compare current market outcomes.⁴⁸⁶ For example, AWS submitted that the Competitive landscape working paper ‘is content with only pointing to the hypothetical possibility that innovation could be even stronger in a hypothetical counterfactual’.⁴⁸⁷
- 3.284 It is not feasible to determine exactly at what level and based on what measures, for example, the quality or extent of innovation would be too low. And even for

⁴⁸¹ We also consider future trends which may impact competition in cloud services in chapter 4: Impact of AI.

⁴⁸² [CC3 \(Revised\)](#), paragraph 103.

⁴⁸³ [CC3 \(Revised\)](#), paragraph 103.

⁴⁸⁴ [CC3 \(Revised\)](#), paragraphs 104 and 126.

⁴⁸⁵ [CC3 \(Revised\)](#), paragraph 127.

⁴⁸⁶ [AWS’ response to the CMA’s updated issues statement and working papers](#), dated 25 June 2024, paragraphs 2, 6 and 23; see further [AWS’ Economic response to the CMA’s competitive landscape working paper](#), paragraphs 108-111. [Microsoft’s response to the CMA’s competitive landscape, committed spend agreements and Egress fees working papers](#), paragraphs 12 and 16.

⁴⁸⁷ [AWS’ Economic response to the CMA’s Competitive landscape working paper](#), paragraphs 108-109.

more measurable outcomes, our guidelines are clear that they should be treated as indicators in the broader context of the overall assessment of the market rather than as precise yardsticks. This means, for example, and as discussed below, that pricing trends can be interpreted in different ways,⁴⁸⁸ depending on the broader market context.⁴⁸⁹

3.285 Despite our view that it would not be appropriate to specify quantitative benchmarks, several cloud providers also suggested comparative benchmarks, none of which we think are appropriate for this market:⁴⁹⁰

- (a) **Comparison with on-premises:** Some cloud providers have submitted that the cloud market produces better outcomes than traditional IT services.⁴⁹¹ We are clear that this is not a helpful comparator. We are not comparing cloud technology with what came before, but rather whether it is more likely than not that the cloud infrastructure services market would be more competitive, securing better outcomes for customers, absent the features we identify.
- (b) **Comparison with the past:** Microsoft and [X] have submitted evidence which they submit shows continuing improvement in relation to particular market outcomes, for example that prices have fallen over time or that numbers of patents have increased.⁴⁹² However, using the past as a benchmark in this way is only instructive to a degree because of the risk that, over time, market outcomes may have been affected by events other than competition including, for example, advancements in technologies outside the relevant markets.

3.286 In the rest of this section, we considered:

- (a) profitability;
- (b) pricing trends;
- (c) quality and innovation;

⁴⁸⁸ Our guidelines note that there may be several factors affecting prices and that we take this into account when considering inferences from this type of analysis. (See [CC3 \(Revised\)](#), paragraph 108.)

⁴⁸⁹ Where there are well-established guidelines to help us interpret evidence on market outcomes, as is the case in relation to profitability, we can use these as tools to aid us in our interpretation, again while bearing in mind the broader market context.

⁴⁹⁰ In addition to the comparators suggested below, our guidelines suggest that, in relation to price, comparisons with other markets such as those for similar products in other countries or for comparable products in the UK, can sometimes be helpful. ([CC3 \(Revised\)](#), paragraph 113.) However, we do not consider that to be the case here, as there do not appear to be other UK markets sharing the distinctive aspects of cloud services and we consider that the UK market for cloud services is likely to operate in a similar way to other geographies (which means the features whose effects we are assessing in the UK are also likely to be present elsewhere, hampering the utility of such other markets as a comparator).

⁴⁹¹ See Market definition section

⁴⁹² [X] submission to the CMA [X]; Microsoft's submission to CMA [X].

- (d) switching; and
- (e) multi-cloud.

Profitability

Role of profitability analysis

3.287 The aim of profitability analysis is to improve our understanding of competitive conditions within a market, by examining the outcomes of that market in terms of the financial performance of the participating firms. Our guidance states that:

‘Firms in a competitive market would generally earn no more than a ‘normal’ rate of profit – the minimum level of profits required to keep the factors of production in their current use in the long run, ie the rate of return on capital employed for a particular business activity would be equal to the opportunity cost of capital for that activity.’⁴⁹³

3.288 The purpose of conducting profitability analysis, therefore, is to understand whether the levels of profitability (and therefore prices) achieved by cloud providers are consistent with the levels we might expect in a competitive market.

3.289 We do not regard profitability that exceeds the cost of capital at a point in time in itself to be a problematic feature of a market. However, a situation where the profitability of firms representing a substantial part of the market has exceeded the cost of capital over a sustained period could be an indication of limitations in the competitive process.⁴⁹⁴

3.290 Our guidance sets out that profitability tends to be a more observable and measurable outcome than outcomes like quality and innovation. An analysis of profitability may be useful in quantifying the extent and nature of competition by examining the outcomes of that market in terms of the financial performance of the participating firms.⁴⁹⁵ Profitability can provide a more holistic approach to assessing outcomes as it takes into account not only prices, but also other factors such as quality and innovation and firms’ relevant costs. However, we have considered the evidence on all outcomes and do not rely on our analysis of profitability alone to reach conclusions on the functioning of the cloud markets and/or market power of cloud providers.

3.291 In this section we set out the role of profitability analysis and assess the profitability of cloud providers. We have considered whether profits among certain

⁴⁹³ CC3 (Revised), paragraph 116.

⁴⁹⁴ CC3 (Revised), paragraphs 117, 118 and 126.

⁴⁹⁵ See CC3 (Revised) paragraphs 104, 114 and 116.

cloud providers reflect a 'normal' rate of return based on the nature of competition in the supply of public cloud infrastructure services.⁴⁹⁶

Stakeholder views on profitability analysis

3.292 We summarise below the submissions received in response to the PDR⁴⁹⁷ and our working paper⁴⁹⁸ on our approach to assessing the profitability of cloud providers. We consider these submissions further in this section, where we set out our approach to assessing profitability and our analysis. More detailed submissions on the measurement of capital employed and calculation of the cost of capital for our analysis are set out and considered in appendix E.

Scope and conceptual approach

3.293 Microsoft submitted that profitability analysis that compares the return on capital employed (ROCE) and weighted average cost of capital (WACC) is well-suited to mature markets such as energy, water and telecoms but unsuitable for the cloud services market where:

- (a) operational excellence and scale requires significant investments multiple years before reaching positive returns (Amazon: 2015; Microsoft: 2016, Google: 2023), and
- (b) it has not reached an equilibrium steady-state but is in a capex spending 'race'.⁴⁹⁹

3.294 AWS and Microsoft did not agree with our view, as expressed in the PDR, that the cloud services markets have been in a steady state over the last five years for the purposes of analysing profitability:

- (a) AWS submitted that the cloud market is not in a steady state because firms keep investing in cloud services, compute capacity, and related IT infrastructure; innovation keeps occurring as a result of such investments; and the cloud industry continues to grow at a double-digit rate.⁵⁰⁰
- (b) Microsoft submitted that the assumption of steady state profitability does not capture the competitive conditions of Azure becoming profitable only in

⁴⁹⁶ We provide additional details on our analysis in appendix E.

⁴⁹⁷ See chapter 3, paragraphs 3.218 to 3.276 of [Provisional decision report](#) and [appendix E: Profitability analysis](#).

⁴⁹⁸ See chapter 6 of [Competitive landscape working paper](#).

⁴⁹⁹ [Microsoft's response to the Competitive landscape, Committed spend agreements and Egress fees working papers dated 1 July 2024](#), paragraph 32.

⁵⁰⁰ AWS' submission to the CMA on quality and innovation investments and profitability dated 16 April 2025, paragraphs 59-60.

FY2020, and the rapid pace of technological change in cloud computing in recent quarters with the developments around AI.⁵⁰¹

- 3.295 Two cloud providers submitted that gross margins are not meaningful or sufficiently informative metrics for comparing profitability between different providers.⁵⁰² One of the cloud providers said that this was due to the different cost structures and the different ways costs are attributed and the assumptions used under different providers' accounting methodologies.⁵⁰³
- 3.296 An academic submitted that there were issues with use of global profitability to assess the competitive landscape in the UK.⁵⁰⁴
- 3.297 We received the following submissions on the relevance of profitability to assessing competition:
- (a) AWS submitted that there is no conclusive relationship between profit margins and effective competition because several factors other than market power determine profit margins, including innovation, product differentiation and higher efficiency.⁵⁰⁵
 - (b) AWS submitted that there is no empirical link between cloud service providers' profitability and their size/growth and/or their share of cloud spend, thus there is no systematic evidence that larger cloud service providers are more profitable.⁵⁰⁶
 - (c) AWS also submitted that our working paper did not define the differential between a cloud provider's ROCE and WACC that would indicate an abnormal level of profits and was of the view that ROCE levels exceeding the WACC do not represent evidence against the market being workably competitive.⁵⁰⁷
 - (d) Microsoft submitted that it is a 'good thing' for competition (and in turn customers), not a 'bad thing', that each of Google, Microsoft and Amazon are generating positive returns which can justify the kinds of capital investment they are making for the future.⁵⁰⁸
 - (e) In response to the PDR, AWS and the International Center for Law and Economics said that profitability alone cannot indicate whether competition in

⁵⁰¹ [Microsoft's response to the CMA's PDR dated 28 January 2025](#), paragraph 71.

⁵⁰² Submissions to the CMA [REDACTED].

⁵⁰³ [REDACTED] submission to the CMA [REDACTED].

⁵⁰⁴ R. Parisi, [The Cloud Services Markets' Competitive Landscape: A contribution to the Competition and Markets Authority](#), page 13.

⁵⁰⁵ AWS' submission to the CMA [REDACTED].

⁵⁰⁶ AWS' submission to the CMA [REDACTED]; AWS' submission to the CMA [REDACTED].

⁵⁰⁷ AWS' submission to the CMA [REDACTED].

⁵⁰⁸ [Microsoft's response to the Competitive landscape, Committed spend agreements and Egress fees working papers](#), paragraph 32.

a sector is effective or be proof of market power. AWS submitted that significant levels of profitability are expected even in the presence of fierce competition in a dynamic industry, where profits are uncertain, and where firms compete through innovation, and that these features apply to the cloud sector.⁵⁰⁹ The International Center for Law and Economics submitted that a high ROCE versus WACC could be observed even in a competitive environment, especially in technology markets, due to recovery of sunk costs, innovation and transient advantage, and contestable market pressure.⁵¹⁰

- 3.298 AWS submitted that comparator analyses with other cloud providers shows that AWS' returns and margins are in line with those of its competitors, including smaller providers. AWS said this is fully consistent with a competitive market in which there are firms of different sizes competing and innovating.⁵¹¹
- 3.299 AWS submitted analysis comparing its ROCE and the differential between its ROCE and WACC to other cloud providers and other IT services providers (such as HP, Tencent and Cisco) and said that this shows that AWS' economic profits are not abnormal relative to others in the industry and, thus, are not evidence of market power. AWS submitted that if the PDR's hypothesis were true, it should be observed that only AWS and Microsoft have persistent positive economic profits at much higher levels than rivals and its analysis confirms the alternative, that economic profits result from innovation consistent with dynamic competition.⁵¹²

Margins and ROCE

- 3.300 AWS submitted that its gross margins are in line with those of other cloud service providers for the period 2021-2023.⁵¹³
- 3.301 In response to our working paper, AWS submitted that ROCE estimates for AWS showed a downward trend, and that with the inclusion of estimates for brand value, goodwill and IPR&D intangible assets, AWS' ROCE ranges between [REDACTED] during 2013-2023. [REDACTED].⁵¹⁴
- 3.302 Subsequently, in response to the PDR, AWS submitted updated ROCE analysis that includes an estimate for intangible assets and ranges between [REDACTED] for 2018 to 2024. AWS also said that the PDR underestimates AWS' WACC, and with

⁵⁰⁹ AWS' response to the CMA's PDR dated 28 January 2025, paragraph 19.

⁵¹⁰ International Center for Law and Economics' response to the CMA's PDR dated 28 January 2025, pages 7-8. The International Center for Law and Economics declared that its funding includes donations from companies with an interest in the investigation.

⁵¹¹ AWS' response to the CMA's PDR dated 28 January 2025, paragraph 19.

⁵¹² AWS' submission to the CMA [REDACTED].

⁵¹³ AWS' submission to the CMA [REDACTED].

⁵¹⁴ AWS' submission to the CMA [REDACTED].

corrections applied AWS' WACC estimate increases to [REDACTED].⁵¹⁵ We consider AWS' more detailed submissions on the calculation of ROCE and WACC in appendix E.

Innovation

3.303 We received the following submissions on the relevance of innovation to interpreting the profitability analysis:

- (a) AWS submitted that the profitability analysis in our Competitive landscape working paper ignores that the cloud segment is characterised by a high degree of innovation consistent with a dynamically competitive environment and that therefore there are pro-competitive reasons for positive margins that benefit consumers.⁵¹⁶
- (b) Two academics and the International Center for Law and Economics also submitted that sustained higher profitability deriving from innovation and greater efficiency (in the past or in ongoing waves) is not evidence that the cloud market is not competitive or that high returns are the result of exploiting consumers.⁵¹⁷ One of the academics characterised the nature of the cloud market as new, risky, uncertain, innovative technology or way of doing business with high demand for investment and noted that in the long run, the profitability or rate of return falls in markets.⁵¹⁸ The International Center for Law and Economics submitted that treating cloud firms' profitability like mature utilities' returns fails to reflect dynamic competition, as the cloud market is driven by innovation and continual change and profits are sustained by delivering superior service in a competitive race.⁵¹⁹

3.304 AWS submitted that, according to economic theory, in conditions of continuous innovation in an industry, economic profits can be persistently positive, ie ROCE can be persistently above WACC. AWS submitted that this implies that sustained economic profits alone do not show that there are limitations in the competitive process, and that sustained economic profits would be expected regardless of whether some cloud providers have market power or not.⁵²⁰

⁵¹⁵ AWS' submission to the CMA [REDACTED].

⁵¹⁶ AWS' submission to the CMA [REDACTED].

⁵¹⁷ R. Parisi, [The Cloud Services Markets' Competitive Landscape: A contribution to the Competition and Markets Authority](#), page 13; Dr George R Barker, [Comment on The UK Competitive Market Authority's \(CMAs\) Cloud Services Market Investigation Three Working Papers on The Supply of Public Cloud Infrastructure Services In the UK Covering The CMA's 1. Competitive Landscape Working Paper; 2. Egress Fee Working Paper; and 3. Committed Spend Agreements Working Paper](#), pages 65-67. International Center for Law and Economics's response to the CMA's PDR dated 28 January 2025, pages 8-9.

⁵¹⁸ Dr George R Barker, [Comment on The UK Competitive Market Authority's \(CMAs\) Cloud Services Market Investigation Three Working Papers on The Supply of Public Cloud Infrastructure Services In the UK Covering The CMA's 1. Competitive Landscape Working Paper; 2. Egress Fee Working Paper; and 3. Committed Spend Agreements Working Paper](#), pages 65-67; Dr George R Barker's response to the CMA's PDR dated 28 January 2025, pages 90-91.

⁵¹⁹ International Center for Law and Economics's response to the CMA's PDR dated 28 January 2025, pages 8-9.

⁵²⁰ AWS' submission to the CMA on quality and innovation investments and profitability dated 16 April 2025, paragraph 62.

- 3.305 AWS submitted that in an industry where there is continuous innovation as a result of continuous investments, economic returns can persist over sustained periods of time, even as long as or longer than 7-10 years, and that continuous innovation is a defining characteristic of the cloud industry.⁵²¹
- 3.306 AWS submitted that ROCE fluctuates according to how successful firms have been at innovating at different points in time to competitively offer services that customers value.⁵²²

Investment

- 3.307 In our working paper,⁵²³ we noted that increasing levels of investment in cloud infrastructure which are likely to be aimed at supporting the development of AI services has contributed to flat or slightly falling ROCE in the most recent years. In response we received the following submissions:
- (a) AWS submitted that the presence of significant investments is consistent with effective competition.⁵²⁴
 - (b) Microsoft submitted that it is incorrect to say falling or flat ROCE are the result of 'increased levels of investment in cloud infrastructure ... [to] [support] the development of AI services' which is 'not ... a result of competitive forces'. Microsoft submitted that it is wrong that its 'capex race' to serve AI has nothing to do with 'competitive forces'.⁵²⁵
- 3.308 In response to the PDR, Microsoft submitted that by treating recent profitability trends as indicative of long-term stability, this misses companies' ongoing investment cycles. Microsoft submitted that its AI-driven capex on GPUs is expected to increase from \$2.7 billion in FY2023 to [REDACTED] in FY2027, a [REDACTED] compound annual growth rate, and that this indicates that cloud services are entering a new investment phase rather than operating in a steady state.⁵²⁶
- 3.309 AWS submitted that the PDR was wrong in asserting that much of AWS' and Microsoft's initial investment has occurred prior to going to the market. AWS cited that its investments in 2022-2024 remain significant in comparison to AWS' revenue and that Microsoft's R&D expenditure and CapEx have accelerated since Azure first launched in 2010.⁵²⁷

⁵²¹ AWS' submission to the CMA [REDACTED].

⁵²² AWS' submission to the CMA [REDACTED].

⁵²³ [Competitive landscape working paper](#), paragraphs 6.41 and 6.42

⁵²⁴ AWS' submission to the CMA [REDACTED].

⁵²⁵ [Microsoft's response to the Competitive landscape, Committed spend agreements and Egress fees working papers](#), paragraph 33.

⁵²⁶ [Microsoft's response to the CMA's PDR](#) dated 28 January 2025, paragraph 73.

⁵²⁷ AWS' submission to the CMA [REDACTED].

- 3.310 AWS submitted that the PDR notes that AWS' and Microsoft's AI-related capex investments are responsible for the recent decline in ROCE, while at the same time predicting that AWS' and Microsoft's ROCE will stabilise around current levels in the future, and that this means the PDR essentially assumes that AWS' and Microsoft's recent investments will be successful.⁵²⁸

Our analysis of profitability

Our approach

- 3.311 We have used a ROCE vs WACC profitability analysis in the cloud market. Whilst Microsoft has submitted that operational excellence and scale in the cloud market requires significant investment several years ahead of reaching positive returns, we do not consider this to be unique to the cloud market. Investment ahead of returns is reflected in the capital employed used in calculating ROCE and the analysis of profitability over a sufficiently long period.
- 3.312 Where large and risky investments have been made, or the industry has experienced a period of growth, we would expect to see a normal level of profitability restored over a relatively long timescale.⁵²⁹ Microsoft and AWS each entered the cloud markets over ten years ago, and much of their initial investment⁵³⁰ in cloud services would have occurred prior to going to market (because much of the infrastructure being used to deliver cloud services was already being used by Microsoft's and AWS' other businesses). Both providers have been investing and operating in the markets for a sufficiently long period of time for their current (ie last five years) profitability to be reflective of market conditions and any trend of ROCC-WACC convergence to be observable, absent any significant market changes.⁵³¹
- 3.313 A lack of 'steady state' in the markets might be relevant if we saw significant fluctuation in ROCE or profit margins, however our analysis finds that margins for Microsoft's and AWS' cloud businesses have been broadly stable or slightly growing, and even with the recent increase in capex spending (discussed further below) Microsoft's and AWS' cloud businesses have consistently achieved ROCE in excess of WACC. We recognise that recent investments in infrastructure to support AI may have impacted Microsoft's and AWS' recent profitability, since Microsoft's and AWS' investments in infrastructure to support AI (as well as revenues and profit contributions from AI cloud services) are included in the

⁵²⁸ AWS' submission to the CMA [X].

⁵²⁹ CC3 (Revised), paragraph 121 and Annex A, paragraph 10.

⁵³⁰ Note: initial investment refers to the investment required to first launch cloud services to customers in the market. We do not consider this to include subsequent ongoing investment in maintaining and/or growing cloud services.

⁵³¹ Note, our assessment here is specific to the state of AWS' and Microsoft's profitability in the cloud market. This should not be taken to mean that the UK cloud market as a whole is in a steady state. We also consider the impact of innovation and investments in AI on our assessment of ROCE later on in this section.

financial reporting of their cloud businesses. We consider the implications of Microsoft and AWS reporting investments, revenues and profits for their cloud businesses inclusive of AI in our analysis below, but do not consider this to reflect profit instability in non-AI cloud services for Microsoft and AWS.

- 3.314 We have also analysed gross margins and earnings before interest and tax (EBIT) margins for cloud providers as indicators of financial performance. We note that two cloud providers consider gross margins to be uninformative metrics for comparing profitability between different providers.⁵³² We benchmark margins as it provides useful context and insight into the comparative profitability of cloud providers, as well as trends in profitability over time. However, we do not typically consider margins by themselves to be a suitable benchmark for determining if profitability exceeds a 'normal' level, if we can determine other measures such as ROCE, and note that margin benchmarking often has some limitations.
- 3.315 In order to develop our profitability assessment, we have reviewed the relevant revenues, costs and capital base of the main public cloud infrastructure services providers operating in the UK.⁵³³ We examine the profitability of cloud services for the three main cloud providers in the UK (Microsoft, AWS and Google) as well as financial information on smaller cloud providers who have been identified as global and UK competitors and where we have been able to obtain cloud services profit margin figures.⁵³⁴
- 3.316 We examined the global profitability of cloud providers in our analysis. Whilst we have defined the geographic scope of the cloud markets for this investigation as UK and EEA, and one stakeholder submitted that the use of global data to assess profitability for the UK cloud market was incorrect,⁵³⁵ for the purposes of assessing profitability we considered it more suitable to assess the global profitability of cloud providers due to (i) the global nature of cloud providers' businesses and the cloud services they provide, and (ii) the global nature of their financial reporting, asset base and capital investment. There is also limited financial data available at a UK or UK and EEA level which limits our ability to do meaningful UK and EEA-level profitability assessment.
- 3.317 We have analysed the ROCE for Microsoft's Azure and Cloud & Enterprise business segments and AWS, as these are the two largest cloud providers in the

⁵³² Submissions to the CMA [X].

⁵³³ Our profitability analysis includes revenues, costs and capital base from IaaS based on accelerated compute which is not separately identifiable in the provider data.

⁵³⁴ We have analysed AWS, Microsoft, Google, Oracle and IBM as we have identified these firms as relevant (and largest, based on cloud services revenues) providers of public cloud infrastructure services in the UK. We also analysed OVHcloud as it publicly reports its public cloud business performance and is Europe-focused, and we consider it to be a reasonable proxy for a mid-sized competitor operating in the UK.

⁵³⁵ R. Parisi, [The Cloud Services Markets' Competitive Landscape: A contribution to the Competition and Markets Authority](#), page 15.

UK and represent [60-80]% of the supply together.⁵³⁶ We have compared the ROCE for these businesses to our estimate for the WACC to assess their profitability.⁵³⁷

- 3.318 We do not include Google in our ROCE analysis as it has only recently (since 2023) reported profits,⁵³⁸ and it has a significantly lower share of cloud services revenues compared to Microsoft and AWS.⁵³⁹ Similarly we do not include other smaller cloud providers which have lower market shares in our ROCE. We considered the profitability trends of Google and other cloud providers in our margin benchmarking for comparative purposes, but as the profitability of firms representing a substantial part of the markets is a useful indicator of competitive conditions in the cloud markets, we are primarily interested in assessing the profitability of the largest incumbent cloud providers in the markets.⁵⁴⁰ We compare the ROCE for Microsoft's and AWS' cloud businesses to the WACC, in order to assess the extent to which these cloud providers earn a 'normal' rate of profit. As noted above, where firms persistently earn substantially in excess of a normal return, this signals that there may be limitations in the competitive process.⁵⁴¹
- 3.319 As set out above, the relevant test for our profitability analysis is whether the profitability of firms representing a substantial part of the market has exceeded the cost of capital over a sustained period.⁵⁴² If firms persistently earned profits in excess of the return required to compensate investors for the risks taken, we would expect these profits to attract entry and/or expansion, which would serve to compete away profits in excess of the cost of capital up until the point where firms cover their total costs, including a market-based cost of capital, and no more. We therefore expect, in the long run equilibrium, the largest firms in a competitive market to have a rate of return on capital employed equal to the opportunity cost of capital,⁵⁴³ and the presence of other firms in the market similarly achieving profits in excess of the cost of capital does not prove that the market is competitive.

⁵³⁶ This is the share of the combined UK IaaS and PaaS markets, of which AWS and Microsoft each had [30%-40]% market share based on revenue in 2024. See appendix D for details.

⁵³⁷ The rationale for benchmarking return on capital with the opportunity cost of capital is that in a competitive market, if firms persistently earned in excess of the return required to compensate investors for the risks taken, we would expect these profits to attract entry and/or expansion. This entry/expansion would serve to compete away profits in excess of the cost of capital up until the point where firms cover their total costs, including a market-based cost of capital and no more. Where firms persistently earn in excess of a normal return, this therefore signals that there may be limitations in the competitive process.

⁵³⁸ [Alphabet 2023 Form 10-K](#), page 87.

⁵³⁹ As set out in Market structure and concentration above, Google is the third largest provider in the combined IaaS and PaaS UK market and has [5-10]% market share compared to AWS and Microsoft's combined [60-80]% market share.

⁵⁴⁰ [CC3 \(Revised\)](#), paragraphs 114 and 116.

⁵⁴¹ For example, the ability to earn profits persistently above the competitive level could indicate the presence of entry barriers or, where a firm with a large market share has earned profits that have been persistently above the competitive level, may indicate significant market power. [CC3 \(Revised\)](#), paragraphs 118 and 119.

⁵⁴² [CC3 \(Revised\)](#), paragraphs 117, 118 and 126.

⁵⁴³ [CC3 \(Revised\)](#), paragraph 116.

- 3.320 However, even if we consider the analysis submitted by AWS comparing AWS' and Microsoft's ROCE for their cloud businesses to other cloud providers, we have identified a number of issues with the analysis which make it uninformative for the extent to which profits exceed the cost of capital (ie the ROCC-WACC differential) for other cloud providers. The analysis does not consider whether there has been a sufficiently long time period to assess the profitability of the other cloud providers.⁵⁴⁴ Besides Microsoft, only three other cloud providers in the analysis (Oracle, IBM and Akamai) have similar ROCE levels to AWS for a period of time and a positive ROCC-WACC differential over the period of 2018 to 2024 in AWS' submitted analysis.⁵⁴⁵ We identified a number of issues with the method for calculating these providers' ROCEs and WACCs, which make these ROCE-WACC differentials insufficiently comparable to our analysis for Microsoft's and AWS' cloud businesses.⁵⁴⁶
- 3.321 We also note that comparison to companies in other IT services markets is uninformative as these markets have not been the subject of our assessment, so we do not have information on whether there is effective competition in these markets, and/or whether the specific companies included could each have explanatory factors for profits exceeding the cost of capital within their main market(s) for the period of analysis submitted by AWS.
- 3.322 The CCIA submitted that our approach to assessing profitability will generally struggle to distinguish between final profits on investments that have succeeded and expected profits on risky investments that could fail.⁵⁴⁷ We recognise that our measure of capital employed could include investment to enable future expected (but uncertain) profits; we therefore consider the appropriate time period over which to examine the persistence of the gap between profitability and the cost of capital.⁵⁴⁸ However, as we use historic data in our analysis, our EBIT measure only includes achieved profits.
- 3.323 As set out above, Microsoft and AWS have been in the cloud markets for a sufficiently long period of time such that current profitability will be informative of whether sustained profits substantially above the cost of capital are being

⁵⁴⁴ For example, as noted earlier in this report, Oracle entered cloud services in 2016 compared to AWS and Microsoft, who started providing cloud services in 2002⁶ and 2008 respectively.

⁵⁴⁵ AWS' submission to the CMA [§].

⁵⁴⁶ For example, the ROCEs [§] Akamai are respectively calculated using a much [§] and Akamai's entire business, [§]. Oracle's capital employed is calculated simply by using Oracle Cloud Services' share of total Oracle revenue to allocate a portion of Oracle group total assets net of accounts payable. The WACCs for comparators are calculating using equity betas for the listed companies overall rather than considering the betas for the companies' cloud businesses specifically (see appendix E for our consideration of this when estimating AWS' and Microsoft's WACC). [§]; OVHcloud Consolidated Financial Statements; IBM Form 10-Ks; [§].

⁵⁴⁷ CCIA's response to the CMA's PDR dated 28 January 2025, page 2.

⁵⁴⁸ CC3 (Revised), paragraph 121.

achieved.⁵⁴⁹ Even if the cloud markets are still evolving, we would expect to see returns converge with the cost of capital over time.

- 3.324 We examined profitability over the last five years but also cross-checked our ROCE analysis over a longer time period (ten years for AWS and seven years only for Microsoft due to issues with the availability of data prior to financial year 2018).⁵⁵⁰
- 3.325 We cross-checked our ROCE analysis over a longer seven to ten-year period to examine whether, if there has been a rate of profit above the 'normal' level at some point, there is a clear downward trend towards the cost of capital that is apparent.

Our analysis

- 3.326 Below we summarise the findings of our profitability analysis.⁵⁵¹ We found that, with regard to gross margins:
- (a) Microsoft, AWS and Google and Oracle have generated gross margins for their cloud businesses which are higher than those of other cloud providers.
 - (b) Microsoft's Cloud and Cloud & Enterprise segments and AWS have consistently had the highest gross margins, in excess of [REDACTED] for AWS and [REDACTED] for both Microsoft segments over the last five years.
- 3.327 We found with regard to EBIT margins that:
- (a) The EBIT margins for AWS were consistently between 24% and 30% for the nine year period of 2015 to 2023 and increased to 37% in 2024.⁵⁵²
 - (b) The EBIT margins for Microsoft's Azure segment have consistently been between [20-30]% and [30-40]% for the last four years and [REDACTED] for the half-year period to December 2024, and for Microsoft's Cloud and Enterprise segment have consistently been between [40-50]% to [50-60]% for the last eight years and [REDACTED] for the half-year period to December 2024.⁵⁵³ EBIT margins for both segments are also on an upward trajectory.
 - (c) Google Cloud became profitable in FY23 and is reporting growing EBIT margins (averaging an annual ten percentage point improvement in margin in

⁵⁴⁹ We note that whilst Microsoft's Azure segment, based on our estimate of EBIT (as this is not tracked by Microsoft), only became profitable in 2020, Microsoft has [REDACTED] despite entering the cloud markets over ten years ago. [REDACTED]. As set out in appendix E, we consider it necessary to assess both Microsoft's 'Azure standalone services' and broader 'Cloud & Enterprise' segment to reach a view on the profitability of Microsoft's cloud services.

⁵⁵⁰ [REDACTED] response to the CMA's information request [REDACTED].

⁵⁵¹ Our full profitability analysis is set out in appendix E.

⁵⁵² CMA analysis of Amazon Form 10-Ks and 10-Qs.

⁵⁵³ CMA analysis of Microsoft Form 10-Ks and Microsoft's responses to the CMA's information requests [REDACTED].

the last 24 months).⁵⁵⁴ The EBIT margins for Google Cloud are currently significantly lower than AWS [REDACTED].

3.328 We estimated a WACC range of 10% to 14% for AWS and Microsoft's cloud businesses, against which we compared the results of our ROCE analysis.⁵⁵⁵ With regard to ROCE analysis, we found that:

- (a) AWS ROCE has consistently been substantially above our estimated WACC for the last ten years. AWS ROCE has been over [10-20]% since 2016 under all sensitivities used in our analysis.
- (b) Microsoft Cloud & Enterprise ROCE has consistently been substantially above our estimated WACC and in excess of [REDACTED] under all sensitivities used in our analysis for the last seven years.
- (c) Microsoft Azure ROCE has been substantially above our estimated WACC for the last four years and has been stable for the last two years. Microsoft Azure ROCE has been in excess of [REDACTED] under all sensitivities used in our analysis for the last three years.

3.329 Our ROCE analysis shows that there has been some annual variation in AWS ROCE, within a range of [REDACTED] to [REDACTED] for our baseline ROCE estimate.⁵⁵⁶

3.330 The ROCE for Microsoft Cloud & Enterprise ROCE has been trending downward since financial year 2021 ([REDACTED] for our baseline ROCE calculations), but still substantially above the WACC. We find this trend to be in large part due to [REDACTED],⁵⁵⁷ and in more recent years the impact of AI (discussed further below). The ROCE for Microsoft Azure has been increasing and then stable in financial years 2023 and 2024.

3.331 Our estimate of forecast Microsoft Cloud & Enterprise and Azure ROCE for financial year 2025 using EBIT and asset data for the half-year to December 2024 indicates that [REDACTED].⁵⁵⁸

3.332 We did not have sufficient data to estimate forecast AWS ROCE using our baseline calculation, However, we estimated forecast AWS ROCE under one of our ROCE sensitivities for financial years 2025 and 2026, using AWS' internal forecasts for EBIT and capex. For the reasons discussed in appendix E, we consider this sensitivity to understate ROCE; however, this estimate indicates that AWS ROCE will [REDACTED] in 2025 and 2026 but remain [REDACTED] the WACC.⁵⁵⁹

⁵⁵⁴ Period to 31 March 2025. CMA analysis of Alphabet Form 10-Qs.

⁵⁵⁵ See appendix E for detail on the calculation of our WACC estimate.

⁵⁵⁶ [REDACTED].

⁵⁵⁷ CMA analysis of [REDACTED].

⁵⁵⁸ Our estimate is based on our baseline ROCE calculation. CMA analysis of Microsoft Form 10-Ks, [REDACTED].

⁵⁵⁹ CMA analysis of Amazon Form 10-Ks and AWS' response to the CMA's information request [REDACTED].

Interpretation of our analysis

- 3.333 Although we recognise that there can be many factors that influence profitability in any given year, our analysis focuses on whether there have been returns in excess of the cost of capital for a sustained period and if there are relevant contextual factors (other than limitations in the competitive process) which might account for it.
- 3.334 It is not necessary to define a specific value for the differential between a cloud provider's ROCE and WACC that would indicate a level of profits above the 'normal' level. We consider whether the differential is 'unequivocally substantial' but also the length of the period over which the differential persists and profit trends.⁵⁶⁰
- 3.335 As noted above, we compare cloud providers' margins in our analysis to provide supporting context and trends to help inform our assessment of profitability. We note we have not put forward any views on whether there is a causal relationship between cloud provider margins and size/growth.

Margins and ROCE

- 3.336 As set out in our analysis, we find Microsoft's Cloud and Cloud & Enterprise segments and AWS have consistently had the highest gross margins. AWS' submission that its gross margins are in line with other cloud providers for the period 2021-2023 is based on the inclusion of companies (eg Adobe, SAP, SPS Commerce) that we do not consider to be directly relevant to the public cloud infrastructure services market which is the focus of this investigation, as well as a gross margin for Oracle which appears to be based on its 'cloud and license' segment (of less relevance as it is likely to include a substantial portion of on-premises services).
- 3.337 In the analysis AWS submitted, [REDACTED].⁵⁶¹
- 3.338 We do not agree with AWS' submitted ROCE ranges based on estimates of capital employed that include intangible assets for which we have not received supporting evidence and have determined do not merit inclusion in capital employed.⁵⁶²
- 3.339 In response to our working paper, AWS submitted that there is a clear downward trend in AWS' ROCE over 2021 to 2023. AWS submitted that this 'implies a convergence between ROCE and WACC which suggests the presence of fierce

⁵⁶⁰ CC3 (Revised), paragraphs 120, 122 and 124.

⁵⁶¹ AWS' submission to the CMA [REDACTED].

⁵⁶² See criteria in CC3 (Revised), annex A paragraph 14. We set out our consideration of intangible assets in more detail in appendix E.

competition despite a temporal competitive advantage granted to AWS through its constant innovation'.⁵⁶³

- 3.340 However, our analysis found that AWS ROCE increased in 2024 to [REDACTED] and remained in excess of the WACC range of 10% to 14%, over the 2021 to 2023 period (ranging from [REDACTED]). The two years of decreasing ROCE in 2022 and 2023 are consistent with annual ROCE variation as noted above and the increase in 2024 is inconsistent with a trend of decreasing ROCE and [REDACTED]. As set out in the following section, our analysis also indicates that the downward trend for AWS ROCE in 2022 and 2023 was largely driven by a new wave of investment in AI which has impacted AWS' capital employed for its overall cloud business.
- 3.341 We discuss our interpretation of the recent movements in AWS ROCE in more detail in the following section, but overall our analysis does not indicate that AWS ROCE is or clearly will be at the level of the WACC.
- 3.342 We consider innovation and investment in the following section.

Considerations when assessing ROCE

- 3.343 As stated above, we have considered reasons why ROCE may exceed WACC for a firm. We have not seen any indications, and none have been suggested to us, that Microsoft Azure and Cloud & Enterprise and AWS' ROCE has been above WACC due to cyclical factors. We also do not consider there to be any transitory price or other marketing initiatives in place driving ROCE and have not seen evidence of superior efficiency justifying ROCE being above WACC, and no evidence has been suggested to us.
- 3.344 As explained above, we consider the last five years to be relevant to assess the profitability of Microsoft's and AWS' cloud businesses using our ROCE analysis. Returns may be subject to investment cycles, and if the cloud market is one where large and risky investments have been made then we may have needed to consider profitability over a relatively long timescale.⁵⁶⁴ We consider our longer ROCE cross-check period of seven to ten years to be sufficiently long to take any investment cycles into account and to observe whether the rate of profit for AWS and Microsoft's cloud businesses has returned to, or is returning to, a 'normal' level. Whilst we recognise that the profitability of some firms may exceed the 'normal' level at particular points in time (as explained above),⁵⁶⁵ seven or more years is a 'sustained' period rather than point in time for the cloud market.
- 3.345 In addition, Microsoft's and AWS' investments in their cloud businesses have been taken into account in our measure of capital employed used in calculating ROCE.

⁵⁶³ AWS' submission to the CMA [REDACTED].

⁵⁶⁴ CC3 (Revised), paragraph 121.

⁵⁶⁵ CC3 (Revised), paragraph 117.

As Microsoft and AWS have characterised their investments in the cloud market as continuous,⁵⁶⁶ (rather than an initial one-off investment), these investments will continually be adding to capital employed and hence are reflected in our ROCE figures.

Impact of innovation

- 3.346 We have also considered whether the profitability of Microsoft's and AWS' cloud businesses may be explained by a high degree of innovation in the cloud market.
- 3.347 We recognise that profitability could exceed the 'normal' level at particular points in time due to a firm earning higher profits as a result of past innovation, in line with our Guidelines.⁵⁶⁷ However, we would expect the competitive advantage conferred by an innovation to be eroded over time by competition such that a past innovation would not explain sustained profits in excess of the cost of capital over the long term. On this basis, absent further explanation, we would consider the ROCE substantially in excess of WACC for Microsoft's and AWS' cloud businesses over a sustained period to be an indication of limitations in the competitive process in the cloud market.
- 3.348 AWS has submitted that the level profitability in the cloud industry is explained by continuous innovation.⁵⁶⁸ We accept that a situation of continuous new innovation (whereby repeated new innovations may result in temporary competitive advantages that give rise to repeated higher returns) could, in principle, then result in continuous periods of temporary profits exceeding the cost of capital, which cumulatively present as sustained profitability above the 'normal' level over time. This is consistent with one of the papers cited by AWS and the hypothetical scenario submitted by AWS.⁵⁶⁹
- 3.349 We note that the other papers cited by AWS, as evidence that economic theory supports that conditions of continuous innovation in an industry can explain ROCE being persistently above WACC, discuss the impact of innovation on the profitability of firms relative to other firms in terms of operating profits or margins, and do not consider profitability in terms of returns relative to the cost of capital.⁵⁷⁰

⁵⁶⁶ See for example: AWS' reference to 'AWS' continuous investment in physical infrastructure' in [AWS response to the CMA's working papers and updated issues statement](#), paragraph 32; and Microsoft's reference to 'record levels of investment, with at least three firms (Amazon, Google and Microsoft) investing tens of billions each and every year' in [Microsoft response to the Competitive landscape, Committed spend agreements and Egress fees working papers](#), paragraph 3.

⁵⁶⁷ [CC3 \(Revised\)](#), paragraphs 117 and 118.

⁵⁶⁸ AWS' submission to the CMA [X].

⁵⁶⁹ AWS' submission to the CMA [X]; Roberts, P. W. (1999) 'Product innovation, product-market competition and persistent profitability in the US pharmaceutical industry', *Strategic Management Journal*, 20(7)), page 656.

⁵⁷⁰ Only one paper (Roberts, P. W. (1999) 'Product innovation, product-market competition and persistent profitability in the US pharmaceutical industry', *Strategic Management Journal*, 20(7)) assesses returns, rather than operating profits or profit margins, and uses a measure of return on assets without reference to a cost of capital. Only this same paper and another paper (Cable, J. R. and Jackson, R. H. G. (2008) 'The Persistence of Profits in the Long Run: A New Approach', *International Journal of the Economics of Business*, 15(2).) consider competition as variables in their analysis. The Cable, J. R. and Jackson, R. H. G paper finds innovation and competition to be non-observable variables.

We consider the findings on operating profits or margins to be uninformative as innovation will also impact a firm's returns through its capital base and cost of capital.

- 3.350 We would expect the impact of any innovation on profitability to be reflected in our ROCE analysis. Any contribution to profitability should be captured in the measure of earnings, and investment in innovation should be captured in the measure of capital employed,⁵⁷¹ particularly due to the ongoing rather than one-off nature of investment that cloud providers have described. For example, neither Microsoft nor AWS have submitted that they made a large one-off past investment in innovation (that is no longer being captured in capital employed) that continues to be a source of competitive advantage-driven profitability. Similarly, any ongoing investment cycles in general (or potential new investment phases, such as for AI as submitted by Microsoft and considered below), will be captured in the measure of capital employed.
- 3.351 We would also expect the risk or uncertainty of investment in innovation that results in a potential profit pay-off for firms (as characterised by Microsoft and AWS as applying to the cloud market)⁵⁷² to be reflected in the cost of capital. Therefore, findings on operating profits or margins will not necessarily be reflective of, or explain, relative ROCE-WACC differentials between firms.
- 3.352 We consider the above proposed scenario of continuous innovation resulting in sustained profits exceeding the cost of capital to be predicated on:
- (a) There being investment in the specific innovation that is introduced; and
 - (b) There being a specific profit contribution from the innovation when introduced (as a result of the temporary competitive advantage arising from the innovation).
- 3.353 We considered whether the above can be established for AWS, given its submissions. As considered further later in this chapter, AWS submitted its total capex and R&D expenditure and its number of patents and new services and features released as evidence of its innovation, along with examples of specific product and process innovations.⁵⁷³ Microsoft has also submitted its R&D spend,

⁵⁷¹ Or if investment is not capitalised, it will be expensed and therefore reflected in the firm's EBIT. We have assumed that innovation in the cloud markets would typically require capital investment, but note that in general, innovation may or may not be capital-intensive (with associated investment expenditure).

⁵⁷² [AWS' response to the CMA's PDR dated 28 January 2025](#), paragraph 19.

⁵⁷³ We note that two of the examples are partnerships with other companies. AWS provided information the customer uptake of the six product innovation examples, but did not provide information the investment or profit contribution for any of the product or process innovation examples. AWS' submission to the CMA on quality and innovation investments and profitability dated 16 April 2025, paragraphs 29-40 and Figures 2 and 3.

capex, number of feature updates and data centre capacity expansion as evidence of innovation.⁵⁷⁴

- 3.354 We note that capex (including for capacity expansion) and R&D expenditure does not necessarily equate to innovation.⁵⁷⁵ Microsoft and AWS also did not provide evidence to demonstrate how these patents and new services and features released are contributing to or driving profitability (particularly relative to their existing services). Our assessment of quality and innovation, set out later in this chapter, also finds that most of these metrics do not capture the degree to which customers are using cited new products and features, and we find there is evidence of some innovations in cloud services being more incremental in nature.
- 3.355 We note that AWS has also submitted that its intangible assets are not reflected in our measure of its capital employed. Intangible assets could relate to investment in innovation, although AWS has not submitted this. We consider AWS' submissions on intangible assets in detail in appendix E, but have not found evidence to support the inclusion of the estimate submitted by AWS in capital employed. Our understanding is that our measure of capital employed for AWS already includes intangible assets that are specifically identifiable and/or allocated to AWS in its public accounts.⁵⁷⁶ We also note that in the ROCE range inclusive of intangible assets submitted by AWS, AWS' ROCE remains at [X] in the seven year period of analysis ie in excess of our WACC range.⁵⁷⁷ As above, AWS also did not provide evidence to demonstrate how the intangible assets represented by its submitted measure are contributing to or driving its profitability (particularly relative to its existing services).
- 3.356 Given the above, we have not found sufficient evidence to demonstrate that AWS is constantly achieving temporary profits in excess of the cost of capital from a rolling release of highly innovative new products, and have not found evidence to support AWS' view that its sustained profit gap (in the form of ROCE above WACC) is explained by continuous innovation. We similarly have not seen evidence of continuous innovation explaining Microsoft's ROCE in excess of WACC.
- 3.357 We therefore do not consider the sustained ROCE in excess of WACC for Microsoft's and AWS' cloud businesses to be explained by the level of innovation in the cloud market.

⁵⁷⁴ Microsoft response to the Provisional Decision Report, paragraph 67(iii) to (v); Microsoft [Response to The Competitive landscape, Committed spend agreements and Egress fees working papers](#), paragraph 12(c).

⁵⁷⁵ For example, capex could be for maintaining existing services, and R&D may not necessarily result in successful innovations.

⁵⁷⁶ We use AWS total assets reported in public accounts as a starting point for capital employed, which only has specified exclusions of 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. Amazon Form 10-Ks.

⁵⁷⁷ AWS' submission to the CMA [X].

Impact of investments in AI

- 3.358 We note that Microsoft and AWS have recently been substantially increasing their investments in their cloud infrastructure, in particular in investments to support the development of their AI services.
- 3.359 Whilst both Microsoft and AWS suggested that their significant investments in cloud reflect competition in the market, we do not accept that investment in and of itself necessarily explains sustained profits in excess of the cost of capital. We have also considered the extent to which recent investments reflect a new investment cycle for AI.
- 3.360 Microsoft and AWS have been making substantial capex investments in supplying accelerated compute (see our assessment of this later in this chapter). For example:
- (a) [REDACTED],⁵⁷⁸ and
 - (b) Microsoft's cloud infrastructure and server capex in financial years 2022 and 2023 was [REDACTED].⁵⁷⁹
- 3.361 This may indicate that a new investment cycle is commencing for these firms' cloud businesses in relation to AI services.
- 3.362 For a new investment cycle, we might expect to see initial upfront investment that then generates returns in the future. Our review of Microsoft's and AWS' public announcements and internal documents indicates that both companies are expecting large, accelerating revenue growth from their AI services.⁵⁸⁰
- 3.363 Microsoft and AWS have noted the [REDACTED] margins resulting from the significant capex associated with AI.⁵⁸¹ Microsoft's internal documents also indicate that:
- (a) [REDACTED].⁵⁸² [REDACTED].⁵⁸³
 - (b) [REDACTED].⁵⁸⁴
 - (c) [REDACTED].⁵⁸⁵
- 3.364 Microsoft's and AWS' increased investment in AI, together with [REDACTED], will have put downward pressure on the ROCE for AWS, Microsoft Cloud & Enterprise and

⁵⁷⁸ CMA analysis of AWS' response to the CMA's information request [REDACTED].

⁵⁷⁹ CMA analysis of Microsoft's response to the CMA's information request [REDACTED]; Microsoft's response to Ofcom's information request [REDACTED].

⁵⁸⁰ See chapter 4 of this report.

⁵⁸¹ AWS' response to the CMA's information request [REDACTED]; Microsoft's response to the CMA's information request [REDACTED].

⁵⁸² [REDACTED] response to the CMA's information request [REDACTED].

⁵⁸³ [REDACTED] responses to the CMA's information request [REDACTED].

⁵⁸⁴ [REDACTED] response to the CMA's information request [REDACTED].

⁵⁸⁵ [REDACTED] response to the CMA's information request [REDACTED].

Microsoft Azure in recent years, as it increases capital employed at a greater rate than the growth in EBIT.

- 3.365 Based on the comments in relation to AI growth in Microsoft's public announcements⁵⁸⁶ and internal documents, Microsoft's increased cloud capex appears to represent an upfront investment ahead of developments that would be expected to generate their own returns in the future.
- 3.366 Microsoft's margin and capex information suggests that if cloud AI services were removed from Microsoft data for financial years 2023-2024, we would expect ROCE to be higher. To the extent that these returns are included in the Azure and Cloud & Enterprise businesses, there could be increased EBIT growth in future years with corresponding impact on ROCE.
- 3.367 Similarly, our analysis of [REDACTED].⁵⁸⁷
- 3.368 We recognise that Microsoft and AWS are making substantial investments in AI services for which the future returns are uncertain and which may not reflect the competitive pressures of the wider cloud market. We consider our ROCE analysis for the last five years (taking into account the downward pressure from AI investment) to be informative of the profitability of the wider cloud market. Irrespective of potential future growth from AI services, our ROCE analysis for Microsoft and AWS includes their AI investment and the comparatively [REDACTED] profit contribution associated with this currently, and ROCE levels have consistently remained above WACC even with this increase in investment.
- 3.369 We do not assume that Microsoft's and AWS' recent investments in AI will be automatically successful, but have considered whether the impact of this AI investment indicates that the ROCE for their cloud businesses is likely to converge with WACC in the future. As noted above, Microsoft's and AWS' public announcements and internal documents indicate that both companies are expecting large, accelerating revenue growth from their AI services.
- 3.370 On that basis, ROCE is and will be in excess of WACC for Microsoft's cloud businesses and AWS' when taking into account of Microsoft's and AWS' recent and continuing investments in AI. This is supported by the forecast data and commentary reviewed, and is consistent with ROCE for each of Microsoft's and AWS' cloud businesses stabilising around current levels in the future, rather than declining.

⁵⁸⁶ [Microsoft Q2 FY24 Earnings Conference Call transcript](#); [Microsoft Q3 FY24 Earnings Conference Call transcript](#).

⁵⁸⁷ CMA analysis of [REDACTED] response to the CMA's information request [REDACTED].

Pricing trends

- 3.371 In markets subject to effective competition, prices are likely to respond to changing supply and demand conditions. Firms will seek to win business by improving their prices and other aspects of their offer. The pattern of prices over time can therefore indicate something about the nature of competition. However, there may be several factors affecting prices and we have taken this into account when making inferences from this type of analysis.⁵⁸⁸
- 3.372 In particular, we have sought to understand the trends in prices over time among Microsoft, AWS and Google and whether that could tell us anything about competition in cloud services. In doing so, we note that prices can be driven by both competitive constraints and other factors such as costs and changes in quality. This means any analysis of price alone is purely indicative, particularly in a growing industry with strong economies of scale where we would expect costs to fall, and it needs to be considered alongside other evidence such as costs – this is why we have also considered profitability above.

Stakeholder views and our assessment

- 3.373 Google submitted that the cloud market is characterised by a downward pricing trend, and the availability of numerous deals and pricing offers.⁵⁸⁹ An academic stated that prices have fallen, not risen.⁵⁹⁰ Another academic stated that cloud computing's prices continue to be deflationary despite economic inflation, which highlights the competition among advanced companies.⁵⁹¹
- 3.374 Both Microsoft and AWS⁵⁹² submitted quantitative analyses of historical prices for some of their services. We consider these below.
- 3.375 During Ofcom's market study, AWS submitted an analysis of global net price trend data and UK list price trend data for [redacted] cloud infrastructure services (S3, EC2 and DT).⁵⁹³ AWS said its analysis shows that:
- (a) global average net prices decreased significantly in both nominal and real terms for all three services;

⁵⁸⁸ CC3 (Revised), paragraph 108

⁵⁸⁹ Google's response to the CMA's Issues Statement dated 17 October 2023, paragraph 9.

⁵⁹⁰ Dr Barker, Comment on The UK Competitive Market Authority's (CMAs) Cloud Services Market Investigation Updated Issues Paper and Working Papers 4-6 on The Supply of Public Cloud Infrastructure Services In the UK Covering The CMA's 1. Updated issues statement on Public cloud infrastructure services market investigation: 2. Licensing Practises Working Paper; 3. Technical Barriers; and 4. Potential Remedies, page 31.

⁵⁹¹ R. Parisi, The Cloud Services Markets' Competitive Landscape: A contribution to the Competition and Markets Authority, page 11.

⁵⁹² Microsoft's submission to CMA [redacted]. Microsoft had also previously submitted pricing analysis to Ofcom. We understand the analysis submitted to us to be an update on the analysis submitted to Ofcom, and have therefore focused on the later submission to us.

⁵⁹³ AWS' submission to Ofcom [redacted].

- (b) its UK list prices have been constant since November 2016 for S3 Standard ([§]) and since December 2014 for data transfer-out (before which both had been declining), consistent with a decline in real prices in recent years; and
- (c) its UK average list prices across all EC2 instance families declined significantly in nominal terms between 2016 and 2022, with no increase in list price for any instance type between 2019 and 2022. It said average EC2 list prices increased in 2022 due to the introduction of new instances.⁵⁹⁴

- 3.376 AWS noted that it has innovated and improved the quality of these three services, implying a further reduction in quality-adjusted real prices.⁵⁹⁵ AWS similarly submitted that, according to our analysis in the PDR, the vast majority of Google's services have also decreased in price and the majority of Microsoft's services have also decreased in price. It submitted that if one were to account for quality, the price decreases would be even more pronounced.⁵⁹⁶
- 3.377 AWS further submitted a revised analysis that addressed all of the concerns we had on the above analysis in our PDR.⁵⁹⁷ AWS submitted that its revised analysis shows that the prices of S3, EC2 and DT have fallen between 2018 and 2024.⁵⁹⁸
- 3.378 We have reviewed AWS' analysis. We consider that the analysis provides evidence that three of AWS' services have fallen in price between 2018 and 2024. However, the analysis does not cover other services, including some with significant associated revenues, nor does it explore price changes at a product level in the same way we do in our analysis below.
- 3.379 Microsoft said its analysis shows that quality-adjusted real prices have fallen.⁵⁹⁹ In particular, Microsoft said that while nominal effective prices (ie nominal prices calculated as total revenues for each product, net of discounts, divided by the number of units sold) for the top five Azure products have remained stable, if not decreased, when controlling for inflation, real effective prices have decreased materially between 2019 and 2022.⁶⁰⁰
- 3.380 In response to our PDR, Microsoft submitted that 'quality-adjusted real effective prices' for the top five Azure products, which constitute [§]% of Azure's total revenue, fell from 2018-22. It submitted that 'real effective prices' for the top five products fell by [§]% as of 2022.⁶⁰¹

⁵⁹⁴ AWS' submission to Ofcom [§].

⁵⁹⁵ AWS' submission to Ofcom [§].

⁵⁹⁶ AWS' submission to the CMA [§].

⁵⁹⁷ See paragraph 3.284 of the PDR.

⁵⁹⁸ S3 stands for Simple Storage Service. EC2 stands for Elastic Compute Cloud. DT stands for Data Transfer. AWS' submission to the CMA on pricing dated 17 April 2025, paragraph 9.

⁵⁹⁹ Microsoft's [response to the Competitive landscape, Committed spend agreements and Egress fees working papers](#), paragraph 12 (a) and (b).

⁶⁰⁰ Microsoft's submission to CMA [§].

⁶⁰¹ Microsoft's submission to the CMA [§].

3.381 We have reviewed Microsoft's pricing analysis and have identified the following limitations:

- (a) Microsoft deflated prices using the consumer price index (CPI). This is a general inflation index that likely overestimate the specific inflation of costs related to the cloud industry. We consider it more appropriate to use the ONS Information Service Activities index as it accounts for inflation specific to cloud services better than CPI.⁶⁰²
- (b) Microsoft's categorisation of prices as simply 'decreasing', 'not changing', or 'increasing' omits any information related to the magnitude of such changes. This means we cannot determine which products are driving any overall changes in price.
- (c) Microsoft analysed unweighted prices, ie it considered each product equally regardless of its importance (eg in terms of customer spend on them). This means that products are considered equally, when in reality changes in price of more popular products affect more customers (and are therefore more significant).
- (d) Microsoft aggregated all products' prices within a service, therefore conflating the effect of price changes of single products and that of new products replacing old ones at different price points. For example, this will overstate price decreases if new products are more expensive than the ones they replace.

3.382 Our own analysis, which we explain below, builds on Microsoft's analysis.⁶⁰³

3.383 AWS submitted that there was no evidence that it prices less competitively for existing customers, which means that any concerns that AWS may take advantage of customers who are 'locked in' are not borne out by the data. It stated that economic theory predicts that, if customers were locked in, existing customers would eventually pay higher prices or receive lower discounts compared to new customers.⁶⁰⁴ Similarly, Microsoft submitted that the value of discounts received by Azure customers remained stable at [REDACTED]% [REDACTED] while CSAs for UK customers [REDACTED].⁶⁰⁵

3.384 We consider that existing customers may still be harmed by being 'locked in' even if they are charged the same prices as new customers. This is because, as the requirements of existing customers evolve and new services become available, they may find it difficult to switch to (and therefore take advantage of) a different

⁶⁰² See [Office for National Statistics Industry Deflators](#), accessed on 15 August 2024.

⁶⁰³ See appendix F for a detailed description of this.

⁶⁰⁴ AWS' submission to the CMA [REDACTED].

⁶⁰⁵ Microsoft's submission to the CMA [REDACTED].

cloud provider or service which better meets their needs, restricting their ability to fully benefit from cloud services.

Our quantitative analysis

- 3.385 We analysed Microsoft and Google data for each of their top five services by revenue in 2022 to describe trends in prices for cloud services and products. We did not include AWS in this analysis as the data we received was not comparable enough to replicate the same methodology in full. AWS provided its own analysis which we discussed above.
- 3.386 Our analysis builds on the Microsoft analysis presented above, and attempts to overcome some of its limitations:
- (a) We have used a different price index to control for inflation. While Microsoft used a general inflation index, the Consumer Price Index (CPI), as set out above, we consider it more appropriate to use the ONS Information Service Activities index.⁶⁰⁶
 - (b) We have considered average prices of products within services weighted by the revenue of each product. We consider this to be an improvement on Microsoft's submission as each product is weighted by its relative importance.
 - (c) We have computed two measures of price trends: (i) the revenue-weighted average prices for each service over time, and (ii) the revenue-weighted average of the percentage change in price of products. The second metric excludes the so-called 'composition effect', ie the effect of changes in the set of products underlying a service over time.⁶⁰⁷
 - (d) When looking at price changes, we have looked not only at whether the price change was positive or negative, but also the magnitude of this change.
 - (e) We also excluded all products that are free. That is, we have excluded products that are priced at zero when they are first and last observed.⁶⁰⁸
- 3.387 Overall, our results show that there is variety in the direction and size of real (ie after controlling for inflation) price changes for each of Microsoft's top five services. While most of the products underlying each of these services decreased

⁶⁰⁶ [Office for National Statistics Industry Deflators](#), accessed on 15 August 2024.

⁶⁰⁷ If, for example, the price of one service decreases over time, that may be due to (a) products under that service changing prices and (b) new cheaper products replacing older more expensive ones. While the revenue-weighted average prices for each service over time captures both, the revenue-weighted average of the percentage change in price of products isolates only (a).

⁶⁰⁸ Microsoft submission to the CMA [38]. This is particularly relevant when taking the percentage change in price as products with a price of zero cannot return a percentage change.

in price, a significant share of products increased in price and some more than doubled.

3.388 Amongst Microsoft's top five services, two have decreased in price, one has stayed stable, and two have increased in price:

- (a) Relative to 2018, the price of [X] and [X] decreased by [30-40%] and [60-70%] respectively.
- (b) Prices for [X] have fluctuated substantially over time. However, [X] prices for [X] were similar to prices in 2018.
- (c) For [X], prices increased by [70-80%] between 2018 and 2022 and the prices for [X] more than doubled between 2018 and 2022.

3.389 For Microsoft, the average price of products within the [X] and [X] services has increased. This is driven by changes in the product mixes underlying these services.

3.390 Within Microsoft's top five services:

- (a) the [X] of products within each service decreased in price;
- (b) approximately [30-40%]-[40-50%] of products in each SL2 increased in price; and
- (c) [0-5%]-[10-20%] of products did not experience any material change in price.⁶⁰⁹

3.391 For Google, most products experienced a small decrease in real price. A minority of products increased or did not change in real price:

- (a) Relative to 2020, prices of four out of five of the top services offered by Google decreased by a range of [30-40%]-[50-60%].
- (b) Prices for one other top service increased by approximately [5-10%] in real terms.
- (c) the majority of products in each service decreased in price;
- (d) [10-20%] of the products in one service [X] increased in price, while no products in the other four services [X] increased in price; and
- (e) Between [0-5%] and [20-30%] of products in each service experienced no material change in price.

⁶⁰⁹ We define 'no change' in price as a price change of less than 1% in absolute value.

- 3.392 We note here that, in response to the PDR, Microsoft further submitted that the CMA has not indicated what proportion of prices are increasing or decreasing.⁶¹⁰ As outlined above, we have considered what proportion of prices are increasing or decreasing above, alongside other relevant metrics, such as the size of those changes.
- 3.393 In interpreting these results, we have accounted for the fact that this analysis does not:
- (a) account for changes in costs. We note that if unit costs were to fall faster than prices, this would indicate an increase in profit margins, an indicator of market power;
 - (b) account for changes in quality of the products; and
 - (c) cover all cloud providers to the same extent, with the analysis focusing only on Microsoft and Google due to lack of comparable data from AWS.⁶¹¹
- 3.394 Consistent with the above, and as mentioned at the beginning of this section, any analysis of price alone is purely indicative and needs to be considered alongside other evidence. We place more weight on our assessment of the cloud providers' profitability than the pricing analysis described above. This is because the former accounts for broader financial indicators (eg revenue, which is determined by prices and other variables, and costs).
- 3.395 The conclusion of our assessment of pricing trends is that, while the prices of most products and services across Microsoft and Google's most used services have fallen, the price of some products and services, including among the top five services has increased over the relevant period. AWS' analysis provides evidence that three of AWS' services have fallen in price between 2018 and 2024. However, the analysis does not cover other services, including some with significant associated revenues, nor does it explore price changes at a product level in the same way we do in our analysis below. This presents a mixed picture of pricing trends that needs to be considered alongside other evidence on market outcomes.

Quality and innovation

- 3.396 Outcomes related to quality, innovation and range may also be useful indicators of competition.⁶¹²
- 3.397 The Jigsaw report indicates that customers value innovation and find value in sharing the burden of architecture design, upgrades and maintenance. Many

⁶¹⁰ Microsoft response to the Provisional Decision Report, paragraph 67(i).

⁶¹¹ See appendix F for more detail.

⁶¹² [CC3 \(Revised\)](#), paragraph 127.

customers valued cloud providers innovating on infrastructure and the user being able to focus on their core business function.⁶¹³

Cloud provider views

3.398 Microsoft, AWS and Google submitted that levels of innovation in the market are high.⁶¹⁴

3.399 AWS submitted that:

- (a) innovation brought by the prospect of growing the market would be an indication of the value added that is created by innovation and that it would increase output, which is a positive outcome;⁶¹⁵
- (b) innovation is omnipresent in the cloud segment⁶¹⁶ and cloud providers are characterised by a high degree of innovation consistent with a dynamically competitive market;⁶¹⁷
- (c) [REDACTED],⁶¹⁸ and
- (d) economic literature confirms high dynamism and innovative activity within the cloud sector.⁶¹⁹

3.400 Microsoft submitted that:

- (a) cloud has advantages relative to other forms of IT models and there have been significant technological developments in the industry;⁶²⁰
- (b) innovation occurred also at the infrastructure level⁶²¹ and that it advanced first- and third-party innovation across a range of services and features, throughout the whole stack and with significant customer uptake;⁶²²
- (c) there are examples of innovation in the cloud services market leading to direct benefits for consumers;⁶²³ and

⁶¹³ Jigsaw Report, page 25.

⁶¹⁴ AWS' submission to the CMA [REDACTED]; AWS' [Economic response to the CMA's Competitive landscape working paper](#) paragraph 110 and 112; [REDACTED]; Google's [response to the CMA's Issues Statement dated 17 October 2023](#), paragraph 9.

⁶¹⁵ AWS' [Economic response to the CMA's Competitive landscape working paper](#), paragraph 116.

⁶¹⁶ AWS' [Economic response to the CMA's Competitive landscape working paper](#), paragraph 110.

⁶¹⁷ AWS' submission to the CMA [REDACTED].

⁶¹⁸ AWS' submission to the CMA [REDACTED].

⁶¹⁹ AWS' submission to the CMA [REDACTED].

⁶²⁰ Microsoft's submission to the CMA [REDACTED].

⁶²¹ Microsoft's submission to the CMA [REDACTED].

⁶²² Microsoft's submission to the CMA [REDACTED].

⁶²³ Microsoft's submission to the CMA [REDACTED].

- (d) cloud computing providers are rapidly innovating and building, both individually and with a host of partners, new AI platform services and development capabilities.⁶²⁴

3.401 Google submitted that the market for cloud services is characterised by continuous innovation.⁶²⁵

3.402 AWS and Microsoft both submitted metrics relating to the degree of innovation in the cloud services market.

3.403 AWS submitted that:

- (a) cloud providers' R&D share of revenue is in line with that of other IT companies and higher than that of some electronics manufacturers like Samsung and Dell. The R&D share of revenue is relatively similar across different, smaller cloud providers;⁶²⁶
- (b) the quality of AWS' innovations is evidenced by various factors eg the fact that the innovations are often directly responsive to customer requests, are aimed at ensuring the quality of a service for customers and experience quick and extensive customer uptake;⁶²⁷ and
- (c) it regularly introduced additional services and features each year between 2011-2023, and new EC2 instance families each year between 2008-2022,⁶²⁸ and [REDACTED].⁶²⁹

3.404 Microsoft submitted that:

- (a) Microsoft's R&D spend and capex has accelerated since 2010, more than tripling in absolute R&D spend, and a 12x rise in absolute capex;⁶³⁰
- (b) its cumulative number of feature updates for the top 20 UK cloud services was large in 2023;⁶³¹ and
- (c) it has expanded the UK capacity of its dedicated Azure data centres by [REDACTED].⁶³²

⁶²⁴ Microsoft's response to the CMA's PDR dated 28 January 2025, paragraph 10.

⁶²⁵ Google's response to the CMA's Issues Statement dated 17 October 2023, paragraph 9.

⁶²⁶ AWS' submission to the CMA [REDACTED].

⁶²⁷ AWS' Economic response to the CMA's Competitive landscape working paper, paragraph 112; AWS' submission to the CMA [REDACTED].

⁶²⁸ AWS' submission to the CMA [REDACTED]; AWS' submission to the CMA [REDACTED].

⁶²⁹ AWS' submission to the CMA [REDACTED]; AWS' submission to the CMA [REDACTED].

⁶³⁰ Microsoft's submission to the CMA [REDACTED].

⁶³¹ Microsoft's response to The Competitive landscape, Committed spend agreements and Egress fees working papers, paragraph 12 (c); Microsoft response to the Provisional Decision Report, paragraph 67.iii.

⁶³² Microsoft's submission to the CMA [REDACTED].

- 3.405 Microsoft submitted that the PDR prioritised static measures of competition, which provide an incomplete picture, instead of focusing on evidence which shows that prices are falling and innovation is prevalent.⁶³³ In relation to this submission, we have considered how a range of indicators have changed over time in our assessment, including pricing and profitability.
- 3.406 AWS submitted that the PDR did not present evidence that would suggest that innovation or quality provision is restricted or ‘too low’ in the cloud segment, nor did the PDR make any prediction of how it would expect innovation and quality provision to be characterised in a competitive market.⁶³⁴

Our assessment

- 3.407 We have considered the measures of quality and innovation provided by cloud providers and we consider that they show there has been innovation in cloud services.
- 3.408 We have considered in detail the quantitative measures provided by cloud providers in relation to quality as well as innovation, including metrics related to quality assurance, customer satisfaction, software updates, number of new services and features of services launched, number of patents and patent citations received, and levels of investment in R&D.⁶³⁵
- 3.409 However, it is difficult to evaluate quantitative evidence related to quality and innovation because:
- (a) considering measures of quality and/or innovation across different time periods can lead to different conclusions;⁶³⁶
 - (b) most metrics do not capture the degree to which cloud providers’ innovations are advancing technologies in the market and are valued by customers;⁶³⁷
 - (c) the data available does not help distinguish between investment in innovations which might (intentionally or otherwise) reinforce customer lock-

⁶³³ Microsoft’ submission to the CMA [§].

⁶³⁴ AWS’ submission to the CMA [§].

⁶³⁵ See appendix G

⁶³⁶ AWS data shows that there have been new patents granted each year for the last six years, but the extent of new patents varies over the years and its growth is highly sensitive to the time period considered. For example, when considering the number of AWS patents and annualised citations in the UK from 2019-2023, the [§] would have been less than [§] and the [§] would be considerably less significant than during the period 2018-2023. This is driven by an increase in the number of patents and citations from 2018 to 2019, with the number of patents granted from 2019 to 2023 only slightly increasing. Specifically, AWS was granted [§] patents in 2018 and [§] patents in 2023, a [§]. It was granted [§] patents in 2019, giving an [§]% from 2018 to 2019 but just [§]% from 2019 to 2023. CMA analysis of AWS’ submission to the CMA [§].

⁶³⁷ For example, AWS cited its number of new features, services and patents as evidence of its innovation, but these metrics do not capture the extent to which customers are using, and benefiting from, these products and features.

in as opposed to innovations which would enable greater interoperability, switching and/or multi-cloud. We consider this further in chapter 6;

- (d) there is no clear way for us to evaluate the significance of these metrics and levels of quality and/or innovation, even when compared against other industries – there may be reasons why the cloud industry should show a higher/lower innovation rate than other industries which cannot be factored in our analysis; and
- (e) some metrics are not specific to cloud services which limits our ability to use them to draw conclusions about the degree of innovation in these markets. For example, the R&D expenditure used by AWS is not specific to its cloud services.⁶³⁸

3.410 Despite these limitations, the evidence suggests there has been some innovation in the market, including amongst larger cloud providers who innovate and invest to win business. This is reflected in some customers' perception of the providers as 'great innovators'.⁶³⁹

3.411 Nonetheless, observing that there is innovation in the market can be consistent with a finding that competition in the market is not working as well as it could. Cloud providers may have incentives to innovate even where customers are inhibited from switching or using multi-cloud if this generates additional revenue, particularly if the additional revenues exceed the cost of innovating. In this respect, even a monopolist could have incentives to invest in innovation, but its level might be lower and its nature different than in a more competitive market.

3.412 Given the limitations in being able to assess whether the observed levels of innovation are indicative of a well-functioning market, we place more weight on our assessment of the cloud providers' profitability than on analysis of quality and innovation indicators described above.

Prevalence of switching and multi-cloud

3.413 We have considered how common it is for customers to switch or multi-cloud.

3.414 The level of switching or use of multiple clouds is unlikely to be probative on its own of the strength of competition. In particular, low levels of switching or limited use of multiple cloud providers may be consistent with at least two different scenarios:

⁶³⁸ See appendix G, section 'R&D Investment'.

⁶³⁹ [Jigsaw Report](#), page 25.

- (a) First, low levels of switching or use of multiple cloud providers may be consistent with a lack of ability or incentive for customers to switch or multi-cloud, and therefore with weaker competition.
- (b) Second, a lack of switching or multi-cloud could be consistent with suppliers responding proactively to the credible threat of switching or multi-cloud from their customers with competitive prices and levels of quality.

- 3.415 Because of this, evidence on rates of switching and multi-cloud must be interpreted alongside other evidence, such as evidence on the ease of switching and on the broader market context (such as profitability). However, we would note that if the level of switching and use of multiple cloud providers is very low, this is less likely to be consistent with customers being able to make a credible threat to switch or multi-cloud.
- 3.416 Microsoft and AWS both submitted quantitative surveys as evidence of the prevalence of switching and multi-cloud. We consider results relevant to each of switching and multi-cloud below in the relevant sections. Below, we considered the overall evidential value we can attribute to these quantitative surveys.⁶⁴⁰

Quantitative surveys on the prevalence of switching and multi-cloud

- 3.417 In general, quantitative surveys have some advantages: for example, if the sample is representative, the results can be generalised to the whole population (although this is not the only condition for being able to rely on a quantitative survey). This means that, in some cases, they can be an important part of our evidence base. The CMA can use quantitative surveys in its inquiries and often draws important insights from them.⁶⁴¹
- 3.418 However, we have concerns about using surveys to assess the prevalence of switching and multi-cloud in the cloud market:
- (a) Ambiguity in terminology: quantitative surveys require customers to make judgements about category responses (ie the customers spoken to must choose one or more options in a given list), but there is little room to follow up or clarify, possibly affecting the validity of their responses.⁶⁴²

⁶⁴⁰ A more detailed analysis of these surveys can be found in the appendix C.

⁶⁴¹ The CMA's good practice on customer surveys for its casework can be found at [CMA78, Good practice in the design and presentation of customer survey evidence in merger cases \(revised\)](#).

⁶⁴² For example, as identified in appendix C, the Public First survey asked 'How many different cloud infrastructure providers does your company currently use?'. However, we consider that the customers spoken to may have interpreted this in as including any of the following:

- using multiple cloud providers for the underlying infrastructure layer;
- using multiple first and third-party providers hosted on the same cloud;
- using both public and private cloud providers for infrastructure; and/or
- using multiple private cloud providers.

For the purposes of estimating the prevalence of multi-cloud in this market, we are only interested in the first category.

- (b) Vulnerability to the type and coverage of the customers spoken to: results from quantitative surveys are crucially dependent on the representativeness of the sample and likely response rate. We have concerns about both in this market – see appendix C. Further, the quality and accuracy of customer record-keeping within public cloud providers is highly variable, meaning their use as a sample frame for a quantitative survey becomes challenging. Alternative robust sample frames for this target population do not appear to us to be available.⁶⁴³

3.419 For these reasons, there are significant weaknesses with using a quantitative research method to estimate the prevalence of switching and multi-cloud in this market. We have therefore placed more weight on the in-depth qualitative research we commissioned ourselves than the quantitative survey evidence submitted to us by cloud providers. We chose to use qualitative research methods for our customer research in this investigation as these are better suited to investigating our key research questions and navigating some of the complexities of customer behaviour in this technical market.⁶⁴⁴

Prevalence of switching

3.420 In this section, we considered evidence on the prevalence of switching. In particular, we set out:

- (a) Cloud provider views;
- (b) Quantitative analysis of switching submitted by one provider;
- (c) Quantitative surveys;
- (d) Customer views;
- (e) Our own analysis of switching between Microsoft, AWS and Google.

Cloud provider views

3.421 AWS said that, due to the way cloud services are designed and priced (pay-as-you-go) compared to alternative IT environments, it had never been easier for

⁶⁴³ Sample bias is also a concern when the customers spoken to are drawn from a panel, in particular from an online panel, where sample recruitment does not rely on randomisation methods. Whilst a panel can be made to look like a random, representative cross-section of consumers in terms of its demographic profile, the characteristics of people who join a panel may be very different from other consumers. The CMA tends to place less evidential weight on surveys involving customer recruitment from panels, though each case is treated on its individual merits. See paragraph 2.29 in [CMA78, Good practice in the design and presentation of customer survey evidence in merger cases \(revised\)](#).

⁶⁴⁴ We have presented the results of these surveys in appendix C. We have also set out supplemental points that apply to the individual surveys and, as a general point, we note that the survey providers for all these surveys did not provide their sampling and methodology, and we do not have access to the underlying data. As a result, we cannot assess the representativeness of the samples.

customers to switch IT provider.⁶⁴⁵ AWS also highlighted its own migration programmes as well as the transfer services of rival cloud providers and other IT companies that offer migration services (eg Accenture, BMC, Capgemini and Deloitte).⁶⁴⁶

- 3.422 Microsoft submitted that it is incentivised to make it as easy as possible for customers to switch to Microsoft (in particular, from AWS) or to multi-cloud as customers focus on diversifying beyond AWS.⁶⁴⁷
- 3.423 Google said it is deeply invested in making sure that customers can use its cloud infrastructure in combination with other third-party cloud services of their choice (in particular AWS and Microsoft as they are usually the primary cloud provider).⁶⁴⁸
- 3.424 Microsoft said that low levels of switching are not necessarily evidence of weak competition as it may suggest that the gains from switching do not outweigh the costs of doing so.⁶⁴⁹ It also said that:
- (a) The CMA's music streaming market study showed that positive outcomes can arise notwithstanding limited switching and some barriers to switching.⁶⁵⁰
 - (b) The market being considered is different from others the CMA has considered as customers are sophisticated, the cloud market is not mature and has a large number of new customers and it is not a digital multi-sided platform with network effects.⁶⁵¹
- 3.425 AWS and Google identified practices that restrict customer choice and make switching more difficult such as certain licensing practices.⁶⁵²
- 3.426 Microsoft said that, given the prevalence of multi-cloud (which we discuss below), it is much less common for customers to switch fully than it is for them to scale up consumption of services on one cloud, which may or may not also result in scaling down of consumption on another. It said the key indication of whether it is 'losing business' to a competitor is when a customer's consumption on the provider's cloud decreases (or stalls/increases at a slower rate than expected).⁶⁵³

⁶⁴⁵ [AWS' response to the Issues Statement](#), 17 October 2023, paragraph 17.

⁶⁴⁶ [AWS' response to the Issues Statement](#), 17 October 2023, paragraph 14.

⁶⁴⁷ [Microsoft's response to the Issues Statement](#), 17 October 2023, paragraph 26.

⁶⁴⁸ Google's response to the CMA's information request [§].

⁶⁴⁹ Microsoft's submission to Ofcom [§]. In addition to the above, an academic submitted that limited switching behaviour alone does not provide strong evidence of limited substitution possibilities. See Dr George R Barker response to the Provisional Decision Report, page 62-63.

⁶⁵⁰ Microsoft's submission to Ofcom [§].

⁶⁵¹ Microsoft's submission to Ofcom [§].

⁶⁵² [AWS' response to the Issues Statement](#), 17 October 2023, paragraph 33; Google, submission to the CMA [§].

⁶⁵³ Microsoft's response to the CMA's information request [§].

- 3.427 Further, in its response to our working papers, Microsoft submitted that the customer evidence suggests that customers are generally satisfied with their cloud provider but would switch if there were a reason to.⁶⁵⁴
- 3.428 However, the three main cloud providers were able to provide only limited data on switching:
- (a) AWS provided the number of UK customers that decreased their total annual spend across all of its services by at least 25%, 50%, 75% and 100% on a year-on-year basis from 2018 to 2022. This analysis implied a low level of switching in terms of revenue – for example, customers that decreased spend by at least 25% in 2022 made up roughly [REDACTED]. Based on our analysis, this is about [5-10]% of AWS' UK revenue in 2022.⁶⁵⁵ AWS said that, while this does not necessarily indicate customer switching, it implies a reduction in spend or a halt in the use of some workloads on its cloud which may constitute switching. AWS added that the analysis would not capture customers simultaneously switching specific workloads from its cloud and other workloads to its cloud, and combined with the increased customer needs over time, this analysis likely understates switching.⁶⁵⁶
 - (b) Google gave an estimate of the number of customers that were not billed in a given year, but billed the year before, over the years 2019-2022. Google stated that this method does not conclusively evidence customer switching, as there are several reasons why customers might record revenues in one year but not the next, including testing its services, temporarily pausing use and the timing of invoice. Based on this method, Google submitted that it lost around [REDACTED] of its UK Google Cloud Platform customers in 2022, altogether accounting for [REDACTED] in revenue in 2021.⁶⁵⁷ Based on our analysis this amounts to [0-5]% of GCP's total UK revenue in 2021.⁶⁵⁸
 - (c) Google also provided its lost global customers in 2022 (based on a customer whose stored data had essentially reduced to zero by December 2022). In total [REDACTED] Google, but there are limitations to what we can take from this given that [REDACTED].⁶⁵⁹ [REDACTED]. Finally, we do not have information on the size of these customers, which makes it hard to interpret this figure in context.
- 3.429 We agree that switching rates should be considered alongside other evidence on wider competitive conditions. That said, we consider it reasonable that switching

⁶⁵⁴ Microsoft's response to the Competitive landscape, Committed spend agreements and Egress fees working papers, paragraph 17.

⁶⁵⁵ CMA analysis of AWS' response to the CMA's information request [REDACTED].

⁶⁵⁶ AWS' response to the CMA's information request [REDACTED].

⁶⁵⁷ Google's responses to the CMA's information request [REDACTED].

⁶⁵⁸ Given that Google's total UK revenue in 2021 was [REDACTED]. CMA analysis of Google's response to the CMA's information request [REDACTED].

⁶⁵⁹ [REDACTED] submission to the CMA [REDACTED].

rates may be correlated with how easy or hard it is to switch, and low switching rates may be associated with some switching barriers, especially if other indicators like profitability or market shares are consistent with some cloud providers holding a significant degree of market power. We consider this alongside customer evidence.

- 3.430 We therefore consider analysis of switching rates in conjunction with other evidence in order to reach our conclusion.

Quantitative analysis from one cloud provider

- 3.431 One cloud provider submitted quantitative analysis of its customer data and said that the results show that customers can and do switch.⁶⁶⁰ According to this analysis, approximately [X] of its customers churned on an annual basis, considering all the cloud services it offered. The provider said that [X] of its customers, accounting for [X] of its revenues, [X] their spend between the first half of 2020 and the first half of 2022. The provider said that these figures indicate that customers can move significant workloads (therefore reducing their spend on the provider's cloud) from that provider's cloud to other cloud providers, which would be evidence of significant switching.
- 3.432 There are limitations to this quantitative analysis on switching.⁶⁶¹ These are related to the influence of very small customers on the churning rates (eg removing customers that spent less than \$500 reduces the churn rate to [X]% – our own analysis excludes customers that spent less than \$1,000), the much lower churning rates for customers with high spend, and limitations as to how customers are identified as having churned. We further note that smaller reductions of spend might be brought about by customers moving to multi-cloud strategies, rather than switching. The results from this quantitative analysis have been interpreted in light of these caveats.

Quantitative surveys

- 3.433 Microsoft and AWS submitted a quantitative survey (Public First) that included an estimate of the prevalence of switching in the market.⁶⁶² In addition, the prevalence of switching was considered as part of Ofcom's quantitative survey. There are also publicly available estimates of the prevalence of switching in the market.⁶⁶³
- 3.434 As discussed above and in appendix C, there are several limitations to these estimates. Even when not considering these methodological limitations, the Public

⁶⁶⁰ [X] response to the CMA's information request [X]; Ofcom's analysis in its Final Report, Annex 3, paragraph A3.37

⁶⁶¹ We include our full assessment of this analysis in appendix H.

⁶⁶² AWS' submission to the CMA [X]; Microsoft's submission to the CMA [X].

⁶⁶³ Our views on quantitative research in cloud services are set out in appendix C.

First survey referenced by Microsoft and AWS still offers only limited insight into the extent of switching by customers between public cloud providers.

- 3.435 For example, the Public First survey asked: 'Have you ever switched one of your cloud infrastructure providers in the past?'⁶⁶⁴ As phrased, the question is not informative of the proportion of workloads that were switched from one cloud provider to another. Customers spoken to could have answered 'yes' even if they had switched only a handful of workloads. This also gives rise to the possibility of the customers spoken to conflating switching with multi-cloud, which adds further uncertainty to any inference that can be drawn from these surveys on the prevalence of customer switching.
- 3.436 We have considered some of the above limitations of these quantitative estimates when considering these findings alongside other qualitative evidence and our own analysis in our conclusions.
- 3.437 The research commissioned by Ofcom included both a qualitative research phase and a quantitative online survey.⁶⁶⁵ The quantitative online survey has the same limitations of validity and sample quality outlined above. The qualitative research is not subject to the same methodological limitations.⁶⁶⁶
- 3.438 We note that Ofcom's qualitative results also indicate that switching is lower than the results reported by Ofcom's quantitative survey. In particular, the former found:
- (a) in some cases, firms were adding additional platforms, rather than switching;
 - (b) few examples of organisations switching away from Microsoft, AWS or Google to another cloud provider;
 - (c) the switching that was described by customers tended to involve a relatively small portion of data and workloads, moving from one minority cloud provider to another;
 - (d) it is still relatively early in the adoption journey for most companies, and they were evaluating progress rather than looking to make significant changes; and

⁶⁶⁴ Public First Poll [38].

⁶⁶⁵ [Cloud Services Market Research - Summary of Findings March 2023](#).

⁶⁶⁶ Context Consulting conducted 50 hour-long depth interviews with a further 14 follow-up interviews with current and potential cloud customers in their qualitative research. Qualitative research methods such as in-depth interviews allow for more discussion, clarification and explanation of customers behaviour compared to quantitative research methods, which use more structured research instruments such as questionnaires that seek to measure very defined category responses. Qualitative research methods can work well in complex and technical markets such as public cloud infrastructure services. [Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw \(2024\) para 1.2](#) provides further detail on the qualitative research approach the CMA commissioned for its customer research on this investigation.

- (e) in most cases, firms were still on their way into, not out, of their IaaS/PaaS environments.

3.439 As above, we take into account the limitations of this quantitative research in our conclusions below.

Customer views

- 3.440 The Jigsaw report states that few of the customers interviewed had chosen to switch between public cloud providers.⁶⁶⁷ It states that, overall, switching not only brought cost and operational risk, but took IT staff away from the customer's core work and typically ended up being more challenging and time consuming than anticipated.⁶⁶⁸ The report also notes that switching cloud providers is seen as the equivalent of moving other kinds of infrastructure, such as 'moving house' or moving a business from one country to another. It is not something to undertake lightly or consider at all unless it leads to significant business benefits long term that override the inherent cost and risk of changing. To an extent, dependency on a current provider(s), or a sense of 'lock in', is a factor across all providers as change brings cost and risk.⁶⁶⁹
- 3.441 AWS submitted that the Jigsaw report underestimates switching by ignoring switching to on-premises solutions and switching to (or from) cloud providers outside the top five UK cloud providers.⁶⁷⁰
- 3.442 We accept that some switching back to on-premises solutions may be due to genuine competition, but we also consider that some switching back to on-premises solutions may reflect that cloud or on-premises solutions are not suitable for some specific workloads. We discussed this when we looked at customer evidence on the suitability of alternative IT solutions and concluded that these pose only a weak competitive constraint on public cloud providers.⁶⁷¹ In any case, we have interpreted the findings from the Jigsaw report as relating to switching among public cloud providers, rather than switching across all markets.⁶⁷²
- 3.443 We further note that our qualitative research approach enabled participants to discuss switching to and/or from alternative IT solutions, even though the initial question focused on public cloud. Additionally, while the research focused on customers of the five largest cloud providers, nothing precluded customers from

⁶⁶⁷ [Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw \(2024\)](#), paragraph 1.3.14

⁶⁶⁸ [Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw \(2024\)](#), paragraph 1.3.16

⁶⁶⁹ [Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw \(2024\)](#), paragraph 1.3.13.

⁶⁷⁰ AWS' submission to the CMA [38].

⁶⁷¹ See Section Market Definition.

⁶⁷² The discussion guide in the Jigsaw report makes clear reference to switching between (public) cloud providers and specifies that switching does not include moving from a non-public cloud IT environment onto the public cloud or vice versa. See [Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw \(2024\)](#), pages 115-116.

discussing situations where they have moved from a smaller cloud provider or planned to move into a smaller cloud provider.

3.444 We also collected customer evidence on switching as part of our assessment of technical barriers and committed spend agreements.

3.445 As discussed in the next chapter, Barriers to switching and multi-cloud, the evidence we have seen shows that many customers anticipate or experience significant technical costs to switch public clouds.⁶⁷³ Customers described the costs as significant either in absolute terms – eg a customer said ‘it would take 12 months and tie up approximately 1,000 employees’⁶⁷⁴ – or in relative terms, eg some customers described technical barriers as the main barrier to switching.⁶⁷⁵ In particular:

- (a) Some of these customers indicated that these costs had stopped them from switching or considering switching.⁶⁷⁶
- (b) Some customers also said that there are a range of operational concerns that arise when they consider switching between public clouds.⁶⁷⁷ This is consistent with the findings of the Jigsaw report.⁶⁷⁸
- (c) Some customers said that, given the similarity of the current offerings by cloud providers from their perspective, the value of switching is low in comparison to the costs.⁶⁷⁹ This is consistent with the findings of the Jigsaw report.⁶⁸⁰

3.446 This last point was also corroborated by other customer evidence. Some customers and other market players (eg professional services suppliers) viewed Microsoft, AWS and Google as having broadly equivalent offerings in terms of products, features and prices.⁶⁸¹ For example:

- (a) A customer said that the capability gap between the ‘three main hyperscalers’ is much reduced now compared to five years ago, and there is little to choose between them outside of some speciality areas and niche use cases.⁶⁸²

⁶⁷³ Responses to the CMA’s information requests issued to [REDACTED]. Notes of meetings with [REDACTED].

⁶⁷⁴ Note of meeting with [REDACTED].

⁶⁷⁵ Responses to the CMA’s information requests [REDACTED]; Notes of meetings with [REDACTED].

⁶⁷⁶ Notes of meetings with [REDACTED].

⁶⁷⁷ Responses to the CMA’s information requests [REDACTED].

⁶⁷⁸ [Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw \(2024\)](#), paragraph 4.2.1

⁶⁷⁹ Responses to the CMA’s information requests [REDACTED]; Note of meeting with [REDACTED].

⁶⁸⁰ [Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw \(2024\)](#).

⁶⁸¹ Responses to the CMA’s information requests [REDACTED]; Notes of meetings with [REDACTED].

⁶⁸² [REDACTED] response to the CMA’s information request [REDACTED].

(b) Other customers said that any innovations in one cloud provider's offerings are matched quickly by the others.⁶⁸³

(c) Another customer said that there are differences in functionality between IaaS/PaaS on different clouds, but the question is whether they are significant enough to switch. It said that for IaaS, the differences are not significant enough and for PaaS it comes down to developer preferences, which largely come from which platform they are familiar with using.⁶⁸⁴

3.447 Where customers mentioned differences in the overall offering of public cloud providers, they said that these are currently relatively minor or cover edge cases.⁶⁸⁵

3.448 Finally, we asked large customers to list any cloud services they were getting from their main cloud provider which they would not be willing to switch to alternative cloud providers, including the proportion of spend that these services accounted for.⁶⁸⁶ This evidence gives an indication of the share of existing demand for Microsoft's and AWS' services that is 'sticky' for this specific group of customers, that is demand over which customers cannot exercise effective choice due to lack of suitable alternatives or barriers to switching.

3.449 Based on our analysis, there was significant variation in customer responses to this question, with the proportion of Microsoft's and AWS' services that is sticky being large for some customers and smaller for others. The proportion was generally higher for Microsoft's services than it was for AWS'.⁶⁸⁷

Our analysis of Microsoft, AWS and Google customer switching

3.450 We conducted data analysis to measure the prevalence of switching amongst customers of Microsoft, AWS and Google.

3.451 Microsoft, AWS and Google provided customer data sets that identified customer names and annual spend on their respective clouds in each year from 2020 to 2023. We matched these data sets to identify customers that are likely to have fully switched from one of these three main cloud providers to another one of these three providers.

3.452 In our baseline results, we identified a customer as having switched if:

⁶⁸³ Note of meeting with [REDACTED]; [REDACTED] response to the CMA's information request [REDACTED].

⁶⁸⁴ Note of meeting with [REDACTED].

⁶⁸⁵ Notes of meetings with [REDACTED]; Responses to the CMA's information requests [REDACTED].

⁶⁸⁶ This evidence is described more in detail in the CSAs chapter.

⁶⁸⁷ Responses to the CMA's information requests [REDACTED].

- (a) they spent more than \$1,000 a year on a cloud provider – otherwise the customer does not appear in all our data sets. This is to discard very small customers and keep our analysis manageable;
- (b) their spend on one cloud provider drops by more than 85% year-on-year; and
- (c) at least 60% of this drop in spend appears as spend increase on another cloud provider.

3.453 The data, methodology, results and limitations of our analysis are set out in detail in appendix H. Here, we highlight some key considerations that we account for when interpreting the results of our analysis:

- (a) We look only at full switching (defined as customers moving at least 85% of their spend from one cloud). Partial switching would be impossible to distinguish from natural fluctuations of customers' spend on cloud.
- (b) Our analysis is limited to the three main cloud providers: Microsoft, AWS and Google. Therefore, we are not able to identify customers who switched to or from other smaller cloud providers. Nonetheless, these three providers cover a large majority of the market, so the amount of switching not captured due to this limitation is likely to be limited.
- (c) Customers are matched using the names in the providers' data set. Inevitably, some customers might not have been matched when they should have, or they might have been matched when they should not have been. We took reasonable steps to minimise these sources of errors, eg conducting sensitivity checks to determine the appropriate threshold for our fuzzy matching 'similarity score' and randomly spot-checking matches.⁶⁸⁸
- (d) Identifying who has switched requires some simplistic assumptions. We identified customers switching based on different thresholds of their year-on-year fall in spend on one provider (similar to the AWS analysis, above) and how much of that spend fall increases in the following year on another provider's data set. To mitigate the risk that the choice of our thresholds drives the results, we performed several sensitivities: our results are not sensitive to most of the thresholds chosen.⁶⁸⁹

3.454 In response to our PDR, cloud providers provided views on our analysis.⁶⁹⁰ We have taken account of the points raised their submissions in our analysis, where appropriate. Below, we discuss some additional points:

⁶⁸⁸ For further detail, see appendix I.

⁶⁸⁹ For full details see appendix H.

⁶⁹⁰ AWS' submission to the CMA [X]; Microsoft's submission to the CMA [X].

- (a) AWS submitted an analysis based on whether a customer has reduced spend by a certain threshold, regardless of whether it has increased elsewhere. AWS submitted that this analysis shows that switching rates increase by more than [REDACTED], unweighted by cloud spend, for the period from 2020 to 2022.⁶⁹¹ However, a reduction in cloud spend only captures customers that have left a cloud (or reduced their spend on it significantly), not necessarily those that have switched to another cloud provider. For example, the following developments would be incorrectly classified as switching:
- (i) A customer becoming more efficient in its cloud usage and reducing its spend, or;
 - (ii) A customer dropping entirely a workload on the cloud.
- (b) AWS submitted that full switching is not a complete metric for assessing how well competition works in the cloud sector as competition takes place at the workload level and not for entire customers. It submitted an analysis that shows when reducing the CMA's 85% spend reduction condition to lower thresholds between 5-80%, the proportion of customers switching between the three main cloud providers increases from 1.5% (85%) to 3.7% (5%).⁶⁹² Here, we note that as the threshold is reduced, at no point does the proportion of customers switching increase rapidly. We consider this shows that our original threshold is robust and appropriate. We further note that at smaller levels of reduction in spend, the analysis increasingly risks picking up smaller fluctuations in usage that are not switching, so is more likely to overestimate the prevalence of switching.
- (c) AWS submitted that, given less complex IT infrastructures and lower spending, the smallest customers should find it easier to overcome any inherent technical difficulties to switching.⁶⁹³ AWS also submitted that most of the 'full' switching occurs with smaller customers, which is consistent with larger customers opting for multi-cloud strategies instead.⁶⁹⁴ We consider that this may equally indicate that larger customers find it harder to switch due to various technical or commercial barriers, rather than necessarily indicating a tendency to multi-cloud instead. We consider the overall findings on the prevalence of switching and multi-cloud below.

3.455 Our analysis suggests that switching between public cloud providers is very uncommon in the cloud markets: about 0.1% of customers by count, and a similar

⁶⁹¹ AWS' submission to the CMA [REDACTED].

⁶⁹² AWS' submission to the CMA [REDACTED].

⁶⁹³ AWS' submission to the CMA [REDACTED].

⁶⁹⁴ AWS' submission to the CMA [REDACTED].

proportion of revenue, switched cloud providers in any given year over the period 2020-2023. The full results of our analysis are outlined in appendix H.

- 3.456 We also looked at the size of those customers we identified as switchers between cloud providers. In particular, we note that, in two of the three years assessed, there were very few customers that fully switched and the unweighted number of switchers is higher than their weighted (by spend on cloud) counterpart, which suggests that most switchers have smaller than average cloud expenditures. Our analysis suggests that among those very few customers that fully switched, customers with an annual spend of around \$1,000 and \$5,000 make up the majority.
- 3.457 This suggests that full switching between public cloud providers is extremely rare. While cloud providers said that low switching rates reflect that customers are satisfied with their providers, we consider such a low level of switching together with high levels of profitability among the three main cloud providers to be consistent with the presence of high barriers to switching. This is supported by customers' views. This indicates that switching costs may outweigh the benefits of changing cloud provider for many customers.

Prevalence of multi-cloud

- 3.458 This section sets out the evidence on the prevalence of multi-cloud use among customers. We have considered the following:
- (a) cloud provider views,
 - (b) AWS' quantitative analysis;
 - (c) publicly available survey data and the results of a survey that Ofcom commissioned during its market study;⁶⁹⁵ and
 - (d) our analysis of customer data provided by cloud providers.

Cloud provider views

- 3.459 Cloud providers generally submitted that they consider that multi-cloud architectures are common and that enabling customers to multi-cloud is part of their business strategy. Some cloud providers [X] said that using multiple clouds is particularly prevalent among large customers. Microsoft submitted that customers multi-cloud 'when it works for them'.⁶⁹⁶ Microsoft also submitted that industry studies confirm that multi-cloud is the new norm when appropriate. It also

⁶⁹⁵ For more information the research that Ofcom commissioned, see [Cloud Services Market Research - Summary of Findings March 2023](#).

⁶⁹⁶ [Microsoft's response to the CMA's Competitive landscape, Committed spend agreements and Egress fees working papers](#), paragraph 3.

said that the working papers suggest only one kind of integrated multi-cloud shows a well-functioning market, even though customers disagree.⁶⁹⁷

- 3.460 Another cloud provider submitted that the use of multi-cloud is widespread.⁶⁹⁸ But this is not consistent with an internal document which states, in the context of launching a new product, that it was not beginning a multi-cloud support strategy. The document said that key decision-makers in private and public companies tended to choose one single cloud provider. In particular, the document states that, in relation to [REDACTED].⁶⁹⁹ This suggests that even large customers tend to use one main cloud provider.
- 3.461 In relation to the type of multi-cloud used by customers, Microsoft said that the lack of use of 'integrated' multi-cloud is because it does not yield significant customer benefits and that there may be good reasons why customers concentrate their spend around a primary cloud provider and/or to run different workloads separately in different clouds.⁷⁰⁰
- 3.462 Our view is that customers' decisions on whether to multi-cloud or not—as well as the level of integration to adopt in doing so—are driven by how they weigh the benefits and disadvantages of adopting multi-cloud setups. However, as noted in the case of switching, low levels of multi-cloud could be consistent with a lack of ability or incentive for customers to multi-cloud driven in part by high costs of multi-clouding (or costs that exceed the benefits of doing so), which may in turn result in weaker competition.
- 3.463 AWS and Google pointed to the use of specific cloud services as an indicator of customers adopting multi-cloud:
- (a) AWS explained that customers can manage their users in another on-premises or cloud directory and then connect them into that provider's cloud through the provider's Identity Access Management (IAM) solution.⁷⁰¹
 - (b) Google said that the fact that a customer is using BigQuery Omni can give an indication that the customer is deploying a multi-cloud strategy.⁷⁰²
- 3.464 The rate of use of these services can indicate the prevalence of multi-cloud in the market but, unlike our approach, it does not give us information on what proportion of workloads is run on each cloud.

⁶⁹⁷ [Microsoft Response to the Competitive landscape, Committed spend agreements and Egress fees working papers](#), paragraph 17.

⁶⁹⁸ [REDACTED] response to the CMA's information request [REDACTED].

⁶⁹⁹ [REDACTED] response to the CMA's information request [REDACTED].

⁷⁰⁰ [Microsoft's response to the Issues Statement](#), paragraph 40.

⁷⁰¹ AWS' submission to CMA [REDACTED].

⁷⁰² Google's response to the CMA's information request [REDACTED].

- 3.465 Some cloud providers have suggested that customers' growing interest in accessing FMs may lead to a higher use of multi-cloud architectures. Conversely, others have suggested that there may be barriers to adopting a multi-cloud strategy for accessing FMs.⁷⁰³
- 3.466 [REDACTED] and [REDACTED] highlighted surveys that sought to estimate prevalence of multi-cloud, and we considered these surveys below. One cloud provider also submitted a quantitative analysis which we consider briefly below, and in detail in the appendix I.

Quantitative analysis from a cloud provider

- 3.467 One cloud provider submitted a quantitative analysis of its tender data to illustrate that customers do not view themselves as 'locked in' to their incumbent cloud provider. This analysis shows that a majority of the tenders this cloud provider participated in from 2009 to 2021 were issued by existing customers. The cloud provider also said that its win rate in tenders for customers with existing workloads in its cloud is [REDACTED] its win rate for other customers. It said that this shows it does not enjoy a significant advantage as an incumbent cloud provider.⁷⁰⁴
- 3.468 The same cloud provider also submitted analysis on the distribution of revenue share of customers in its opportunity data by the number of cloud providers the customers awarded tenders to between 2018 and 2022.⁷⁰⁵ This analysis showed that many of its customers [REDACTED] used the cloud provider after awarding tenders to at least one other cloud provider between 2018 and 2022. The cloud provider submitted that this is an indication that these customers were using multiple clouds.⁷⁰⁶
- 3.469 The cloud provider said that its analysis likely understates the prevalence of multi-cloud because:⁷⁰⁷
- (a) many customers acquire IT services without a tender process. These customers would not be recorded in the opportunity data set;
 - (b) the cloud provider did not participate in all tenders issued by customers;
 - (c) it is often not clear who won the tender. Conservatively, the analysis only flags customers as having awarded a tender elsewhere if the cloud provider knows the identity of the other competitor that won the tender; and

⁷⁰³ See chapter 4, section 'Impact of cloud customers use of FMs on switching and multi-cloud'.

⁷⁰⁴ [REDACTED] submission to CMA [REDACTED].

⁷⁰⁵ [REDACTED]. [REDACTED] response to Ofcom's information request [REDACTED].

⁷⁰⁶ [REDACTED] response to the CMA's information request [REDACTED].

⁷⁰⁷ [REDACTED] submission to CMA [REDACTED].

- (d) some customers may have awarded tenders before or after the sample period.

3.470 The analysis has several conceptual and technical limitations: sample selection bias, arising because the analysis inherently excludes customers that are locked in and do not go to tender including through customer inertia (ie lack of credible options to tender to), wherein a customer may consider its current cloud provider the only credible option; and limited market coverage. We understand that the majority of customers do not use a competitive tender process to acquire cloud services. Tenders are rarely used outside of the public sector, so this analysis is only concerned with the behaviour of a specific subset of customers. In light of these caveats, we interpret this analysis with care and⁷⁰⁸ place less weight on it than our own assessment, set out below.

Quantitative surveys on the prevalence of multi-cloud

3.471 Microsoft and AWS submitted that independent surveys and industry reports show that using multiple clouds is common.⁷⁰⁹

3.472 We discussed the general limitations we found in the use of quantitative surveys to estimate the prevalence of switching in the section above and appendix C. The same applies to estimating the prevalence of multi-cloud.

3.473 Further, we do not consider this evidence to be fully informative of the extent of multi-cloud among customers because it provides no information on the split of workloads deployed on different public clouds by customers that multi-cloud. We think it is important to consider not only the proportion of customers that multi-cloud, but also the extent to which they use both clouds when they multi-cloud. Other pieces of evidence, eg our own analysis, can tell us more about this than these quantitative surveys do.

3.474 Our own analysis goes further than just assessing the prevalence of multi-cloud at a customer level, as a quantitative survey otherwise would. Our analysis complements customer-level findings with a breakdown of customers' spend across cloud providers. As explained below, this shows that when multi-cloud occurs, most of the spend (75%-80%) tends to concentrate on one cloud provider only. Our own analysis also explores the extent to which customers split their spend across different clouds.

⁷⁰⁸ We explore these limitations further in appendix I.

⁷⁰⁹ [redacted] submission to the CMA [redacted]; Microsoft's response to the CMA's information request [redacted]; AWS' submission to the CMA [redacted]; Microsoft's submission to the CMA [redacted].

Customer views

- 3.475 The Jigsaw report notes that, amongst those customers that participated, single-cloud and siloed multi-cloud (where a customer runs multiple clouds that do not interact with one another) are the main operating models. Other operating models were found to be less common.⁷¹⁰
- 3.476 Only a few participants were described as having models of multi-cloud that were not siloed:⁷¹¹
- (a) Mirroring: There were no cases of cloud architecture being simply duplicated on two or more public clouds, but there were some examples of plans being in place to easily transfer to another service in the case of significant change in their business need or cloud provider failure eg through use of ‘open-source’ rather than provider’s proprietary services to enable easier portability if needed, use of tools that help define infrastructure as code, or actively maintaining disaster recovery plans and/or low-level relationships with other cloud providers (sometimes as a regulatory requirement).
 - (b) Integrated: There were a few specific cases of choosing to integrate services across public cloud providers. The main reason for this was the perceived performance benefits, with the most mentioned examples being use of BigQuery from Google or Microsoft IAM alongside infrastructure from a different public cloud provider as these services were superior to those offered by their main cloud provider. Most businesses avoided integrating across public cloud providers as this presented unnecessary challenges.
 - (c) Commodity use: There were a few examples of firms using five or more public cloud providers at once. This was another form of siloed use, but where the behaviours were price-driven, sometimes temporary or for short periods and considered separately to their main cloud provider(s). The two kinds of use uncovered were start-ups using services because they were free or heavily discounted via start-up credits, even though they did not intend to continue to use these suppliers once credits ended, or companies ‘cloud bursting’ (using public cloud only when their private data centre was at capacity, allocating work automatically based on cloud provider availability and price for uses like machine learning).
- 3.477 Large customers told us that they reviewed the option to integrate public clouds but concluded that the benefits did not outweigh the technical costs of doing so for their current use cases.⁷¹²

⁷¹⁰ [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), paragraph 3.5.2.

⁷¹¹ [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), paragraph 3.5.5.

⁷¹² Responses to the CMA’s information requests [§]; Note of meeting with [§].

3.478 Other customers said that they viewed the benefits of integrating multiple clouds as being too low, but did not mention whether this was in comparison to the technical costs.⁷¹³

3.479 The Jigsaw report highlighted that, overall, there is a preference towards using as few public cloud providers as possible. Many favour the simplicity of one provider that covers their business needs.⁷¹⁴

Our analysis of multi-cloud across Microsoft, AWS and Google

3.480 We have built our own estimate of multi-cloud prevalence using customer data from Microsoft, AWS and Google, that identified customer names and annual spend on their respective clouds for 2020-2023.⁷¹⁵ By analysing these data sets to identify customers using multiple cloud providers, we avoid any potential issue of customers misunderstanding what it is to multi-cloud, as we define it for the purposes of the investigation. For example, customers using both private cloud and public cloud would not be counted as using multiple clouds in our analysis, but such customers may have responded in surveys that they use multiple clouds.

3.481 We highlight some key considerations that we account for when interpreting the results of our analysis:⁷¹⁶

- (a) Our method counts customers as using multiple clouds in a binary manner: customers are counted as using multiple clouds if they spend over \$1,000 on another cloud, irrespective of the size of that workload. This method may identify some customers as adopting multi-cloud even in cases where they have limited usage of their secondary cloud.
- (b) We look only at customers who spent more than \$1,000 a year on a cloud provider. Customers with lesser levels of spend do not appear in our data. This excludes very small customers and keeps our analysis manageable.
- (c) We only match customers from Microsoft, AWS and Google. This is due to limitations to the data available. Nonetheless, we have incorporated in our results an estimate for the prevalence of multi-cloud for customers using other cloud providers too (ie 'scaled up'): we assumed that 50% of customers (by cloud spend) who do not use Microsoft, AWS and Google (ie they are

⁷¹³ Responses to the CMA's information requests [X].

⁷¹⁴ [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), paragraph 3.6.4.

⁷¹⁵ AWS submitted that the revenue data sets from AWS, Microsoft, and Google include total customer revenues including SaaS revenues. We have removed SaaS revenues from AWS' customer revenues here and elsewhere based on our own categorisation of its services into IaaS, PaaS and SaaS. See AWS' submission to the CMA on multi-cloud and switching dated 16 April 2025, paragraph 48.

⁷¹⁶ The data, methodology and limitations of our analysis are fully considered in appendix I.

with other cloud providers) use multi-cloud.⁷¹⁷ We then added these customers to our prevalence of multi-cloud statistics to get an estimate of the share of revenue in the wider market that is using multi-cloud.⁷¹⁸

- (d) Customers are matched across data sets using an approximation technique called fuzzy matching which allows for small variations in the way customers' names are recorded.⁷¹⁹ This can lead to false positives (two different customers on two different data sets being matched because their names are similar) or false negatives (the same customer on two different data sets not being matched, for example because the customer's name is recorded in a very different way).

3.482 In response to our working papers and PDR, cloud providers provided views on our analysis.⁷²⁰ We have incorporated some of their submissions in our analysis, where appropriate. Below, we discuss some additional points:

- (a) AWS submitted revised figures for prevalence of multi-cloud based on our analysis.⁷²¹ However, these estimates include revenues from SaaS and are therefore based on a different market definition than the one adopted by us (see the section on market definition). This has the effect of AWS' revised analysis overestimating the prevalence of multi-cloud. These estimates also exclude the smallest customers (ie those with annual spend under USD 10,000), which AWS submitted are unlikely to multi-cloud for economic reasons.⁷²² Our analysis, detailed below, looks at the prevalence of multi-cloud by spend band and includes most of these customers (those with annual spend over USD 1,000). This way, we can see how these smaller customers behave separately from the wider customer group.
- (b) AWS submitted that our claims of low sample sizes for higher-bucket-spend customers are also unjustified, and that results for larger customers and smaller customers are equally reliable.⁷²³ We consider that, in any analysis, small sample sizes are more sensitive to the actions of individuals within, and therefore at greater risk of not being representative of a wider population over

⁷¹⁷ We consider 50% a reasonable assumption as evidence we have seen suggests that smaller cloud providers are more likely to be used as a 'secondary' cloud, so a higher proportion of their revenues than the three largest cloud providers is likely to be using multi-cloud. If anything, this leads to an overestimate of revenue-weighted multi-cloud.

⁷¹⁸ But not the unweighted share of customers.

⁷¹⁹ See appendix I for further details.

⁷²⁰ [AWS' response to the CMA's updated issues statement and working papers](#), paragraph 16; [redacted] submission to the CMA [redacted]; [AWS' response to the CMA's PDR dated 28 January 2025, paragraph 26](#); AWS' submission to the CMA [redacted]; [Microsoft response to the Competitive landscape, Committed spend agreements and Egress fees working papers](#), paragraph 40; Microsoft's submission to the CMA [redacted].

⁷²¹ AWS' submission to the CMA [Economic response to the CMA's Competitive landscape working paper](#), paragraph 56; AWS' submission to the CMA [redacted].

⁷²² AWS' submission to the CMA [redacted].

⁷²³ AWS' submission to the CMA [Economic response to the CMA's Competitive Landscape working paper](#), paragraph 54

time (eg basing conclusions off two customers in 2019 might not reliably predict the behaviour of 20 customers in 2025).

- (c) AWS submitted that even if one attempts to account for other cloud providers (by ‘scaling up’ multi-cloud revenues based on estimates of the prevalence of multi-cloud for other cloud providers, as we have discussed above), such an adjustment may still be arbitrary due to the lack of actual visibility of the customers of other cloud providers.⁷²⁴ As noted in appendix I, we make this assumption as we consider customers of smaller cloud providers are more likely to be multi-clouding than those of larger providers.
- (d) AWS submitted that increasing multi-cloud and switching rates is not an end in itself.⁷²⁵ We agree, and note that we are considering the prevalence of switching and multi-cloud in order to consider it alongside evidence on whether customers find switching and/or multi-cloud difficult, and evidence on wider barriers to switching and multi-cloud. This will allow us to form a judgement on whether barriers to switching and/or multi-cloud are restricting or distorting competition.

3.483 We present the results of our analysis in Table 3.9 below which shows the proportion of customers, by count and spend, identified as using multi-cloud between 2020 and 2022. While the number of customers using multi-cloud is small, they account for a significant share of spend.

Table 3.9: prevalence of multi-cloud, unweighted and weighted by spend, 2020-2023

<i>Prevalence of multi-cloud</i>	<i>2020</i>	<i>2021</i>	<i>2022</i>	<i>2023</i>
Unweighted (%)	6.8%	7.2%	7.4%	7.5%
Weighted by revenue (%)	37.6%	38.5%	40.0%	40.8%

Source: CMA analysis of customer data provided by Microsoft, AWS and Google


3.484 This table shows that 7.5% of customers in 2023 (accounting for 40% of spend) engaged in at least some form of multi-cloud between Microsoft, AWS and/or Google. Conversely, 92.5% of customers (accounting for 60% of all spend) engaged in no multi-cloud at all across these three providers.

3.485 To explore this further, we looked at the prevalence of multi-cloud across customer spend bands amongst the three largest cloud providers: Figure 3.1 below shows the prevalence of multi-cloud weighted by customer spend on cloud, broken down by customer spend band. This is still subject to the same caveat as above: customers are counted as using multiple clouds if they spend over \$1,000 on another cloud, irrespective of the size of that workload. In appendix I, Prevalence of multi-cloud, we show that these results are very similar if customers are not weighted by their cloud spend within spend bands.

⁷²⁴ AWS’ submission to the CMA [38].

⁷²⁵ AWS’ submission to the CMA on multi-cloud and switching dated 16 April 2025, paragraph 28.

Figure 3.1: Prevalence of multi-cloud, weighted by spend, split by spend band, 2020-2023



Source: CMA analysis of customer data provided by Microsoft, AWS and Google.

Note: This graph does not include the adjustment made to account for 'other' providers as we have no visibility over their customers' levels of annual spend and so cannot assign them a spend band.

3.486 The above results do not show the level of integration of customers' multi-clouds. Below, in Table 3.10, we look at the customers' average spend split across cloud providers. We note that, while a customer's spend split might be indicative of the level of integration they have adopted within their multi-cloud setup, we recognise that this is highly dependent on the type of architecture adopted by customers.

3.487 Nonetheless, we observe below how customers typically allocate their spend among cloud providers when they use multi-cloud.

Table 3.10: Average proportion of spend on primary cloud for customers that multi-cloud, by spend band, 2020-2023 (%)

Spend band	2020	2021	2022	2023
Less than 10k	67.1	67.1	66.6	66.4
10K – 1M	82.4	82.7	82.7	83.1
1M – 5M	88.8	90.7	91.3	90.3
5M – 10M	86.7	89.8	91.6	92.5
10M – 20M	94.3	82.0	83.4	84.6
Over 20M	79.3	85.4	86.3	85.2

Source: our analysis of customer data provided by Microsoft, AWS and Google

3.488 The table indicates that, among customers that multi-cloud, for those who spend more than \$10,000 a year on cloud (which is the group of customers with the highest proportion of multi-cloud), their spend is typically concentrated on one primary provider. These results suggest that, in general, the larger the customer is,

the more likely they are to concentrate their spend on their 'primary' cloud. This, together with our findings from Figure 3.1, suggests that those customers which multi-cloud tend to spend more than \$10,000 and to concentrate the majority of their spending on one primary cloud provider.

3.489 Overall, our analysis indicates the following:

- (a) There is demand from customers for some form of multi-cloud. This suggests that at least some customers see the benefits of multi-cloud being greater than the disadvantages.
- (b) Many larger customers multi-cloud, but most smaller and medium-sized customers do not. This might be consistent with larger customers benefitting more from multi-cloud and/or the barriers they face to doing so being relatively lower. However, our evidence does not allow us to reach a firm view on the types of multi-cloud and the level of integration that customers adopt.
- (c) Customers that multi-cloud often have one primary cloud provider accounting for the significant majority of their expenditure. This means that customers' secondary clouds see only a fraction of overall multi-cloud demand, which together with the above findings points to limited multi-cloud on the whole.

3.490 We interpret this evidence in a context of an evolving market: as the offering of public cloud infrastructure services evolves and customers mature in their use of cloud, customers are likely to consider moving towards multi-cloud strategies as they seek to select the best-performing elements from each cloud service. If so, barriers to multi-cloud would become even more important to the competitive process. As noted above, customers' growing interest in accessing FMs may change the way they use multi-cloud architectures.

3.491 Microsoft and AWS both submitted that there is no benchmark or 'well-founded counterfactual' against which to gauge whether observed multi-cloud levels are 'high' or 'low' and therefore signal weak competition or bad outcomes. AWS added that without a benchmark, it is not possible to state with any degree of confidence whether observed levels of multi-cloud or switching rates are lower than what would be expected in a competitive market.⁷²⁶ It also added that not all customers wish to multi-cloud, as it would not be efficient for some of them to do so.⁷²⁷

3.492 In line with our approach to market outcomes, we do not consider that it is necessary or appropriate to specify quantitative benchmarks or thresholds or estimate the 'right' level of multi-cloud in order to assess whether there may be

⁷²⁶ AWS' submission to the CMA [38].

⁷²⁷ AWS' submission to the CMA [38]; [Microsoft response to the Competitive landscape, Committed spend agreements and Egress fees working papers](#), paragraph 41.

barriers to multi-cloud. We have interpreted the evidence on the prevalence of multi-cloud by considering factors such as customer views and the broader context of the market, rather than by using a quantitative benchmark.

Conclusions

- 3.493 There are separate product markets for IaaS based on standard compute infrastructure (which is a separate market from IaaS based on accelerated compute infrastructure) and PaaS, each having a European (UK and EEA) geographic scope. SaaS and alternative IT models (traditional IT and private cloud) do not form part of the same product markets as IaaS and PaaS.
- 3.494 The IaaS market is highly concentrated and the PaaS market is concentrated. The PaaS market has a longer tail of smaller competitors but most are not vertically integrated with IaaS. Therefore, any sale they make may also be an IaaS sale for one of the IaaS providers, as IaaS is an input to PaaS. This weakens the constraint which PaaS-only providers exert on Microsoft or AWS.
- 3.495 Microsoft is now the largest provider in both IaaS and PaaS markets in the UK+EEA and its share is increasing. AWS is the second largest provider in both markets. Google's share in each is much lower than either Microsoft or AWS and while it has grown over the period 2020-2024, it has been more stable in recent years. Other providers are much smaller, particularly in IaaS.
- 3.496 In terms of market outcomes, Microsoft's and AWS' cloud businesses have had sustained profits (as measured by the return on capital employed) substantially above our estimated cost of capital for several years.
- 3.497 Evidence on market shares and profitability, taken together, indicates that Microsoft and AWS each holds a strong position in the IaaS and PaaS markets. Google is the largest of their competitors, but is significantly smaller than them. Other providers of either IaaS or PaaS services are smaller and provide only a weak competitive constraint on Microsoft, AWS and Google. Forward-looking metrics suggest this market structure and outcomes are likely to endure.
- 3.498 Some cloud services have fallen in price while others have seen their prices increase. This is consistent with a market with the potential to be more competitive.
- 3.499 Cloud service providers offer both quality and innovation to their customers who value this, but given the limitations in being able to assess whether levels of innovation are indicative of a well-functioning market, we place more weight on our assessment of the cloud providers' profitability. This is because profitability accounts for broader financial indicators (eg revenue, which is determined by

prices and other variables, and costs including costs to improve quality and innovate).

- 3.500 The prevalence of switching and multi-cloud is helpful in understanding the balance between benefits and costs to customers of switching and multi-cloud, and the current levels of interest in multi-cloud and switching.⁷²⁸
- 3.501 Full switching between cloud providers is extremely rare in cloud services markets. Such a low level of switching together with high levels of profitability among the largest cloud providers is consistent with the presence of high barriers to switching and also indicates that switching costs likely outweigh the benefits of changing cloud provider for most customers.
- 3.502 The prevalence of multi-cloud usage amongst customers suggests that barriers to multi-cloud are lower than for switching. Multi-cloud is used by many larger customers, but most smaller and medium-sized customers do not multi-cloud. But its overall prevalence in conjunction with customer evidence suggesting that the costs of multi-cloud generally outweigh the benefits, indicates that some barriers exist and that, in particular, customers' ability to integrate workloads on more than one cloud is subject to barriers. We note that where multi-cloud is occurring, customers' spend tends to be concentrated on a 'primary' cloud, especially for larger customers.

⁷²⁸ Our consideration of such evidence as relevant context does not need to specify particular quantitative thresholds for the 'right' levels of switching and multi-cloud. As noted above in market outcomes, in relation to market outcomes more generally, we consider that such an approach is neither necessary nor helpful.

4. The impact of AI on cloud services

- We have found that Microsoft, AWS and Google are all strong vertically integrated providers of AI-related cloud services, including through partnerships with FM developers, and that other providers are significantly weaker.
- Whilst AI-related cloud services are likely to be an increasingly important part of the offer from cloud service providers, current evidence from customers and providers suggests that AI is yet to materially alter the competitive dynamics in cloud services.
- **Providers' supply of IaaS based on accelerated compute** is not currently significantly impacting competition in cloud infrastructure services. This is because there is a lack of demand and supply side substitutability between standard and accelerated compute services and significant barriers to entry and expansion in IaaS. Any impact is further limited by the fact that IaaS based on accelerated compute is being used by a small number of customers who are primarily AI model developers.
- **Revenue from the provision of access to FMs is rising** for Microsoft, AWS and Google but it is still a relatively small part of their businesses. This means that the scope for demand for access to FMs to translate into greater demand for other cloud services is currently limited, although we expect demand to grow over time.
- AI capabilities currently play a limited role in driving customers' choice of cloud provider. However, these capabilities are likely to become more important to customers over time and so AI could impact competition in cloud services to a greater degree in the future.
- AI's impact, including on the relative positions of Microsoft, AWS and Google in cloud services, is likely to depend on how customers' demand for AI capabilities across IaaS, PaaS and SaaS evolves and the extent to which increased customer demand for AI capabilities leads to increased demand for other cloud services from the same provider.

Introduction and background

- 4.1 This chapter considers the impact of the growth of AI on competition in cloud services. Cloud providers have told us that AI is changing how competition works in cloud services.

- 4.2 AI⁷²⁹ has become more prominent in public awareness over the past few years, particularly due to generative AI⁷³⁰ services that are powered by foundation models (FMs)⁷³¹.
- 4.3 AI is becoming increasingly important for cloud services, such as:
- (a) IaaS specifically designed for meeting the computational requirements of AI development and deployment.^{732,733} As discussed in our market definition section above, some customers require IaaS based on accelerated compute (ie involving accelerator chips such as GPUs) for their AI workloads.
 - (b) At the PaaS layer there are different platforms designed to assist developers in building AI models themselves⁷³⁴ or provide access to existing AI models (including FMs) with which to build/enrich applications.⁷³⁵ Other PaaS services may also be powered or enriched by AI.⁷³⁶
 - (c) Beyond cloud services, SaaS can be powered by or embedded with AI. This includes consumer, enterprise, and industry specific-applications, and AI assistants.⁷³⁷
- 4.4 Some cloud providers develop their own AI models ('first-party FMs') and make these available on their platforms: for example, AWS has Amazon Nova FMs,⁷³⁸

⁷²⁹ Artificial intelligence (AI) refers to technologies that can perform tasks that would otherwise require human brainpower to accomplish. Machine learning and generative AI are both fields of AI. [Data science and AI glossary | The Alan Turing Institute](#), accessed on 9 May 2025.

⁷³⁰ Generative AI refers to AI technologies that generate text, images or other media in response to user prompts; AI chatbots, such as OpenAI's ChatGPT, are examples of generative AI. [Data science and AI glossary | The Alan Turing Institute](#), accessed on 9 May 2025.

⁷³¹ Foundation Models (FMs) are AI models trained on large amounts of data, enabling them to perform a range of tasks; large language models (LLMs) are a type of FM used to power chatbots. In addition to large amounts of data, developing (or 'training') an FM requires a large amount of computational power and therefore cloud services are typically used for this process. [AI Foundation Models: initial review](#). See page 34 of the Initial Report published on 18 September 2023.

⁷³² For example, see [Azure AI Infrastructure | Microsoft Azure](#), accessed on 23 April 2025, [AI Infrastructure on AWS – Artificial Intelligence Innovation Capabilities – AWS](#), accessed on 23 April 2025, [AI Infrastructure ML and DL Model Training | Google Cloud](#), accessed on 23 April 2025, and [AI Infrastructure | Oracle United Kingdom](#), accessed on 23 April 2025.

⁷³³ Development of AI models is also referred to as 'training'. AI models can then be 'deployed' in new or existing products or services. Each time one of these products or services is used and the model is called upon, a process called 'inference' is performed, which also requires the use of compute. [AI Foundation Models: Initial Report](#), paragraph 2.18.

⁷³⁴ Examples include: [What is Azure Machine Learning? - Azure Machine Learning | Microsoft Learn](#), accessed on 23 April 2025, [Machine Learning Service - Amazon SageMaker AI - AWS](#), accessed on 23 April 2025, and [Vertex AI with Gemini 1.5 Pro and Gemini 1.5 Flash | Google Cloud](#), accessed on 23 April 2025.

⁷³⁵ We refer to cloud providers' services that are used to access a range of existing AI models (often on a metered basis) as 'managed platforms'. Examples include: [Azure OpenAI Service | Microsoft Azure](#), accessed on 23 April 2025, [What is Azure AI Foundry? - Azure AI Foundry | Microsoft Learn](#), accessed on 23 April 2025, [Build Generative AI Applications with Foundation Models - Amazon Bedrock - AWS](#), accessed on 23 April 2025, [Vertex AI with Gemini 1.5 Pro and Gemini 1.5 Flash | Google Cloud](#), accessed on 23 April 2025, [Generative AI Service | Oracle United Kingdom](#), accessed on 23 April 2025, and [IBM watsonx](#), accessed on 23 April 2025.

⁷³⁶ Examples include: [Introduction to Azure AI Search - Azure AI Search | Microsoft Learn](#), accessed on 23 April 2025, [Machine Translation - Amazon Translate - AWS](#), accessed on 23 April 2025, and [Introduction to Vertex AI Search | AI Applications | Google Cloud](#), accessed on 23 April 2025.

⁷³⁷ Examples include: [Microsoft Copilot in Azure | Microsoft Azure](#), accessed on 23 April 2025, [AI Assistant - Amazon Q - AWS](#), accessed on 23 April 2025, and [Gemini for Google Cloud \(formerly Duet AI\) | Google Cloud](#), accessed on 23 April 2025.

⁷³⁸ [Amazon Nova - Generative Foundation Model - AWS](#), accessed on 1 May 2025.

Microsoft has Phi models,⁷³⁹ Google has Gemini models⁷⁴⁰ and IBM has Granite models.⁷⁴¹ These firms are vertically integrated across the AI value chain, from the supply of IaaS based on accelerated compute to AI development and deployment. Oracle does not provide a first-party FM and instead partners with FM developer Cohere to develop generative AI services.⁷⁴²

- 4.5 Cloud providers have told us that AI is changing competition in cloud services and we have considered the impact of AI in different parts of our competitive assessment:
- (a) in market definition we found that the market for IaaS based on standard compute is separate from IaaS based on accelerated compute;
 - (b) in market outcomes we considered the impact of investment in AI on AWS' and Microsoft's profitability;
 - (c) in barriers to entry and expansion, we considered the levels of investment that Microsoft, AWS and Google are making into AI infrastructure;
 - (d) in barriers to switching and multi-cloud, we considered whether the growing use of AI may increase how much influence egress fees have on customer switching and multi-cloud decisions or potentially affect technical barriers; and
 - (e) in licensing we have considered whether AI is impacting modernisation away from Microsoft software, AI as a factor which may reduce the importance of Microsoft software inputs, and the provision of AI services as a potential counterstrategy to Microsoft's licensing practices.
- 4.6 We have received submissions from cloud providers on additional points regarding how AI may impact competition in cloud infrastructure services, in particular:
- (a) AWS and Microsoft submitted that investment in AI is enabling the entry and expansion of smaller providers, such as CoreWeave and Oracle, increasing competition in cloud services.⁷⁴³
 - (b) Microsoft submitted that AI is quickly becoming a key factor in cloud customers' decision-making and that Google has advantages in AI that will strengthen its position as a competitor to AWS and Microsoft.⁷⁴⁴

⁷³⁹ [Phi Open Models - Small Language Models | Microsoft Azure](#), accessed on 1 May 2025.

⁷⁴⁰ [Google models | Generative AI on Vertex AI | Google Cloud](#), accessed on 1 May 2025.

⁷⁴¹ [Granite | IBM](#), accessed on 1 May 2025.

⁷⁴² [Oracle to Deliver Powerful and Secure Generative AI Services for Business | Oracle United Kingdom](#), accessed on 12 May 2025.

⁷⁴³ [AWS' response to the CMA's PDR dated 28 January 2025](#), paragraphs 13 – 15, and [Microsoft's response to the CMA's PDR dated 28 January 2025](#), paragraphs 9 – 12.

⁷⁴⁴ [Microsoft's response to the CMA's PDR dated 28 January 2025](#), paragraphs 17 and 88 – 90.

- 4.7 In this chapter we have considered the potential impact of AI on how competition works in cloud infrastructure services by assessing:
- (a) cloud providers' global revenue related to AI;
 - (b) the impact of competition in the provision of IaaS based on accelerated compute on cloud infrastructure service markets;
- 4.8 how cloud providers compete to supply access to FMs and whether customers' use of FMs impacts their ability to switch and multi-cloud; and
- 4.9 the extent to which cloud providers' AI capabilities⁷⁴⁵ impact customers' decisions regarding cloud infrastructure services, as well as cloud providers' relative strengths in competing to meet customer demand for AI.

Cloud provider revenue related to AI

- 4.10 In order to contextualise our assessment of how AI may impact competition within cloud services, we have examined how it is impacting global cloud revenues.
- 4.11 Our evidence relates to global revenues as the financial reporting, asset base and capital investment of cloud providers is generally global in nature and information on specific services (such as AI-related services) is more readily available at a global level.⁷⁴⁶ Further, the submissions made by cloud providers related to the impact of AI in a global context, and so we consider it appropriate to contextualise our analysis with global data. However, as noted below, customer spend on Microsoft, AWS and Google AI-related cloud services is concentrated with a small number of customers, which are not headquartered in the UK or EEA. Therefore, trends in terms of proportion of revenue and revenue growth attributable to AI may not be reflected in the UK or UK+EEA.

Cloud services' AI-related revenue

- 4.12 Cloud providers have submitted their revenues from cloud services related to AI in financial year 2024:⁷⁴⁷
- (a) AWS provided global revenue data, including services categorised as 'GenAI' and 'AI/ML' IaaS and PaaS services that together amounted to [£] in 2024 (representing approximately [%] of total revenue, [%] in 2023).⁷⁴⁸ Below

⁷⁴⁵ As discussed further below, AI capabilities refer to a range of products and services spanning IaaS, PaaS and SaaS. We have focused on IaaS and PaaS, but considered SaaS where relevant.

⁷⁴⁶ As set out in appendix E on our Profitability analysis

⁷⁴⁷ We note that this analysis is based on stated revenues figures provided by Microsoft, AWS and Google and the services and/or reporting segments included do not necessarily align with those used in our profitability assessment (in our report and appendix E). We considered these figures to provide broad contextual information and trends but have relied on own analysis based on cloud segments for our profitability assessment.

⁷⁴⁸ AWS' response to the CMA's information request [£].

we have referred to AWS' 'GenAI' and 'AI/ML' revenue as its AI-related revenue.

- (b) Microsoft said that 'Azure AI revenue' was \$[REDACTED] in FY2024 and expected to be \$[REDACTED] in FY2025.⁷⁴⁹ For FY2024, this represents approximately [REDACTED] of Microsoft's Azure standalone services revenue and [REDACTED] of Microsoft's Cloud & Enterprise segment revenue (which includes Azure as well as other cloud services such as SaaS and licensing).⁷⁵⁰ Microsoft expects Azure AI revenue to represent [REDACTED]% of Azure standalone services revenue in FY2025.
- (c) Google provided an internal document from September 2024 that showed expected 'GCP AI revenue' for 2024 and 2025, globally, to be \$[REDACTED] and \$[REDACTED], respectively.⁷⁵¹ These expectations represent approximately [REDACTED]% and [REDACTED]% of total Google Cloud revenue in 2024 and 2025, respectively.

4.13 While there are some differences in the basis of these figures so they cannot be compared directly to each other, according to these submissions, cloud services related to AI accounted for a small percentage of their revenues in 2024.

4.14 In its quarterly public accounts for 2024, Microsoft's stated revenue for Azure and other cloud services grew between 29% and 33% in each quarter, with 7 to 13 points (ie approximately two fifths) of this growth being due to AI services;⁷⁵²

- (a) Google said that the provision of AI products and services to both digital natives and traditional enterprises represented [REDACTED];⁷⁵³
- (b) AWS revenue data showed that services categorised as "GenAI" and "AI/ML" accounted for approximately [REDACTED] of revenue growth in 2024.⁷⁵⁴
- (c) There is also some evidence that this growth is expected to continue. Google's submission referred to an article by Goldman Sachs that forecasted generative AI would account for 10-15% of cloud (including SaaS) spending by 2030, with a July 2024 survey of IT buyers showing respondents expected 9% of their budgets to be potentially allocated to generative AI in three years.⁷⁵⁵

⁷⁴⁹ Microsoft's response to the CMA's information request [REDACTED]. Microsoft said that Azure AI revenue includes [REDACTED]. Microsoft's response to the CMA's information request [REDACTED].

⁷⁵⁰ We consider it relevant to consider Azure AI revenue as a proportion of Cloud & Enterprise revenue, given it includes revenue from a range of IaaS, PaaS, and SaaS services as noted in the footnote above.

⁷⁵¹ Google's response to the CMA's information request [REDACTED].

⁷⁵² [Microsoft Fiscal Year 2025 Second Quarter Earnings Conference Call, 29 January 2025](#).

⁷⁵³ Google's response to the CMA's information request [REDACTED].

⁷⁵⁴ CMA analysis of AWS' response to the CMA's information request [REDACTED].

⁷⁵⁵ [Cloud revenues poised to reach \\$2 trillion by 2030 amid AI rollout | Goldman Sachs](#) 4 September 2024, accessed on 14 May 2025.

Breakdown of cloud services AI-related revenue by service and customer

- 4.15 We examined which services and customers were contributing the most to AWS, Microsoft and Google's AI-related cloud services revenue.
- 4.16 For each of AWS, Microsoft and Google, IaaS based on accelerated compute was a significant contributor to their stated AI-related revenue:
- (a) AWS submitted that [REDACTED] of its revenue from services categorised as “GenAI” and “AI/ML” is from its IaaS accelerated compute offer (EC2 Gen AI also known as EC2 GPU).⁷⁵⁶
 - (b) Microsoft submitted that around [REDACTED] of Azure revenue associated with AI in FY2024 was from OpenAI's spend on Azure and this was expected [REDACTED] in FY2025. It submitted that revenue derived from other AI infrastructure made up a further [REDACTED]% of Azure AI-related revenue in FY2024, expected to [REDACTED] in FY2025.⁷⁵⁷
 - (i) Google submitted that, of GCP revenue associated with AI in 2024, it expected [REDACTED]% to be from ‘AI IaaS’ (wholly IaaS based on accelerated compute⁷⁵⁸) and the remainder from ‘AI Cloud’ (including both PaaS and SaaS). [REDACTED].⁷⁵⁹
- 4.17 Customer spend on IaaS based on accelerated compute services from AWS, Microsoft and Google is from a small number of customers, relative to how spend is distributed in cloud infrastructure services more widely:
- (a) AWS provided data that showed [REDACTED] of revenue from “EC2 GenAI” in 2024 was from [REDACTED] customers, [REDACTED] of whom are involved in developing FMs.⁷⁶⁰
 - (b) Microsoft said that around [REDACTED] of Azure revenue associated with AI in FY2024 was from OpenAI's spend on Azure, this was expected to [REDACTED]% in FY2025.⁷⁶¹
 - (c) Google said that around [REDACTED] of the growth in GCP's revenue associated with AI in 2024 was driven by [REDACTED] customers, [REDACTED] of whom are FM developers].⁷⁶²

⁷⁵⁶ AWS' response to the CMA's information request [REDACTED].

⁷⁵⁷ Over [REDACTED]% of AI infrastructure revenue was from IaaS based on accelerated compute. [REDACTED].

⁷⁵⁸ Google's response to the CMA's information request [REDACTED].

⁷⁵⁹ Google's response to the CMA's information request [REDACTED].

⁷⁶⁰ AWS' responses to the CMA's information requests [REDACTED].

⁷⁶¹ Microsoft's response to the CMA's information request [REDACTED].

⁷⁶² Google's response to the CMA's information request [REDACTED].

- 4.18 For AWS and Microsoft, PaaS services that supply access to FMs or facilitate building and using AI models were the next largest contributors to AI-related revenue.⁷⁶³
- (a) AWS' service for building and using AI models (SageMaker) and accessing existing FMs (Bedrock), accounted for [%] and [%] of the revenue for services categorised as "GenAI" and "AI/ML") in 2024, respectively.⁷⁶⁴
 - (b) Microsoft's Azure OpenAI service (which has AI solutions built using OpenAI's FMs) contributed [%] of Azure AI-related revenue in FY2024, expected to [%] in 2025.⁷⁶⁵
- 4.19 AWS and Microsoft's revenue from providing access to FMs has been increasing, but remains a relatively small proportion of their overall revenue:
- (a) AWS provided data that showed global revenue from Bedrock (which became generally available in September 2023) [%] in 2023 to \$[%] in 2024. For context, the same submission stated that AWS' global revenue in 2024 was \$[%].⁷⁶⁶ In an August 2024 planning document, AWS projected that the revenues from Bedrock would rise to [%] and reach [%].⁷⁶⁷
 - (b) Microsoft provided data that showed global revenue from its Azure OpenAI Service was \$[%] in FY24 and estimated to be \$[%] in FY25. For context, the same submission stated that Azure's global revenue was \$[%] in FY24 and estimated to be \$[%] in FY25.⁷⁶⁸

Our assessment of cloud provider revenue related to AI

- 4.20 The revenue data from Microsoft, AWS and Google shows that cloud services related to AI have been growing. This growth has largely been driven by:
- (a) the provision of IaaS based on accelerated compute to customers (eg those developing FMs). This is concentrated amongst a small number of customers; and
 - (b) to a lesser extent, the provision of FMs to cloud customers. Whilst the evidence shows that this is forecast to continue growing, it is currently a small proportion of their revenue.

⁷⁶³ Based on the available evidence, we could not compare similar revenue data for Google.

⁷⁶⁴ AWS' response to the CMA's information request [%].

⁷⁶⁵ Microsoft's response to the CMA's information request [%].

⁷⁶⁶ AWS' response to the CMA's information request [%].

⁷⁶⁷ AWS' response to the CMA's information request [%].

⁷⁶⁸ Microsoft's response to the CMA's information request [%].

- 4.21 As the market evolves over the longer-term, revenue from AI-related cloud services may become more significant and these services may become important for a greater number of customers.

The impact of competition in the provision of IaaS based on accelerated compute on cloud services markets

- 4.22 AWS and Microsoft submitted that entry and expansion of specialised providers of IaaS based on accelerated compute is impacting competition in cloud services.⁷⁶⁹
- 4.23 To understand how competition in the provision of IaaS based on accelerated compute works, and whether it is impacting competition in cloud infrastructure services markets, we have examined:
- (a) The providers of IaaS based on accelerated compute and their relative strengths, including how these compare to the providers of IaaS based on standard compute. To do so, we have examined the relative position and strength of Microsoft, AWS and Google, as well as other providers of accelerated compute, including Oracle, IBM, CoreWeave, Lambda Labs, Nebius and Nscale.
 - (b) The barriers to entry and expansion in the provision of IaaS based on accelerated compute, namely cloud provider investments and partnerships between cloud providers and AI developers.

Providers of IaaS based on accelerated compute and their relative strengths

- 4.24 We have considered the providers of IaaS based on accelerated compute and their strengths, cloud providers' accelerated compute capacity and evidence from cloud providers' internal documents on the competitive position of providers.

Providers of IaaS based on accelerated compute

- 4.25 Microsoft, AWS, Google, Oracle and IBM all supply IaaS based on accelerated compute, alongside their IaaS offerings based on standard compute.⁷⁷⁰ As with IaaS based on standard compute, these cloud providers offer a range of software and tools for managing accelerated compute workloads.

⁷⁶⁹ AWS' response to the CMA's PDR dated 28 January 2025, paragraphs 13 – 15, and Microsoft's response to the CMA's PDR dated 28 January 2025, paragraphs 9 – 12.

⁷⁷⁰ See: [Cloud Compute Instances – Amazon EC2 Instance Types – AWS](#), accessed on 16 May 2025, [Virtual Machine series | Microsoft Azure](#), accessed on 16 May 2025, [Cloud GPUs \(Graphics Processing Units\) | Google Cloud | Google Cloud](#), accessed on 16 May 2025, [GPU, Virtual Machines and Bare Metal | Oracle United Kingdom](#), accessed on 16 May 2025 and [GPU Cloud Server | IBM](#), accessed on 16 May 2025.

- 4.26 Some smaller cloud providers have also expanded to offer accelerated compute, for example, Vultr.⁷⁷¹ There are also several smaller cloud providers that specialise in the provision of accelerated compute:
- (a) CoreWeave was established in 2017 as a company focused on cryptocurrency applications.⁷⁷² CoreWeave describes itself as a cloud provider that specialises in GPU-accelerated workloads.⁷⁷³ CoreWeave went public in March 2025, listing on the NASDAQ; its filing showed that CoreWeave's revenue in 2024 was \$1.9billion (net loss \$863million).⁷⁷⁴
 - (b) Lambda Labs was established in 2012 as a company focused on facial recognition technology but in 2018 launched a GPU cloud, pivoting to Cloud and Software Services in 2021.⁷⁷⁵ It was reportedly valued at \$1.5billion in February 2024.⁷⁷⁶
 - (c) Crusoe Cloud (launched in 2023) supplies AI infrastructure.⁷⁷⁷ Crusoe Cloud is owned by Crusoe, which also provides data centre and energy solutions.⁷⁷⁸ In December 2024, Crusoe raised \$600million in Series D funding at a valuation of \$2.8billion.⁷⁷⁹
 - (d) Nebius is an AI infrastructure company owned by the Nebius Group; it provides a 'GPU Cloud' and some AI PaaS services.⁷⁸⁰ The Nebius Group's revenue in 2024 was \$117.5million (net loss \$396.9m), with over half of the Group's total revenue coming from the AI infrastructure business.⁷⁸¹
 - (e) Nscale is an AI infrastructure company launched in 2024.⁷⁸² Nscale offers access to GPU clusters and an API that supplies access to FMs; its website states that it plans to offer additional services for fine-tuning and evaluating FMs.⁷⁸³ In December 2024, Nscale raised \$155million in Series A funding.⁷⁸⁴
- 4.27 Specialised providers' offerings tend to be targeted at a narrower group of customers (for example, the 'solutions' offered on CoreWeave's website are for 'AI Model Training', 'AI Inference', and 'VFX & Rendering') and therefore offer a more

⁷⁷¹ [Vultr rides cloud GPU wave, widens ambitions, July 17 2023](#), accessed on 1 May 2025.

⁷⁷² [CoreWeave came 'out of nowhere.' Now it's poised to make billions off AI with its GPU cloud](#), accessed on 25 August 2024.

⁷⁷³ CoreWeave's response to the CMA's information request [§].

⁷⁷⁴ [COREWEAVE, INC. filing with the U.S. Securities and Exchange Commission March 3, 2025](#), accessed on 24 April 2025.

⁷⁷⁵ [About Lambda](#), accessed on 20 September 2024.

⁷⁷⁶ [AI Compute Startup Lambda Hits \\$1.5B Valuation After Massive \\$320M Raised](#), accessed on 20 September 2024.

⁷⁷⁷ [The Future Of Computing Can And Should Be Climate-Aligned and Cloud](#), accessed on 1 May 2025.

⁷⁷⁸ [Data Centers and Energy](#), accessed on 1 May 2025.

⁷⁷⁹ [Crusoe Closes \\$600M in Series D Round at \\$2.8 Billion Valuation to Power AI](#), accessed on 1 May 2025.

⁷⁸⁰ [Nebius Group Investor Presentation October 18, 2024](#), accessed on 1 May 2025.

⁷⁸¹ [Nebius Group N.V. announces fourth quarter and full-year 2024 financial results](#), accessed on 1 May 2025.

⁷⁸² [London AI hyperscaler Nscale in \\$155m funding round - UKTN](#), accessed on 1 May 2025.

⁷⁸³ [Overview - Nscale](#), accessed on 1 May 2025.

⁷⁸⁴ [AI hyperscaler Nscale secures \\$155 million in Series A | Press Release | Nscale](#), accessed on 1 May 2025.

limited range of cluster management options, including more basic management software or none at all ('bare metal').⁷⁸⁵

- 4.28 In some instances, large public cloud providers procure bare metal services from other cloud providers.⁷⁸⁶
- 4.29 In addition to providing IaaS based on accelerated compute, some of the specialised providers also supply a limited IaaS based on standard compute offering and a narrow range of AI PaaS and AI SaaS offerings (for example, AI development platforms or access to third-party FMs), as per their websites cited above. These offerings are more limited than those offered by larger cloud providers (referenced in the Introduction and background section above). Furthermore, we consider that these specialised providers generally do not serve the wider cloud services market, with many targeting customers who are developing and/or deploying AI models.

Cloud providers' accelerated compute capacity

- 4.30 Given the global nature of financial reporting, asset base and capital investment of cloud providers, we have not been able to construct revenue-based shares of supply specific to IaaS based on accelerated compute. Here, we have instead examined cloud providers' accelerated compute capacity (measured in TFLOPS, as of mid-2024)⁷⁸⁷ and set out available information regarding global revenues.
- 4.31 Microsoft, AWS and Google each have approximately [REDACTED] the amount of accelerated compute capacity than CoreWeave which in turn has approximately [REDACTED] the accelerated compute capacity of each of Oracle and Lambda Labs. IBM had approximately [REDACTED] less accelerated compute than the other providers.⁷⁸⁸
- 4.32 We have not examined the TFLOP capacity of other smaller providers. However, according to public statements, Nebius planned to have deployed 20k GPUs by the end of 2024 and Nscale expected to deploy 50k GPUs by the end of 2025.⁷⁸⁹ Based on internal documents we have seen, these figures (20-50k) represent approximately [REDACTED].⁷⁹⁰
- 4.33 A cloud provider's accelerated compute capacity may indicate its strength in the provision of IaaS based on accelerated compute, although there is not necessarily a straightforward correlation. In particular, some providers may use accelerated

⁷⁸⁵ [AI Infrastructure Explained](#), accessed on 30 September 2024.

⁷⁸⁶ [REDACTED].

⁷⁸⁷ Teraflops (TFLOPS) are used to measure computational performance. One TFLOPS is equal to one trillion floating point operations per second.

⁷⁸⁸ Responses to the CMA's information requests [REDACTED].

⁷⁸⁹ [Nebius Group Investor Presentation October 18, 2024](#), page 8, accessed on 1 May 2025. [Nscale LinkedIn Post, December 2024](#), accessed on 1 May 2025.

⁷⁹⁰ Responses to the CMA's information requests [REDACTED].

compute capacity to serve their own internal demand (eg to develop first-party FMs) or that of their parent company's business.

- 4.34 AWS, Microsoft and Google each provided data on global revenues associated with the supply of IaaS based on accelerated compute, which indicated that these revenues are around a similar level.⁷⁹¹
- (a) In 2024, AWS revenue from its Amazon EC2 'accelerated computing' instance types was \$[REDACTED]. Of this revenue, [REDACTED]% was categorised as [REDACTED].⁷⁹²
 - (b) In FY24, Microsoft Azure revenue from IaaS based on accelerated compute was \$[REDACTED]⁷⁹³ (this includes revenue from [REDACTED]). Of this revenue, at least [REDACTED]% was categorised by Microsoft as being [REDACTED].⁷⁹⁴
 - (c) In 2024, Google expected its revenue from AI IaaS (wholly IaaS based on accelerated compute⁷⁹⁵) to be [REDACTED].⁷⁹⁶

Cloud provider internal documents

- 4.35 We examined cloud providers' internal documents to understand who they consider to be their main competitors and their perceptions of the competitive strengths of different providers.
- 4.36 Microsoft's internal documents demonstrate that it considers its principal competitors in the supply of accelerated compute to be [REDACTED], rather than [REDACTED]. In some internal documents, [REDACTED] are described as [REDACTED]. Microsoft procures [REDACTED].⁷⁹⁷ This suggests that, in some documents, these providers were viewed as being [REDACTED] to Microsoft [REDACTED].
- 4.37 In an email, Microsoft compared the [REDACTED].⁷⁹⁸
- 4.38 An AWS [REDACTED] stated that AWS competes with 'a large and growing number of specialized AI cloud services including CoreWeave, Lambda Labs, San Francisco Compute, Genesis Cloud, Denvr Dataworks, G42, Omniva, Cirrascale, Crusoe Cloud, Gcore, and more.'⁷⁹⁹

⁷⁹¹ We note that the data provided did not contain sufficient detail to be able to determine if the services included were exactly like-for-like.

⁷⁹² AWS' responses to the CMA's information requests [REDACTED].

⁷⁹³ Due to potential overlap in the categories of the data provided, this figure is likely to be an overestimate.

⁷⁹⁴ Microsoft's responses to the CMA's information requests [REDACTED].

⁷⁹⁵ Google's response to the CMA's information request [REDACTED].

⁷⁹⁶ Google's response to the CMA's information request [REDACTED].

⁷⁹⁷ Microsoft's response to the CMA's information request [REDACTED].

⁷⁹⁸ Microsoft's response to the CMA's information request [REDACTED].

⁷⁹⁹ AWS' response to the CMA's information request [REDACTED].

- 4.39 In a September 2024 strategy document, Google said [REDACTED]. The document noted that [REDACTED].⁸⁰⁰
- 4.40 The same Google [REDACTED].⁸⁰¹
- 4.41 A pricing document from a cloud provider compares its prices of accelerated compute offerings and associated discounts to other large cloud providers – the cloud provider considered when and whether to match the competitors’ discounts.⁸⁰² Smaller specialised providers were not mentioned.
- 4.42 In another internal document from a cloud provider, this provider noted that some of its large competitors, including Oracle and Azure, were offering ‘super deep’ discounts on accelerated compute by ‘cross subsidizing their cloud businesses’.⁸⁰³

Barriers to entry and expansion

- 4.43 We have considered whether newer or smaller cloud providers will be able to enter or expand in the supply of IaaS based on accelerated compute in such a way that impacts competition in cloud services. In particular, we have considered:
- (a) Cloud providers’ views;
 - (b) Cloud provider investments in the supply of accelerated compute; and
 - (c) Partnerships between cloud providers and AI developers.

Cloud provider views

- 4.44 Cloud providers submitted that new and existing providers of accelerated compute have entered and expanded due to investment in AI, in particular:
- (a) Microsoft submitted that the rapid increase in demand for accelerated compute had accelerated entry and expansion of suppliers of specialised cloud infrastructure services focused on large-scale GPU-accelerated workloads.⁸⁰⁴ Microsoft further submitted that investment into generative AI is ‘lifting providers old and new’ such as Oracle and ‘startups with infrastructure optimised for AI workloads’ including Lambda Labs, Crusoe, CoreWeave and NScale.⁸⁰⁵ Microsoft separately said that, to access additional GPU capacity

⁸⁰⁰ Google’s response to the CMA’s information request [REDACTED].

⁸⁰¹ Google’s response to the CMA’s information request [REDACTED].

⁸⁰² [REDACTED] response to the CMA’s information request [REDACTED].

⁸⁰³ [REDACTED] response to the CMA’s information request under section 174 notice issued 7 April 2025, question 2, Annex 2.1, page 15.

⁸⁰⁴ Microsoft’s response to the CMA’s information requests [REDACTED].

⁸⁰⁵ [Microsoft’s response to the CMA’s PDR dated 28 January 2025](#), paragraph 12.

for AI workloads, it has agreements with smaller cloud providers – these include agreements with CoreWeave and Lambda Labs.⁸⁰⁶

- (b) AWS submitted that numerous smaller cloud providers have entered or expanded to offer compute for AI workloads and that ‘emerging’ cloud providers such as CoreWeave, Lambda Labs, Vultr, and Nebius provide ‘compelling alternatives to larger IT providers’ and reportedly raised billions of dollars in 2024.⁸⁰⁷ AWS also submitted that substantial investments by venture capital and private equity firms into small and emerging cloud providers demonstrate the prospects of smaller and upcoming providers.⁸⁰⁸
- (c) One cloud provider told us that Nvidia is actively promoting the entry of CoreWeave and Lambda Labs by, for example, granting them preferential access to Nvidia’s GPUs.⁸⁰⁹ It has been reported that CoreWeave and Lambda Labs were two of the first cloud providers to get general access to Nvidia’s H100 chip when it launched in 2023.⁸¹⁰

4.45 [X] also said that the expanding data centre capacity from new entry and expansion in the supply of IaaS based on accelerated compute, as well as from private investment, is an addition of IT capacity that will drive more competition for all types of workloads. It added that, given the shared infrastructure and know-how required to supply IaaS based on standard compute, these providers can be expected to expand the IT services they offer (as [X] had) and compete ‘even more fiercely’ with it.⁸¹¹

4.46 However, the investments AWS and Microsoft are making into accelerated compute, which amount to tens of billions of dollars annually (as discussed below) are much larger than the investment being made in these specialised providers. CoreWeave, who AWS and Microsoft cite frequently in their submissions as a competitor (and who, as set out above, is the fourth largest provider of IaaS based on accelerated compute by capacity, according to evidence we have seen) is dependent on larger cloud providers as customers, in particular, Microsoft.⁸¹²

⁸⁰⁶ Microsoft’s response to the CMA’s information request [X].

⁸⁰⁷ [AWS’ response to the CMA’s PDR dated 28 January 2025](#), paragraph 15.

⁸⁰⁸ The noted investments included Vultr’s \$333 million raise from AMD and LuminArx Capital, and Nvidia’s and Accel’s \$700 million investment in Finnish data centre provider Nebius. AWS’ submission to the CMA [X].

⁸⁰⁹ [X] response to the CMA’s information request [X].

⁸¹⁰ ‘What is ‘compute’, and why does it matter?’, [Computational Power and AI - AI Now Institute](#), page 32, accessed on 30 September 2024; CoreWeave has also received a \$100 million investment from Nvidia (ibid.). [CoreWeave came ‘out of nowhere.’ Now it’s poised to make billions off AI with its GPU cloud](#), accessed on 25 August 2024. [Lambda Cloud Deploys NVIDIA H100 Tensor Core GPUs](#). See also [X] response to the CMA’s information request [X].

⁸¹¹ [X] response to the CMA’s information request [X].

⁸¹² CoreWeave disclosed in its IPO filing that, in 2024, Microsoft accounted for 62% of its revenue and that ‘any negative changes in demand from Microsoft, in Microsoft’s ability or willingness to perform under its contracts with us, in laws or regulations applicable to Microsoft or the regions in which it operates, or in our broader strategic relationship with Microsoft would adversely affect our business, operating results, financial condition, and future prospects.’ [COREWEAVE, INC. filing with the U.S. Securities and Exchange Commission March 3, 2025](#), accessed on 24 April 2025.

[REDACTED].⁸¹³ We therefore considered that it is more appropriate to view CoreWeave as primarily a supplier to large cloud providers rather than a competitor to them.

Cloud provider investments in the supply of accelerated compute

4.47 This section covers cloud providers' capex investments in supplying accelerated compute and in AI accelerator chips.⁸¹⁴

4.48 Several cloud providers have highlighted that demand for accelerated compute exceeds supply.⁸¹⁵ AWS, Microsoft and Google said that they are making substantial investments to expand their accelerated compute capacity. In particular:

- (a) [REDACTED], AWS noted that [REDACTED].⁸¹⁶ AWS since submitted that in 2024 it spent [REDACTED] on 'GenAI' server capex.⁸¹⁷
- (b) In a finance review submitted to Microsoft's Board of Directors in July 2023, Microsoft's server GPU capex was forecast to [REDACTED] server CPU capex was forecasted to [REDACTED].⁸¹⁸
- (c) In a 2023 strategy document, Google outlined how its GPU and TPU capex grew from [REDACTED] in 2022 and was expected to grow again to [REDACTED] in 2023 and [REDACTED] in 2024.⁸¹⁹ Google since told the CMA that it expects to spend [REDACTED] on AI servers (those that incorporate a GPU or TPU⁸²⁰) for Google Cloud in 2025.⁸²¹ In a 2024 strategy document, Google stated that [REDACTED].⁸²²

4.49 A January 2024 document from a cloud provider noted that 'the scale required [to provide training infrastructure for FMs] will be out of reach for all but very few companies. As a result, the primary competitors in this segment of the market will be the large cloud service providers.'⁸²³ Another cloud provider submitted to the

⁸¹³ [REDACTED].

⁸¹⁴ In addition to accessing accelerator chips, cloud providers also require access to additional hardware and infrastructure to build data centres, including memory chips, network technologies, cooling systems and power supplies in sufficient volume. They also need software to programme the AI accelerator chips and integrate them with other cloud services and applications. [Computing Power and the Governance of Artificial Intelligence](#), accessed on 30 September 2024; [Computational Power and AI - AI Now Institute](#), accessed on 30 September 2024.

⁸¹⁵ [REDACTED] response to the CMA's information request [REDACTED]; [AWS earnings conference call for Q3 2024](#); [Oracle's earnings conference for Q3 of the FY24, Oracle announces fiscal 2024 third quarter financial results](#), accessed on 30 September 2024.

⁸¹⁶ AWS' response to the CMA's information request [REDACTED].

⁸¹⁷ AWS' response to the CMA's information request [REDACTED].

⁸¹⁸ Microsoft's response to the CMA's information request [REDACTED].

⁸¹⁹ Google's response to the CMA's information request [REDACTED].

⁸²⁰ Google's response to the CMA's information request dated 23 April 2025, question 1, page 2.

⁸²¹ Google's response to the CMA's information request under section 174 notice issued 7 April 2025, question 1, page 2.

⁸²² Google's response to the CMA's information request under section 174 notice issued 7 April 2025, question 2, Annex 2.1, pages 1 and 2.

⁸²³ [REDACTED] response to the CMA's information request [REDACTED].

CMA that ‘barriers to meeting [FM developers’ compute] demands relate primarily to capital required to install and operate the compute’.⁸²⁴

4.50 AWS submitted that, with the advance of small and more efficient FMs, such as DeepSeek’s R1, consumers are increasingly able to deploy FMs on devices and personal computers using consumer grade chips.⁸²⁵

4.51 In the context of the significant capex investment required to grow accelerated compute capacity, Microsoft and AWS internal documents have noted the importance in returning revenue on these. For example:

(a) A March 2024 CFO update to Microsoft’s Board said [REDACTED].⁸²⁶

(b) In a November 2023 AWS update [REDACTED] discussing generative AI, AWS said that [REDACTED].⁸²⁷

4.52 Cloud providers’ internal documents also noted the low or negative margins resulting from the significant capex associated with AI.

(a) In a March 2023 update [REDACTED], AWS noted that [REDACTED].⁸²⁸

(b) A September 2023 CFO update to Microsoft’s Board said that [REDACTED].⁸²⁹

(c) In a September 2024 strategy document, Google noted that gross margin associated with AI (included IaaS) was expected to [REDACTED].⁸³⁰

Investments in AI accelerator chips

4.53 Cloud providers are making investments to increase their accelerated compute capacity, including by procuring AI accelerator chips. We have considered two main sources of supply of AI accelerator chips.⁸³¹

(a) established chip providers, namely Nvidia, AMD and Intel;⁸³² and

⁸²⁴ [REDACTED] response to the CMA’s information request [REDACTED].

⁸²⁵ [AWS’ response to the CMA’s PDR dated 28 January 2025](#), paragraph 16

⁸²⁶ Microsoft’s response to the CMA’s information request [REDACTED].

⁸²⁷ AWS’ response to the CMA’s information request [REDACTED].

⁸²⁸ AWS’ response to the CMA’s information request [REDACTED].

⁸²⁹ Microsoft’s response to the CMA’s information request [REDACTED].

⁸³⁰ [REDACTED] response to the CMA’s information request [REDACTED].

⁸³¹ [Computational Power and AI - AI Now Institute](#), accessed on 30 September 2024.

⁸³² AI accelerator chips can also be sourced from specialist chip startups (such as Graphcore, Tenstorrent and Cerebras); however, based on public sources, we are not aware of any larger cloud providers currently using AI accelerator chips developed by specialist chip startups.

- (b) self-supply (or ‘custom silicon’), ie cloud providers developing their own ASICs,⁸³³ which can enable systems-level optimisations and hence a lowering of costs.⁸³⁴

Investments in AI accelerator chips by third parties

- 4.54 Evidence from cloud providers shows that the majority of their accelerated compute capacity is powered by Nvidia’s GPUs.⁸³⁵ Evidence from AWS and Microsoft indicates both the importance and the cost of being able to offer Nvidia chips at scale. For example:
- (a) AWS’ internal documents indicate that, given the popularity of Nvidia’s GPUs, [REDACTED].⁸³⁶ In a July 2023 internal document, AWS noted that ‘NVIDIA GPU costs represented about [REDACTED] of its total server cost’. AWS noted that [REDACTED].⁸³⁷
 - (b) Microsoft submitted that it had spent approximately [REDACTED] on GPUs for model training and inference globally in 2023,⁸³⁸ [REDACTED].⁸³⁹
 - (c) A Google internal document setting out its AI strategy for 2025 outlined its plans to [REDACTED].⁸⁴⁰
- 4.55 The supply of AI accelerator chips (such as GPUs) has been unable to keep pace with increasing demand associated with AI over the past three years:
- (a) One cloud provider noted that a delay in Nvidia chip delivery in 2024 had been a key factor leading to lower estimated AI revenue for 2025.⁸⁴¹
 - (b) One cloud provider submitted that ‘any barriers to supplying accelerated compute continue to be lowered by industry innovations’. It said that the ‘increasing variety of credible accelerator chips available enables cloud providers to diversify their accelerated compute offerings’ and that ‘innovations in compute-efficient FM training’ are ‘reducing the amount of compute needed to achieve comparable performance’.⁸⁴²

⁸³³ As explained in the market definition section of this chapter, ASICs are accelerator chips which are hardwired for a specific application. ASICs developed for FM training and inference are typically hyper-specialised GPUs, with modifications that increase the efficiency of specific AI workloads. The AI accelerator chips developed by Google (TPUs), Amazon (Trainium and Inferentia) and Microsoft (Maia 100) are all ASICs. However, there are other manufacturers of ASICs chips in the market aside from these cloud providers.

⁸³⁴ [With a systems approach to chips, Microsoft aims to tailor everything ‘from silicon to service’ to meet AI demand - Source](#), accessed on 30 September 2024; [Google AI Infrastructure Supremacy: Systems Matter More Than Microarchitecture \(semianalysis.com\)](#), accessed on 30 September 2024; [Announcing Cloud TPU v5e and A3 GPUs in GA | Google Cloud Blog](#), accessed on 30 September 2024.

⁸³⁵ Responses to the CMA’s information requests [REDACTED].

⁸³⁶ AWS’ response to the CMA’s information request [REDACTED].

⁸³⁷ AWS’ response to the CMA’s information request [REDACTED].

⁸³⁸ [REDACTED] these GPUs were used for model training.

⁸³⁹ Microsoft’s response to the CMA’s information request [REDACTED].

⁸⁴⁰ Google’s response to the CMA’s information request [REDACTED].

⁸⁴¹ [REDACTED] response to the CMA’s information request [REDACTED].

⁸⁴² [REDACTED] response to the CMA’s information request [REDACTED].

Cloud provider self-supply of AI accelerator chips

- 4.56 Microsoft, AWS and Google are also investing in developing their own proprietary AI accelerator chips for internal use and providing to cloud customers.⁸⁴³
- (a) Google's series of Tensor Processing Units (TPUs) designed for machine learning workloads including FM development and deployment.⁸⁴⁴ While TPUs are used predominantly for internal Google workloads (such as training its FMs), they are also available on the Google Cloud Platform and are used by external FM developers including AI21,⁸⁴⁵ Anthropic,⁸⁴⁶ and Midjourney.⁸⁴⁷
 - (b) Amazon's 'Trainium' and 'Inferentia' series of AI accelerator chips, which are only available to AWS customers. Customers include Anthropic and Snap.⁸⁴⁸
 - (c) Microsoft announced its first custom AI accelerator chip, the 'Maia 100', in November 2023.⁸⁴⁹ Microsoft said in 2024 that it planned for the Maia 100 to power Microsoft's services such as Copilot and Azure OpenAI Service.⁸⁵⁰ [REDACTED].⁸⁵¹
- 4.57 Self-supply of chips may reduce cloud providers' reliance on third-party suppliers, and this may be especially important in the context of GPU supply constraints.
- 4.58 Cloud providers consider it important to have proprietary AI accelerator chips:
- (a) Google considers its custom silicon (TPUs) as a differentiator compared to other cloud providers and as important to achieving its aims of reducing training and inference costs.⁸⁵²
 - (b) AWS considers that its custom silicon (Trainium and Inferentia chips) are beginning to 'differentiate [its] infrastructure offering in the GenAI space' [REDACTED].⁸⁵³

⁸⁴³ Oracle said that it has [REDACTED]. Oracle's response to the CMA's information request [REDACTED]. IBM said that its research division has invested in several prototype chip projects, but that IBM Cloud does not currently commercially offer access to IBM-build accelerators to customers. IBM's response to the CMA's information request [REDACTED].

⁸⁴⁴ [Tensor Processing Units \(TPUs\)](#), accessed on 30 September 2024. Google TPUs are developed in collaboration with Broadcom, who manufactures and supplies the chips. Google's response to the CMA's information request [REDACTED].

⁸⁴⁵ [Building the most open and innovative AI ecosystem](#), accessed on 20 September 2024.

⁸⁴⁶ [Cloud TPU v5e is generally available](#), accessed on 20 September 2024.

⁸⁴⁷ [Building the most open and innovative AI ecosystem](#), accessed on 20 September 2024.

⁸⁴⁸ [AWS AI chips powering Amazon's partnership with Anthropic](#), accessed on 30 September 2024. [AI Chip - AWS Inferentia](#), accessed on 30 September 2024. In its [REDACTED] report for 2024, AWS stated that: [REDACTED]. AWS' response to the CMA's information request [REDACTED].

⁸⁴⁹ Microsoft's response to the CMA's information request [REDACTED]; [Microsoft Azure delivers purpose-built cloud infrastructure in the era of AI](#), accessed on 30 September 2024.

⁸⁵⁰ Microsoft's response to the CMA's information request [REDACTED]; [Microsoft Azure delivers purpose-built cloud infrastructure in the era of AI](#), accessed on 30 September 2024.

⁸⁵¹ Microsoft's response to the CMA's information request [REDACTED].

⁸⁵² Google's response to the CMA's information requests [REDACTED].

⁸⁵³ AWS' response to the CMA's information request [REDACTED].

(c) Microsoft's internal documents note [REDACTED].⁸⁵⁴

4.59 AWS and Microsoft have noted the large investments required to develop custom AI accelerator chips. For example:

(a) AWS said that it had invested [REDACTED] over five years (2019-23) in order to develop its custom AI accelerator chips Trainium and Inferentia, and the accompanying software to program them, Neuron.⁸⁵⁵ AWS said that it [REDACTED].⁸⁵⁶ AWS said that it is also working to meet the growing demand for its Trainium and Inferentia chips.⁸⁵⁷

(b) Microsoft said that between July 2021 and June 2024 (including forecast spend), it expects to have spent [REDACTED] in developing its custom AI accelerator chip, Maia.⁸⁵⁸

(c) Microsoft estimated that Microsoft, AWS and Google had invested between \$139billion and \$323billion on cloud computing between 2018 and 2022, including investments to develop new hardware and software, such as improved processing units capable of running AI-intensive workloads.⁸⁵⁹

4.60 Self-supply also requires development of the accompanying software (that is used to programme the AI accelerator chips), which requires significant investment.⁸⁶⁰

(a) A cloud provider stated [REDACTED], 'the biggest asset for NVIDIA is its superior software environment, which gives developers and scientists the flexibility to innovate and optimise performance'. [REDACTED].⁸⁶¹

Partnerships between cloud providers and FM developers

4.61 Some FM developers have entered into a commercial 'compute partnership' with a cloud provider whereby the FM developer gains preferential or discounted access to accelerated compute in exchange for supplying the cloud provider with preferential or exclusive access to its FMs.⁸⁶²

4.62 To understand whether these partnerships could raise barriers to entry and expansion for new or smaller providers of accelerated compute, we asked cloud providers about their experiences of such partnerships. Overall, providers had

⁸⁵⁴ Microsoft's response to the CMA's information request [REDACTED].

⁸⁵⁵ AWS' response to the CMA's information request [REDACTED].

⁸⁵⁶ AWS' response to the CMA's information request [REDACTED].

⁸⁵⁷ AWS' earnings conference for Q1 of the FY24, April 30th 2024, [Amazon.com announces first quarter results](#), accessed on 30 September 2024.

⁸⁵⁸ Microsoft's response to the CMA's information request [REDACTED].

⁸⁵⁹ Microsoft's response to the CMA's information request [REDACTED].

⁸⁶⁰ Notes of meetings with [REDACTED].

⁸⁶¹ [REDACTED] response to the CMA's information request [REDACTED].

⁸⁶² [AI Foundation Models: initial review](#). See the Update Paper published on 11 April 2024 and the Technical Update Report published on 16 April 2024.

mixed views with some describing them as pro-competitive or neutral, and others highlighting certain characteristics that could lead some partnerships to negatively impact competition:

- 4.63 One cloud provider told us that partnerships do not offer any competitive advantage for AI developers or cloud providers with regards to the purchase of compute capacity, given the broad range of options available for developers to source this capacity.⁸⁶³
- (a) Another cloud provider submitted that partnerships are pro-competitive because it they will allow for ‘rapid distribution and use of AI-infused technologies’ such that ‘competition intensifies at the [AI] application layer’, which will in turn ‘drive even greater competitive pressure at the cloud and [FM] layer, with downstream application providers demanding lower cost inputs from providers of infrastructure and [FMs]’.⁸⁶⁴
 - (b) These two cloud providers submitted that partnerships can accelerate the design, development and deployment of providers’ accelerated compute infrastructure, increasing capacity in the market and expanding customer choice.⁸⁶⁵ However, one of these providers noted that ‘these benefits are not exclusive to Partnerships, and could be generated through other means (eg, ordinary course commercial relationships with no Partnership)’.⁸⁶⁶
 - (c) A third cloud provider said that partnerships that ‘confer upon the cloud provider the ability to materially influence the strategic behaviour of the FM developer/provider in the marketplace, including its ability to freely choose its cloud procurement and FM commercialisation and distribution strategies (...) have the potential to materially influence competitive dynamics in public cloud’ because such partnerships restrain an FM developer’s ability to ‘freely choose from a wide and increasing range of cloud services to train and host their FMs, (...) freely switch between cloud providers and adopt multi-cloud strategies and freely choose or multi-cloud from a wide and increasing range of AI platforms and model hubs’.⁸⁶⁷
 - (d) A fourth cloud provider said that, given the projected increase in demand for accelerated compute, ‘tight relationships between some providers and AI developers will not only be advantageous but necessary for the development of large FMs’. However, it cautioned that the ‘concentration and dependency on compute resources from a few providers would certainly influence competitive dynamics for other cloud providers, who may not be able to

⁸⁶³ [REDACTED] response to the CMA’s information request [REDACTED].

⁸⁶⁴ [REDACTED] response to the CMA’s information request [REDACTED].

⁸⁶⁵ Responses to the CMA’s information requests issued [REDACTED].

⁸⁶⁶ [REDACTED] response to the CMA’s information request [REDACTED].

⁸⁶⁷ [REDACTED] response to the CMA’s information request [REDACTED].

procure and provide (due to tight supply and/or required size of investment) the large amount of required compute capacity'.⁸⁶⁸

Our assessment of investments in AI accelerator chips

- 4.64 We have considered cloud providers' supply of IaaS based on accelerated compute and found that:
- (a) Supplying IaaS based on accelerated compute is a growing revenue source globally for cloud providers. Demand is currently concentrated with a small number of large customers who are typically FM developers;
 - (b) Microsoft, AWS and Google each has a strong position in the supply of IaaS based on accelerated compute globally. They each have significant accelerated compute capacity and are investing to further increase this.
 - (c) Other providers offer IaaS based on accelerated compute, including smaller specialised providers which have emerged to meet AI demand. However, we have found that they have much smaller accelerated compute capacity than Microsoft, AWS and Google and are investing significantly less. Further, specialist providers tend to target their offerings at a narrower group of customers. We also note that some of these providers are partners or suppliers to the large cloud providers rather than competitors.
 - (d) Some specialised providers also supply a limited IaaS based on standard compute offering and a narrow range of AI PaaS and AI SaaS offerings. However, we considered that these specialised providers generally do not serve the wider cloud services markets, with many targeting customers who are developing and/or deploying AI models. Therefore, the cloud providers supplying IaaS based on accelerated compute, and their relative strengths, differ from the market for IaaS based on standard compute.
 - (e) There are specific and significant barriers to entry and expansion in the supply of IaaS based on accelerated compute.
 - (f) We have found that there is significant capex investment associated with growing accelerated compute capacity which requires scale. In particular, Microsoft, AWS and Google are investing significant sums to procure AI accelerator chips from third parties – predominantly from Nvidia – and develop first-party chips. The emergence of new models that are less computationally intensive, such as DeepSeek, may reduce this barrier to entry and expansion in future; however, this is so far having limited impact.

⁸⁶⁸ [redacted] response to the CMA's information request [redacted].

Microsoft, AWS and Google continue to outspend smaller cloud providers on expanding their accelerated compute capacity.

- (g) Partnerships between larger cloud providers and FM developers could negatively affect competition if they restrict an FM developer from choosing other cloud providers as suppliers of accelerated compute. However, partnerships may also lead to benefits for customers by, for example, promoting innovation in the development and deployment of accelerator chips.⁸⁶⁹

4.65 The impact of the supply of IaaS based on accelerated compute on competition in cloud services is currently limited. It is unlikely that new providers of IaaS based on accelerated compute could meaningfully enter into markets for IaaS based on standard compute or PaaS.

4.66 The competitive position of cloud providers who already compete in the markets for IaaS based on standard compute or PaaS is impacted to only a limited extent by their competitive position in IaaS based on accelerated compute. This is because:

- (a) there is currently limited demand and supply side substitutability between IaaS based on accelerated and standard compute;
- (b) there are significant barriers to entry and expansion in the markets for IaaS based on standard compute and PaaS; and
- (c) IaaS based on accelerated compute is being demanded by a small number of customers who are primarily AI model developers. This limits the ability of providers of IaaS based on accelerated compute to cross-sell other cloud services to these customers.

Cloud providers' supply of access to FMs to their customers

4.67 By supplying customers with access to FMs alongside additional tools to support the building of FM-enabled applications, cloud providers enable their customers to deploy FMs either internally (eg to analyse data) or to their customers in commercialised products and services.

4.68 These services are PaaS. As set out in chapter 3, for the purposes of this market investigation we consider PaaS services in aggregate.

4.69 We have considered how the provision to customers of access to FMs could impact competition in the supply of cloud services. In particular, we consider:

⁸⁶⁹ Responses to the CMA's information requests issued [X].

- (a) how cloud providers compete to provide customers with access to FMs;
- (b) whether this affects competition in the supply of cloud services;
- (c) the impact of cloud customers' use of FMs on switching and multi-cloud; and
- (d) providers' relative strengths in the provision of access to FMs.⁸⁷⁰

How cloud providers compete to supply customers with access to FMs

4.70 Cloud providers compete to supply customers with access to FMs by:

- (a) developing and commercialising their own ('first-party') FMs; and
- (b) offering a range of third-party FMs through their managed platforms.

4.71 Microsoft, AWS and Google's FMs are often only available on their own cloud platforms. In contrast, Microsoft, AWS and Google offer customers access to a wide range of third-party FMs through their managed platforms.⁸⁷¹ Microsoft is currently the only provider to offer access to OpenAI's FMs.⁸⁷²

4.72 A cloud provider's internal document from August 2024 states that the effort to expand access to third-party FMs is driven by 'strong customer demand for using multiple models from different providers'. [REDACTED] noted that 'more than [REDACTED] of [REDACTED] customers are using multiple models as they want the flexibility to choose the right model for each task and capitalise on the latest advancements.' To address this, the cloud provider said it is was 'expanding model choice [REDACTED]'.⁸⁷³

4.73 Another cloud provider's internal strategy document from September 2024 noted that to be competitive in the 'AI era' it needed to [REDACTED].⁸⁷⁴

4.74 Cloud providers publicly promote the quantity and range of FMs available on their platforms:

⁸⁷⁰ We do not consider the supply of wider FM-enabled 'AI services' via cloud platforms, ie the integration of FMs into existing or new products or services, including SaaS. Examples of FM-enabled AI services include AI chatbots and assistants (eg Microsoft Copilot or Google Gemini) or tools and APIs that enable cloud customers to build AI-enabled applications. Supply of these AI services is nascent and evolving, and we have not examined their impact on competition in the supply of cloud services.

⁸⁷¹ Examples of cloud providers' managed platforms include: [Azure OpenAI Service | Microsoft Azure](#), accessed on 23 April 2025, [Build Generative AI Applications with Foundation Models - Amazon Bedrock - AWS](#), accessed on 23 April 2025, [Model Garden on Vertex AI | Google Cloud](#), accessed on 23 April 2025, [Generative AI Service | Oracle United Kingdom](#), accessed on 23 April 2025, and [IBM watsonx](#), accessed on 23 April 2025.

⁸⁷² Responses to the CMA's information requests [REDACTED].

⁸⁷³ [REDACTED] response to the CMA's information request [REDACTED].

⁸⁷⁴ Microsoft's response to the CMA's information [REDACTED].

- (a) AWS announced in December 2024 that over 100 foundation models were available on its Amazon Bedrock Marketplace.⁸⁷⁵ [REDACTED].⁸⁷⁶
- (b) In its earnings conference for Q1 FY24, Microsoft expressed its aim to continue attracting the best selection of proprietary and open-source models to the Azure ecosystem.⁸⁷⁷ As of May 2025, the Azure AI Model Catalog said it provided access to 1700 foundation models.⁸⁷⁸
- (c) Google announced in May 2025 that over 200 models were available through its Vertex AI Model Garden.⁸⁷⁹

4.75 Cloud providers' partnerships with FM developers often confer rights to offer access to FM developers' FMs on their cloud platform; this may include restrictions on offering the FMs on other cloud platforms. In particular:

- (a) Microsoft's partnership with OpenAI [REDACTED].⁸⁸⁰
- (b) AWS' partnership with Anthropic allows AWS to offer its customers early access to fine-tuning capabilities for Anthropic's models [REDACTED].⁸⁸¹
- (c) Google's partnership with [REDACTED] FM developers [REDACTED] permit Google to optimise their FMs with its cloud infrastructure.⁸⁸²

Impact of the supply of access to FMs on competition in the supply of cloud services

4.76 In this section, we considered evidence from cloud providers regarding how supplying access to FMs could potentially impact competition in the supply of cloud services.

Cloud provider views

- 4.77 Cloud provider views on the importance of offering access to FMs in competing for cloud customers have evolved over the course of our investigation.
- 4.78 When asked ahead of our working papers, AWS, Microsoft and Google said that it was not a crucial factor in competing for customers' cloud workloads:

⁸⁷⁵ [Amazon Bedrock Marketplace: Access over 100 foundation models in one place | AWS News Blog](#), accessed on 13 May 2025.

⁸⁷⁶ [REDACTED] response to the CMA's information request [REDACTED].

⁸⁷⁷ Microsoft's earnings conference for Q3 of FY24, 25 April 2024, [Microsoft FY24 third quarter earnings conference](#), accessed on 30 September 2024.

⁸⁷⁸ [Azure AI Platform—Cloud AI Platform | Microsoft Azure](#), accessed on 13 May 2025.

⁸⁷⁹ [Welcome to Google Cloud Next '25 | Google Cloud Blog](#), accessed on 13 May 2025.

⁸⁸⁰ [REDACTED] response to the CMA's information request [REDACTED] Annex 029-2023.

⁸⁸¹ [REDACTED] response to the CMA's information request [REDACTED].

⁸⁸² [REDACTED] response to the CMA's information request [REDACTED].

- (a) AWS said that customers generally seek the best tool for each of their IT needs and this means that they select IT service providers on a workload-by-workload basis; for example, if a customer is looking for a storage solution, the ability to provide FMs is irrelevant to their needs. For this reason, AWS did not consider the ability to provide customers with access to FMs to be a key aspect of competing for customers in supplying public cloud.⁸⁸³
- (b) Microsoft said that, although it expected AI would become an important feature of its cloud offering, it would not be the driving force of competition for public cloud services in the foreseeable future.⁸⁸⁴ Microsoft also said that the availability of a particular FM does not drive the adoption of a provider's broader services, as access to FMs is not dependent on customers' choice of public cloud providers.⁸⁸⁵
- (c) Google said that cloud providers do not need to provide customers with a large range of FMs to compete effectively in the supply of cloud infrastructure services.⁸⁸⁶

4.79 Following publication of our working papers:

- (a) Google separately submitted to the CMA that FMs are relevant to a customer's choice of not only AI/FM platform, but also wider cloud services. It further submitted that the availability of third party FMs on the cloud provider's managed platform impacts the cloud provider's direct ability to compete for PaaS customers. It added that customers may purchase other PaaS services as complementary to their access to an FM, such as data and analytics tools and cloud infrastructure services (such as storage, network and compute).⁸⁸⁷
- (b) Two smaller cloud providers submitted that supplying access to FMs is becoming increasingly important in competing for cloud customers:
 - (i) One cloud provider considered that having an FM-related cloud offering, such as a managed platform, is currently 'very important' to compete for IaaS and PaaS customers. The cloud provider submitted that

⁸⁸³ AWS' response to the CMA's information request [§].

⁸⁸⁴ In particular, Microsoft, referenced the following statistics: (a) In 2023 [§] of its cloud revenue was accrued from AI related services [§]. (b) As of March 2024, [§] of all Azure customers had some monthly spend on its AI services and [§] spent more than [§] per month on certain third party AI solutions in any month between January 2023 and March 2024 [§]. [§] response to the CMA's information request [§].

⁸⁸⁵ Microsoft, said that FMs are generally made available through APIs that can be called from anywhere, whether it be a different public cloud or from a customer's premises [§]. Microsoft's response to the CMA's information request [§].

⁸⁸⁶ Google, said that this was because 'the scale of opportunity presented by existing on-premises workloads from traditional enterprises exceeds that of workloads relating to the development and deployment of FMs and GenAI applications' and there is a wide and growing range of FM providers competing in the provision of FMs, with 'FM customers often procuring FM access and related AI development and operations tools from these independent FM providers on an infrastructure-agnostic basis'. Google's response to the CMA's information request [§].

⁸⁸⁷ Google's response to the CMA's information request [§].

‘customers are looking for more and more AI enabled services’ and ‘an AI offering is very important to compete’.⁸⁸⁸

- (ii) Another provider said that the ability to provide customers with FMs will be an important aspect of competing for customers, as ‘demand for AI services has exploded since the release of ChatGPT in 2022 and many clients are asking for these services from [cloud providers]’.⁸⁸⁹

4.80 Since publication of our PDR, Microsoft has submitted that AI capabilities are quickly becoming the key functional consideration customers base purchasing decisions on.⁸⁹⁰

4.81 Although Microsoft did not specify access to FMs in its submissions, evidence set out above (when discussing cloud provider revenue related to AI) demonstrates providing cloud customers with access to OpenAI’s FMs is an important component of Microsoft’s AI offering. As discussed above, the evidence shows that AWS and Microsoft’s revenue from providing access to FMs has been increasing and is forecast to continue increasing, but that for the foreseeable future this is likely to remain a relatively small proportion of their overall revenue.

4.82 Microsoft and Google’s internal documents and public statements discuss how rising demand for access to FMs has or could translate into increased demand for other cloud services:

- (a) In its Q1 FY24 earnings conference, Microsoft stated that ‘AI just doesn’t sit on its own’, while ‘AI projects obviously start with calls to FMs, [...] they also use a vector database. [...] We are seeing adjacent services in Azure that get attached to AI.’⁸⁹¹ In Microsoft’s Q2 FY24 earnings call it noted that over one-third of Azure AI customers were new to Azure over the past 12 months and in recent earnings calls Microsoft has noted the demand for AI services continues to be higher than available capacity.⁸⁹²
- (b) A September 2024 memo from the CEO of Microsoft to the Board of Directors and Senior Leadership Team regarding Microsoft’s Growth Strategy, Microsoft noted that [REDACTED].⁸⁹³
- (c) In an October 2023 internal presentation [REDACTED], Microsoft noted that it would continue to [REDACTED].⁸⁹⁴

⁸⁸⁸ [REDACTED] response to the CMA’s information request [REDACTED].

⁸⁸⁹ [REDACTED] response to the CMA’s information request [REDACTED].

⁸⁹⁰ [Microsoft’s response to the CMA’s PDR dated 24 January 2025](#), paragraph 17.

⁸⁹¹ [Microsoft’s earnings conference for Q3 of FY24](#).

⁸⁹² [Microsoft’s earnings conference for Q2 FY24](#), pages 4 and 16; [Microsoft’s earnings conference for Q4 FY24](#), page 17; [Microsoft’s earnings conference for Q2 FY25](#), page 17

⁸⁹³ Microsoft’s response to the CMA’s information request [REDACTED].

⁸⁹⁴ Microsoft’s response to the CMA’s information request [REDACTED].

- (d) In a September 2023 [X], Google expressed its aim to [X].⁸⁹⁵ In a September 2024 strategy document, Google claimed that [X].⁸⁹⁶

Impact of cloud customers' use of FMs on switching and multi-cloud

- 4.83 In this section we consider how cloud customers' use of FMs could impact their ability or incentives to adopt a multi-cloud strategy, for example to access FMs across different providers.

Cloud provider views

- 4.84 AWS and Microsoft said that customers' growing interest in accessing FMs may lead to a higher use of multi-cloud architectures:
- (a) AWS said that cloud providers are generally incentivised to support the interoperability of FMs produced by third party developers with their own services, as if they cannot reasonably interoperate with the third-party FMs customers wish to use, customers will either stay with their current services provider or choose an alternative.⁸⁹⁷
 - (b) Microsoft said that 'Generative AI applications are particularly suited to multi-cloud strategies as the cost of processing data by querying an FM is larger than the cost of transmitting the data across clouds'. Microsoft stated that 'this decreases the relative cost of hosting an application on [a given provider's cloud] while integrating AI features through an independent FM developer (...), or instead through another cloud provider's service'.⁸⁹⁸
 - (c) Microsoft also said that, due to the fast pace at which FMs are developing, customers are [X].⁸⁹⁹
- 4.85 In contrast, other cloud providers said that there may be barriers to adopting a multi-cloud strategy for accessing FMs:
- (a) One cloud provider said that 'increased demand for multiple AI solutions should, in theory, push more customers to implement multi-cloud solutions'. However, 'pre-existing business decisions relating to AI services that artificially impede customer choice may prevent this ideal becoming a reality, without meaningful intervention'.⁹⁰⁰ [X] also said that if a cloud provider were

⁸⁹⁵ Google's response to the CMA's information request [X].

⁸⁹⁶ Google's response to the CMA's information request [X].

⁸⁹⁷ As examples of the variety of actions AWS was taking to ensure the interoperability of FMs, AWS referenced how Amazon SageMaker provides customers with a broad set of tools (ie Amazon SageMaker Data Wrangler and Amazon SageMaker Groundtruth) to build, train and deploy FMs efficiently and enable the integration of open-source solutions. See: AWS' response to the CMA's information request [X].

⁸⁹⁸ Microsoft's response to the CMA's information request [X].

⁸⁹⁹ Microsoft's response to the CMA's information request [X].

⁹⁰⁰ [X] response to the CMA's information request [X].

to ‘continue to move toward a more exclusive business model’ and prevent competitors from accessing certain FMs unless using that specific cloud provider’s products then ‘the flywheel effects of having exclusive access and an entrenched customer base would lead to lock-in’.⁹⁰¹

- (b) One cloud provider said that ‘as and when the use of AI becomes more prevalent among traditional enterprises’, customer decisions about which AI services to use will be driven by where their legacy data and applications are hosted ie their existing cloud provider. It added that the business processes it expects to be ‘augmented by AI’ include ‘core enterprise functions, such as human resources, customer service, supply chain, finance and accounting’, which are built on customers’ ‘existing data, systems and workflows that are either on-premises or hosted in the cloud’.⁹⁰²
- (c) The same cloud provider said that Microsoft’s use of exclusivity arrangements in its partnership with OpenAI has prevented OpenAI from distributing FMs through any other third-party channel, limiting customers’ ability to multi-cloud.⁹⁰³
- (d) In an October 2023 internal document, one cloud provider discussed customer feedback that said ‘once [a product is built] on one [FM] stack that works, it becomes very hard to convince [developers] to try something else’. However, the same document noted that ‘most customers are building some kind of abstraction layer to allow them to plug and play between the different foundation models that are out there’.⁹⁰⁴
- (e) A different cloud provider said that co-locating customers’ workloads and data with the FMs they want to use will likely benefit customers.⁹⁰⁵

Customer views

- 4.86 The Jigsaw report found that, although customers said that competition is driving innovation in AI, some also felt that eventually they would need to commit to a given provider’s AI service. Customers expressed concerns that this will mean they will be locked into the chosen AI service and the cloud where the AI service sits. In particular, one customer mentioned that the optimisation of certain AI models on certain clouds would make it very difficult from their understanding to move somewhere else.⁹⁰⁶

⁹⁰¹ [redacted] response to the CMA’s information request [redacted] submission to the CMA [redacted].

⁹⁰² [redacted] submission to the CMA [redacted].

⁹⁰³ [redacted] response to the CMA’s information request [redacted].

⁹⁰⁴ [redacted] response to the CMA’s information request [redacted].

⁹⁰⁵ [redacted] response to the CMA’s information request [redacted].

⁹⁰⁶ [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), paragraph 8.1.7

- 4.87 Separately, one customer said that accessing OpenAI's FMs from Azure 'has limited the choice of other cloud vendors for integration reasons'.⁹⁰⁷ Another customer said that 'it is easier to implement OpenAI/Microsoft FMs because we use Microsoft cloud service provider'.⁹⁰⁸
- 4.88 When asked about whether they would consider switching to, or increasing their use of, Azure (whether IaaS or PaaS), for the purposes of accessing – or better accessing – OpenAI's FMs, one customer said that if it were to start to 'rely more heavily on OpenAI's FM, then that would definitely be a factor in selecting or shifting to a new cloud service provider such as Azure. There are advantages to using third-party FM tools integrated and hosted within the same cloud service provider as the datastore you are using'.⁹⁰⁹

Cloud provider internal documents

- 4.89 Providers' internal documents discuss how the supply of access to FMs may raise barriers to switching or multi-cloud amongst customers:
- (a) In an undated competitive analysis of the supply of FMs, Microsoft identified [REDACTED].⁹¹⁰ Additionally, an October 2023 presentation on Microsoft's [REDACTED].⁹¹¹
 - (b) Similarly, an August 2024 AWS planning document for its ML infrastructure and GenAI businesses noted that [REDACTED].⁹¹²
 - (c) In another 2024 AWS planning document, AWS said that that customer investments in generative AI often [REDACTED].⁹¹³

Cloud providers' relative strength in the provision of access to FMs

- 4.90 We have assessed evidence from customers and cloud providers' internal documents regarding providers' relative strengths in the supply of access to FMs.

Microsoft's strength in the provision of access to FMs

- 4.91 Microsoft's internal documents indicate that it views its own strengths in the provision of access to FMs as being:

⁹⁰⁷ [REDACTED] response to the CMA's information request [REDACTED].

⁹⁰⁸ [REDACTED] response to the CMA's information request [REDACTED].

⁹⁰⁹ [REDACTED] response to the CMA's information request [REDACTED].

⁹¹⁰ Microsoft's response to the CMA's information request [REDACTED].

⁹¹¹ Microsoft's response to the CMA's information request [REDACTED].

⁹¹² AWS' response to the CMA's information request [REDACTED].

⁹¹³ AWS' response to the CMA's information request [REDACTED].

(a) [REDACTED].⁹¹⁴ [REDACTED].⁹¹⁵ [REDACTED].⁹¹⁶

(b) [REDACTED].⁹¹⁷

- 4.92 FM developers and cloud providers' revenues from the supply of first-party FMs and ancillary services demonstrate that OpenAI's FMs account for a [REDACTED] share of FM revenues in 2023 and up to June 2024.⁹¹⁸
- 4.93 Anthropic and Google were OpenAI's next largest competitors, in terms of revenue from the provision of FMs in 2023 and up to June 2024.⁹¹⁹ Customers also considered OpenAI's FMs as the best FMs currently available, with Anthropic, Google and Meta's FMs being the closest competitors.⁹²⁰
- 4.94 Since 2024, Google has reduced prices of the first-party FMs offered on its cloud platform.⁹²¹ In a September 2024 strategy document, Google noted that [REDACTED].⁹²²

AWS' strength in the provision of access to FMs

- 4.95 AWS' internal documents indicate that AWS views its own strengths in the provision of access to FMs as being:
- (a) Model choice - [REDACTED].⁹²³ [REDACTED].⁹²⁴ [REDACTED].⁹²⁵
- (b) Its emphasis on 'strong security [REDACTED].⁹²⁶
- 4.96 In its September 2023 update to Amazon's Board of Directors, AWS stated that, although [REDACTED].⁹²⁷
- 4.97 However, in an August 2024 planning document, AWS noted [REDACTED].⁹²⁸

Google's strength in the provision of access to FMs

- 4.98 Google's internal documents indicate that it perceives itself to have a strong offering in the supply of access to FMs, predominantly due to its own family of

⁹¹⁴ Microsoft's response to the CMA's information request [REDACTED].

⁹¹⁵ Microsoft's response to the CMA's information request [REDACTED].

⁹¹⁶ Microsoft's response to the CMA's information request [REDACTED].

⁹¹⁷ Microsoft's response to the CMA's information request [REDACTED].

⁹¹⁸ [REDACTED]. Responses to the CMA's information requests [REDACTED].

⁹¹⁹ Responses to the CMA's information requests [REDACTED].

⁹²⁰ Customers were asked to score other FM developers in terms of how close of a competitor their FMs are to OpenAI's.

Meta provides open access to its FMs. Responses to the to the CMA's information requests [REDACTED].

⁹²¹ [Lower costs and 100+ new languages coming to Gemini 1.5 | Google Cloud Blog](#), accessed on 14 July 2025

⁹²² Google's response to the CMA's information request [REDACTED].

⁹²³ AWS' response to the CMA's information request [REDACTED].

⁹²⁴ AWS' response to the CMA's information request [REDACTED].

⁹²⁵ AWS' response to the CMA's information request [REDACTED].

⁹²⁶ AWS' response to the CMA's information request [REDACTED].

⁹²⁷ AWS' response to the CMA's information request [REDACTED].

⁹²⁸ AWS' response to the CMA's information request [REDACTED].

FMs, Gemini, which is often benchmarked against [REDACTED] models.⁹²⁹ However, the document also noted that [REDACTED].⁹³⁰

- 4.99 Regarding its third-party FM offering, Google noted in an internal document discussing the 'strategic rationale' for its partnership with Anthropic that [REDACTED].⁹³¹
- 4.100 Microsoft said that Google has data advantages (including from its Search, Android and YouTube products as well as through deals with third-party data providers) that confer unique competitive advantages in developing first-party FMs.⁹³²

Smaller providers' strength in the provision of access to FMs

- 4.101 We understand the supply of access to FMs by smaller cloud providers to be very limited. For example:
- (a) Oracle does not provide a first-party FM and only provides access to a limited third-party FMs developed by Cohere and Meta.⁹³³
 - (b) IBM has developed its own series of 'open' FMs (Granite models).⁹³⁴ IBM provides access to a limited number of third-party FMs developed by Mistral, Meta, Google and a small number of other developers.⁹³⁵
 - (c) CoreWeave does not provide access to FMs.⁹³⁶
 - (d) Lambda Labs provides access to a limited number of open source FMs.⁹³⁷

Our assessment of cloud provider' relative strengths in the provision of access to FMs

- 4.102 We have considered the competitive landscape for the provision of access to FMs and our evidence shows that:
- (a) Providing customers with access to FMs is contributing to the growth in Microsoft, AWS and Google's global revenues, although these revenues are currently relatively low.

⁹²⁹ Google's response to the CMA's information request [REDACTED].

⁹³⁰ Google's response to the CMA's information request [REDACTED].

⁹³¹ Google's response to the CMA's information request [REDACTED].

⁹³² Microsoft's response to the CMA's information request [REDACTED]; Microsoft's submission to the CMA [REDACTED].

⁹³³ [Pretrained Foundational Models in Generative AI](#), accessed on 23 April 2025.

⁹³⁴ [Granite | IBM](#), accessed on 1 May 2025.

⁹³⁵ [Foundation Models - IBM watsonx.ai](#), accessed on 23 April 2025.

⁹³⁶ [The AI Hyperscaler for GPU Cloud Computing | CoreWeave](#), accessed on 27 May 2025.

⁹³⁷ [Inference | Lambda](#), accessed on 27 May 2025.

- (b) Cloud providers compete to provide customers with access to FMs by developing their own exclusive models and by improving the range and performance of third-party models offered through their managed platforms.

4.103 Microsoft, AWS and Google have strong competitive positions, each supplying several FM developers as well as developing their own FMs. The competitive constraint from smaller cloud providers is limited.

- (a) The competitive conditions in the supply of access to FMs could evolve in a number of different ways:
 - (i) Partnerships between FM developers and cloud providers could affect competitive conditions in the supply of access to FMs. These partnerships could negatively affect competition in the supply of access to FMs if cloud providers restrict access to a particularly popular FM (or offer access only to an inferior version of the FM) with limited comparable alternatives.
 - (ii) We have heard a range of different views about the likely impact of access to FMs on multi-cloud, with some cloud providers and customers expressing concerns about access to FMs increasing barriers to operating a multi-cloud strategy.

4.104 We note that, as set out above, whilst the demand for access to FMs has been rising, the revenues associated with this are currently relatively low. We considered that this means that the scope for demand for access to FMs to translate into greater demand for other cloud services is currently limited. We recognise that, in the future, access to FMs may become important to a greater number of customers. We therefore consider how the role of AI (and access to FMs) in customers' decision-making and the impact of AI capabilities on competition in cloud infrastructure services may change in the future in the next section.

The impact of AI capabilities on competition between cloud providers

4.105 We have examined the extent to which cloud providers' AI capabilities impact customers' decisions on cloud services, as well as the providers' abilities to meet customers' demand for AI. Our assessment covers all layers of cloud services – this includes accelerated compute and the provision of FMs as well as AI-enabled PaaS and SaaS services.

4.106 We have considered evidence from cloud providers that relates to their UK and global businesses alongside evidence received from customers specifically relating to their public cloud computing usage in the UK.

The impact of AI on how customers choose their cloud provider

Cloud provider submissions

4.107 Cloud providers' submissions varied regarding whether and how AI impacts customers' decision-making relating to cloud:

- (a) Microsoft submitted that AI capabilities are quickly becoming the key functional consideration customers base purchasing decisions on.⁹³⁸ Microsoft separately submitted that AI capabilities and the ability to modernise workloads on a cloud provider are increasingly important factors customers consider when selecting cloud provider and that 'Google's leadership in AI is helping it win customer migrations'.⁹³⁹
- (b) Google submitted that the opportunities offered by AI will not fundamentally alter the contestability of the cloud infrastructure market in the near future.⁹⁴⁰ It submitted that 'traditional enterprise migration will remain the primary driver of cloud demand in the short-to-medium term' and that although 'traditional enterprises are the key driver of overall cloud demand, traditional enterprises are spending [redacted] on AI services'. The cloud provider added that customers for whom AI is likely a key functional consideration of their cloud purchasing decisions, such as 'AI natives and AI-first digital natives', represent [redacted] the overall addressable market.⁹⁴¹
- (c) Google also said that traditional enterprise customers are more likely to purchase AI services from the cloud provider that is already hosting the relevant apps and workflows because they are widely expected to use generative AI to augment their existing apps and processes.⁹⁴²
- (d) Oracle said that customers are still navigating AI and trying to understand the value of AI to their business and their data, but that AI will impact competition in SaaS and PaaS because modernisation and migration from on-premises is happening as part of organisations' AI strategies. Oracle added that, generally, existing strength in cloud services may enable providers to win in AI because their customers will look to their current providers for new AI services; Oracle noted that this would especially be the case if the data the customer wanted to use was stored within their current provider's cloud.⁹⁴³

⁹³⁸ Microsoft's response to the CMA's PDR dated 24 January 2025, paragraph 17.

⁹³⁹ Microsoft's submission to the CMA [redacted].

⁹⁴⁰ Google's submission to the CMA [redacted].

⁹⁴¹ Google submitted that, according to [redacted], traditional enterprises were [redacted]% of the worldwide cloud services total addressable market in 2024 ([redacted]% in UK). Google's submission to the CMA [redacted].

⁹⁴² Google's response to the CMA's information request [redacted].

⁹⁴³ Note of meeting with Oracle [redacted].

- (e) AWS submitted that AI has the potential to dramatically change cloud providers' market positions over the next one to five years.⁹⁴⁴ AWS also said that it is impossible to predict how successful AI entrants would be, but that AI is driving a lot of competition and is going to disrupt the cloud space.⁹⁴⁵
- (f) OVHcloud submitted that AI is expected to be a key driver of the cloud market's growth in the coming years.⁹⁴⁶

4.108 Cloud providers also submitted industry reports that discussed the impact of AI on customers' decisions relating to cloud services, including:

- (a) Microsoft submitted a report from cloud solutions company, Rackspace, that outlined findings from a survey it had conducted of IT teams, including that 84% of respondents reported taking steps to integrate their AI and cloud strategies to enhance operational efficiency or improve the insights generated by data analytics.⁹⁴⁷ We consider that the option of 'taking steps to integrate AI and cloud strategies' used in this survey is very broad and potentially captures a wide variety of different actions a customer could take with some being more significant than others.

4.109 Microsoft also submitted an internal presentation discussing quantitative surveys relating to the impact of AI on customers' views and decisions relating to cloud services. The document set out Microsoft staff's findings and views in relation to this topic, including that [REDACTED].⁹⁴⁸

Customer views

4.110 The Jigsaw report found that the general expectation amongst customers was that the main public cloud providers will – if they do not already – offer an AI PaaS⁹⁴⁹ solution and compete strongly for this AI PaaS business.⁹⁵⁰ AI was also mentioned by the majority of participants in Jigsaw's research, either because it is a field their business is directly involved in, or because it is seen as an emerging technology that they are looking at and expecting to make use of in the near future. Regardless of the context, these participants all said that AI was likely to be the most significant change to public cloud services over the next five years or so.⁹⁵¹

⁹⁴⁴ [Summary of hearing with AWS, 12 March 2025](#), paragraph 10.

⁹⁴⁵ Transcript of hearing with AWS [REDACTED].

⁹⁴⁶ [OVHcloud's response to the CMA's PDR dated 28 January 2025](#), paragraphs 6 and 7.

⁹⁴⁷ [New Research by Rackspace Technology Reveals Hybrid Cloud and AI Integration as Key Drivers for IT Innovation in 2025](#), accessed on 14 May 2025.

⁹⁴⁸ Microsoft's submission to the CMA [REDACTED].

⁹⁴⁹ The Jigsaw report did not provide a definition of AI PaaS in this context, however, as set out above in the Introduction and background, AI PaaS can refer to different platforms designed to assist developers in building AI models themselves or using existing AI models to build/enrich applications. Other PaaS services may also be powered or enriched by AI.

⁹⁵⁰ [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), page 89.

⁹⁵¹ [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), paragraph 8.1.1.

- 4.111 The Jigsaw report said that AI was discussed by participants in the context of potentially being, in future, the most common reason why an organisation who is either operating a single cloud model, or a siloed multi-cloud model, may wish to move to a more integrated multi-cloud model. The Jigsaw report outlined that the reasons for this would be if an organisation's sole/main public cloud provider is different to the provider of the AI PaaS service they wish to use.⁹⁵²
- 4.112 However, these customers also recognised that there are certain challenges to this approach, including the large amounts of data transfer needed (and the corresponding spend on egress fees) and the potential technical barriers preventing interoperability between clouds.⁹⁵³ The Jigsaw report noted that the implications of the adoption of an AI solution outside of a business' main cloud provider is not as simple as deciding who has the best AI offer or machine learning capabilities.⁹⁵⁴
- 4.113 The Jigsaw report found that the prevailing view among participants was that they will largely be locked into the AI service they have already chosen. One customer noted that 'in order to roll out AI as quickly as possible, you almost have to use the machine learning services from AWS to get you there very quickly. That means that you are locked into AWS'.⁹⁵⁵
- 4.114 Earlier in our investigation, we separately heard from large customers of public cloud providers that, generally, they regarded AI capabilities as a material but not very important factor when choosing their main cloud provider. However, some large customers said that AI capabilities were expected to become increasingly important in the future.⁹⁵⁶ Some customers said that, where necessary, they can access the AI services from providers other than their main one.⁹⁵⁷
- 4.115 Given the potential for AI within cloud services to have evolved over the course of our investigation, we asked the same customers after the publication of our PDR about the importance of providers' AI capabilities when choosing their main cloud provider (the following discussion of customer views relates to this more recent evidence). We also asked these customers what they considered to be 'AI capabilities' and how important it was to their organisation that the AI services that they purchase are provided by their main public cloud provider.⁹⁵⁸ We found that:

⁹⁵² [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), paragraph 8.1.2.

⁹⁵³ [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), paragraphs 8.1.2-8.1.3.

⁹⁵⁴ [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), paragraph 8.1.4.

⁹⁵⁵ [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), paragraphs 8.1.7-8.1.8

⁹⁵⁶ Responses to the CMA's information requests [REDACTED].

⁹⁵⁷ Responses to the CMA's information requests [REDACTED].

⁹⁵⁸ Responses to the CMA's information requests [REDACTED].

- (a) Generally, these customers considered AI capabilities to be about access to FMs, as well as AI-enabled SaaS products and AI PaaS services. Some considered AI capabilities to include AI infrastructure.
 - (b) These customers' views on the importance of AI capabilities when choosing their main cloud provider generally had not changed since they were asked at an earlier stage of the investigation, although they were more important for some customers today relative to when they were previously asked.
- 4.116 Some customers that said AI capabilities were a more important factor when choosing their main cloud provider explained that this was because AI was of increasing strategic importance to their organisation.⁹⁵⁹
- 4.117 The customers that said AI capabilities were not a very important factor cited a range of reasons, including that:
- (a) AI capabilities do not drive adoption of a main cloud provider and/or other factors are more important;⁹⁶⁰
 - (b) All major cloud providers have strong/similar AI capabilities and/or are advancing at pace, so AI capabilities are not a determining factor;⁹⁶¹
 - (c) AI is in its infancy and is (currently) of limited use to their organisation;⁹⁶² and
 - (d) Their organisation does not predominantly source AI services from cloud providers.⁹⁶³
- 4.118 We also asked these customers whether it was important to their organisation that AI services they purchase are provided by their main public cloud provider. When purchasing AI services, these customers had mixed views on this: some customers found it important, others did not. Some of these customers referenced proximity to data/applications and/or operational integration as a reason to consider purchasing AI services from their main cloud provider.⁹⁶⁴
- 4.119 We also consider AI as a potential driver for modernisation in chapter 7: Licensing. We find that evidence from providers and customers shows that AI services so far have had a limited effect in making modernisation away from Microsoft software significantly easier.
- 4.120 The evidence discussed here regarding how AI capabilities impact cloud decisions is based on large customers' views regarding decisions they make about their

⁹⁵⁹ Responses to the CMA's information requests [X].

⁹⁶⁰ Responses to the CMA's information requests [X].

⁹⁶¹ Responses to the CMA's information requests [X].

⁹⁶² Responses to the CMA's information requests [X].

⁹⁶³ Response to the CMA's information request [X].

⁹⁶⁴ Responses to the CMA's information requests [X].

main cloud provider (ie at the customer level). We recognise that this may not cover those instances where customers consider demand only for specific workloads.⁹⁶⁵ However, we consider the findings presented above are also relevant to considering the workload level because some of the reasons given by large customers for considering AI capabilities to be not very important would also be applicable at the workload level.⁹⁶⁶ Furthermore, some customers said it was important to their organisation that the AI services they purchase are provided by their main cloud provider – this would apply to new workloads.

- 4.121 To understand further the extent to which competition occurs at the workload level, we have considered providers' strengths in competing for customer demand for AI below.

Cloud providers' relative strengths in AI and the impact on their provision of other cloud services

- 4.122 In this section, we have examined cloud providers' strength in competing for customer demand for AI (outside of strength in supplying accelerated compute and access to FMs, which have been considered in earlier sections). In particular, we have considered potential advantages arising from cloud providers being vertically integrated across cloud services and adjacent sectors.

Cloud provider views

- 4.123 Microsoft submitted that Google, AWS, and Microsoft had all pivoted the focus of their cloud strategies to AI and that each company had unique strengths, with Google being the only one with a 'fully-integrated first-party AI stack'.⁹⁶⁷ Microsoft also submitted that Google has 'control' of 'critical' AI training data and consumer distribution channels and is vertically integrated including AI hardware and FMs; Microsoft said that these 'AI advantages' contribute to a 'competitive advantage that neither Amazon nor Microsoft enjoy'.⁹⁶⁸
- 4.124 Google said that, to the extent AI grows in importance, it was wrong for anyone to suggest Google is particularly well-positioned to benefit.⁹⁶⁹ Google also said that it faces strong competition at each layer of the 'AI stack' and [REDACTED].⁹⁷⁰

⁹⁶⁵ As discussed in the Licensing chapter, some demand is competed at a customer level and some at a workload level (ie on a workload-by-workload basis).

⁹⁶⁶ Reasons given by large cloud customers that may apply at the workload level include: all major cloud providers have strong/similar AI capabilities and/or are advancing at pace, so AI capabilities are not a determining factor; AI is in its infancy and is (currently) of limited use to their organisation; and their organisation does not predominantly source AI services from cloud providers

⁹⁶⁷ Microsoft's submission to the CMA [REDACTED].

⁹⁶⁸ [Microsoft's response to the CMA's PDR dated 28 January 2025](#), paragraph 18.

⁹⁶⁹ [Summary of hearing with Google, 17 March 2025](#), paragraph 18.

⁹⁷⁰ Google's submission to the CMA [REDACTED].

- 4.125 Google submitted that the opportunities offered by AI will not fundamentally alter the contestability of the cloud infrastructure market in the near future:⁹⁷¹
- (a) It submitted that ‘AI technology is still very nascent’ and that ‘both cloud providers and their customers are still exploring how they can gainfully adopt AI’, referencing a January 2025 report by McKinsey that claimed ‘only 1 percent of leaders call their companies “mature” on the deployment spectrum, meaning that AI is fully integrated into workflows and drives substantial business outcomes’. [REDACTED].
 - (b) [REDACTED]. It referenced the release of the efficient and cheaply-priced Chinese FM DeepSeek, which caused a drop in the market value of Nvidia, Microsoft and Google.⁹⁷² It said that customers are already spending [REDACTED] on AI services, given the innovations in cheaper models.⁹⁷³
- 4.126 We note that the CMA’s AI Foundation Models initial review found that both Google and Microsoft were active across all layers of the FM value chain, including holding strong positions in compute, data, FM development and release, and deployment in downstream markets.⁹⁷⁴
- 4.127 Microsoft compared its own announcement to invest \$80billion in FY2025 to build AI-enabled data centres with Google’s planned \$75billion of capex for 2025 ‘to accelerate its AI ambitions’.⁹⁷⁵ However, Google said that the \$75billion capex announcement related to all of Alphabet, with [REDACTED] expected to be allocated to Google Cloud technical infrastructure investments.⁹⁷⁶
- 4.128 OVHcloud submitted that there are risks linked to the emergence of AI with ‘potential to distort competition in the market for cloud services’:
- (a) ‘hyperscalers’ have privileged access to AI inputs;
 - (b) replication of unfair practices observed in the cloud market (such as reinforcement of cloud credits programs, mechanisms to prevent the migration of AI models to other CSPs once they have been trained or fine-tuned on an initial CSP or forcing AI solutions to their customers, with such solution being deployable solely on their cloud infrastructures); and
 - (c) partnerships giving a company unfair access to key AI inputs at the expense of other market players.⁹⁷⁷

⁹⁷¹ Google’s submission to the CMA [REDACTED].

⁹⁷² BBC, [Nvidia shares sink as Chinese AI app spooks markets](#) (27 January 2025)

⁹⁷³ The Information, [Some Large Cloud Customers Slow Down AI Spending as Prices Drop](#) (29 March 2025).

⁹⁷⁴ [AI Foundation Models: initial review](#). See the Update Paper published on 11 April 2024

⁹⁷⁵ [Microsoft’s response to the CMA’s PDR dated 28 January 2025](#), paragraph 67.

⁹⁷⁶ Google’s response to the CMA’s information request [REDACTED].

⁹⁷⁷ [OVHcloud’s response to the CMA’s PDR dated 28 January 2025](#), paragraphs 6 and 7.

Cloud provider internal documents

- 4.129 Microsoft and Google internal documents demonstrated that they see benefits arising from vertical integration in relation to AI:
- (a) A September 2024 memo from the CEO of Microsoft to the Board of Directors and Senior Leadership Team regarding Microsoft's Growth Strategy stated that [REDACTED].⁹⁷⁸
 - (b) In a 2024 strategy document, Google said that [REDACTED].⁹⁷⁹
- 4.130 Microsoft, AWS and Google internal documents also demonstrated that these cloud providers seek to upsell customers of AI services other cloud workloads:
- (a) A September 2024 memo from the CEO of Microsoft to the Board of Directors and Senior Leadership Team regarding Microsoft's Growth Strategy stated that that [REDACTED].⁹⁸⁰
 - (b) In a 2024 planning document, AWS said that customer usage of its FM managed service, [REDACTED]. The document added that [REDACTED]. However, the document also noted AWS' [REDACTED].⁹⁸¹
 - (c) In the same 2024 planning document, AWS highlighted that [REDACTED].⁹⁸²
 - (d) In a 2024 strategy document, Google said that [REDACTED].⁹⁸³ In the same document, Google said that [REDACTED].⁹⁸⁴
- 4.131 In a September 2024 strategy document, Google Cloud said that [REDACTED].⁹⁸⁵

Customer views

- 4.132 The Jigsaw report found that customers perceive the competitive landscape for AI as being 'fiercely competed'. Although these customers perceived Microsoft and Google as currently being ahead of competitors in this space, they also said that their advantage is not fixed and may shift in the future.⁹⁸⁶ Some participants thought of Google as superior – in particular to AWS – in terms of its AI offering.⁹⁸⁷
- 4.133 As discussed in the Industry Background chapter, ahead of our working papers we heard from large customers of public cloud providers that, generally, they did not

⁹⁷⁸ Microsoft's response to the CMA's information request [REDACTED].

⁹⁷⁹ Google's response to the CMA's information request [REDACTED].

⁹⁸⁰ Microsoft's response to the CMA's information request [REDACTED].

⁹⁸¹ AWS' response to the CMA's information request [REDACTED].

⁹⁸² AWS' response to the CMA's information request [REDACTED].

⁹⁸³ Google's response to the CMA's information request [REDACTED].

⁹⁸⁴ Google's response to the CMA's information request [REDACTED].

⁹⁸⁵ Google's response to the CMA's information request [REDACTED].

⁹⁸⁶ [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), paragraph 8.1.6.

⁹⁸⁷ [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), page 30.

view Google as a fully effective alternative cloud provider to AWS or Azure. Following publication of our PDR, we asked the same large customers about Google's suitability and only one of these customers said their view of Google had improved (citing the availability of Google's first-party FMs as a reason).⁹⁸⁸

Our assessment of cloud providers' relative strengths in AI and the impact on their provision of other cloud services

- 4.134 We have considered how AI is impacting cloud customer decisions and spending. Our evidence shows that the migration and modernisation of on-premises workloads is, in part, being driven by customers' desire to use AI PaaS and SaaS services (for example, querying their company data using an FM).
- 4.135 However, we have found that, overall, AI considerations currently have a limited impact on customers' choice of main cloud provider.
- (a) Many large customers told us that AI capabilities were material but not very important when choosing their main cloud provider, including because other factors were more important or because their organisation does not source AI services from cloud providers.
 - (b) Whilst some large customers told us earlier in our investigation that they expected AI to be more important when selecting their main cloud provider in future, we have found that these customers' views on the importance of AI have not markedly changed over the past two years. We have also found that large customers' views of Google as an effective main cloud provider have been largely unchanged over the past two years, even as demand for AI has increased.
 - (c) Our analysis of revenues shows that AI-related revenue is growing substantially. However, we have found that IaaS based on accelerated compute, the largest driver of AI-related revenue growth, is concentrated amongst a relatively small number of customers, suggesting that these services are only currently important to a limited number of customers. The provision of FMs to cloud customers was, to a lesser extent, another key driver. We found that whilst this has been growing, it is currently a small proportion of revenue.
- 4.136 We have also found mixed evidence on the extent to which customers are likely to make greater use of multi-cloud to access AI services hosted on a cloud other than their main provider.

⁹⁸⁸ [REDACTED] response to the CMA's information request [REDACTED].

- (a) In our customer research AI was commonly submitted by customers as a reason for moving to a more integrated multi-cloud model if their sole/main public cloud provider is different to the provider of the AI PaaS service they wish to use. However, customers recognised the challenges with this approach including the technical and commercial barriers we have considered.
- (b) Large customers had mixed views regarding the importance of whether the services are provided by their main cloud provider. Some customers told us that it was important to their organisation that the AI services they purchase are provided by their main cloud provider.

We therefore consider the evidence we have collected from customers suggests that cloud providers' AI capabilities currently play a limited role in driving customers' choice of cloud provider.

- 4.137 Our evidence shows that Microsoft, AWS and Google are the strongest cloud competitors in AI, and each have different strengths. There is evidence that they benefit from being vertically-integrated (ie with comprehensive IaaS, PaaS, and SaaS offerings, as well as large non-cloud businesses) as well as having an ability to invest at scale, cross-subsidise, accept lower margins, and up-sell to general cloud services. Whilst Google has a comparatively weaker competitive position relative to AWS and Microsoft in cloud infrastructure services, it appears to be an equally strong competitor to them in AI related services.

Some cloud providers and customers told us that AI capabilities were likely to become more important to customers over time. However, we found that the factors which customers perceive to be relevant to AI capabilities differ to some extent across customers. Whilst customers generally agreed that AI capabilities include AI PaaS services (including access to FMs) and AI-enabled SaaS products, only some also considered it to include IaaS based on accelerated compute. Therefore, the impact AI capabilities will have on competition in cloud services, including the relative competitive positions of Microsoft, AWS and Google, will depend on how customers' demand across IaaS based on accelerated compute, AI PaaS services and AI-enabled SaaS evolves, and consequently the extent to which increased customer demand for particular AI capabilities leads to increased demand for other cloud services from the same provider.

- 4.138 Cloud providers may be able to attract additional (or retain existing) revenue on cloud services from customers who are interested in accessing a particular AI capability (eg a particular FM or FM service). However, this will depend on a range of factors, including:

- 4.139 Whether customers prefer to use a multi-cloud approach (eg by placing AI specific workloads on clouds other than their main provider), and the extent to which technical or commercial barriers to multi-cloud prevent this from happening.
- 4.140 Whether customers prefer to obtain AI capabilities from their main cloud provider, and if so, whether AI capabilities incentivise customers to switch or whether technical or commercial barriers to switching prevent this from happening.
- 4.141 The extent to which there is demand and supply side substitutability of the AI services that we have found to be separate to the cloud infrastructure services market (AI-enabled SaaS and IaaS based on accelerated compute) in future. For example, if developments in making FMs smaller and more effective caused a meaningful shift in demand from accelerated compute to standard compute, cloud providers that are strong in the provision of IaaS based on accelerated compute, such as Google, may be able to compete more effectively for cloud infrastructure services.
- 4.142 We consider that these developments mean that if, as the evidence appears to indicate, AI capabilities become a more important factor for customers, then we may see changes in competitive dynamics (potentially including a relative strengthening of Google's position in competing for new workloads) over time.

Conclusions

- 4.143 Microsoft, AWS and Google are all strong in the provision of AI-related cloud services. However, this is yet to materially alter competition in cloud services and any potential impact will take time to materialise.
- 4.144 AI capabilities are likely to become more important to cloud customers over time and so the role of AI in driving competition in cloud services could grow in the future. However, the timing of this change and whether it increases or decreases the relative strength of Microsoft, AWS and Google is dependent on a number of factors, including how customer demand evolves and the extent to which any increase in demand for AI capabilities from cloud providers leads to increased demand for other cloud services from the same provider.

5. Barriers to entry and expansion

- We have found substantial barriers to entry and expansion in the provision of cloud services, particularly for IaaS. Market entry and expansion in the supply of IaaS requires significant capital investment in fixed assets, which for many asset types would be largely irrecoverable upon exit. This combines with economies of scale, with larger providers incurring lower average costs compared to smaller providers. Unless a new entrant (or company seeking to expand) is willing to make investments of a similar magnitude to those of the largest suppliers, it is likely to face higher average costs to provide an equivalent level of service and so may struggle to compete effectively across a sufficiently wide range of cloud services. This is likely to be disincentivising IaaS market entry and expansion.
- Furthermore, given the scale of investment and expansion that the three main cloud providers have made in IaaS, any new entrant (or company seeking to expand) would need to invest substantially more than the existing large providers in order to gain competitively significant market share. The levels of investment that Microsoft, AWS and Google are making in IaaS and in AI infrastructure in the coming years may raise these barriers even higher.
- A range of investment firms have recently announced significant investments in energy infrastructure and data centres. These investments are focused on increasing the supply of data centres to support AI.
- Investment in infrastructure is a requirement for entry and expansion into IaaS, but not PaaS. However, competition in PaaS using third party IaaS likely involves sourcing IaaS from firms that are also competing in the supply of PaaS. This limits the scope for PaaS providers to compete effectively.
- While we recognise that investments by cloud providers can have pro-competitive effects and benefit customers, this does not preclude them also having the effect of deterring market entry and expansion.
- The wide product portfolios of the larger cloud providers also contribute to the barriers to entry and expansion in both IaaS and PaaS markets because range of services is an important consideration for customers selecting a cloud provider. ISVs also value access to the wider customer bases of the larger cloud providers and prioritise developing services compatible with cloud services offered by these cloud providers. In addition, customers place some importance on a provider's reputation and this contributes to overall barriers to entry and expansion in cloud services.
- Microsoft and AWS are the largest providers to the public sector. We considered that their leading positions with public sector customers are likely driven by the barriers to entry and expansion, and barriers to multi-cloud and switching we have identified

elsewhere in this report, as distinct from public procurement practices. We suggest that, as public sector use of cloud services grows, the UK government should continue to monitor the outcomes of public sector procurement of cloud services and drive best practice in the application of procurement frameworks.

- 5.1 This chapter assesses barriers to entry and expansion in cloud services markets, focusing on the extent to which the largest providers benefit from irrecoverable sunk costs, economies of scale and a large portfolio of cloud services, compared to new market entrants and smaller competitors.
- 5.2 Entry or expansion by firms will often stimulate competition in a market and the prospect of entry or expansion within a short time can countervail factors that are harming competition.⁹⁸⁹ A major source of competitive discipline in a market is therefore generally reduced or eliminated if there are barriers to market entry and expansion, whether an absolute barrier or some other form of restriction. Barriers to entry and expansion give at least some incumbent firms an advantage over efficient potential firms or rival incumbent firms, either by reducing the expected profits, or increasing the expected costs, of entry or expansion.⁹⁹⁰
- 5.3 This chapter sets out our assessment of barriers to entry and expansion in the cloud services markets. To do this, we assessed whether:
- (a) large sunk cost investments by the largest providers deter entry into the IaaS market;
 - (b) larger cloud providers benefit from economies of scale when compared to smaller competitors; and
 - (c) having a large portfolio of cloud services gives cloud providers strategic advantages over their competitors.
- 5.4 We also assessed potential barriers to entry and expansion arising from ‘cloud credits’ offered by larger cloud providers, reputational barriers that could favour the largest providers and regulatory barriers that may inhibit the growth of cloud service providers.

Sunk investment costs

- 5.5 Economies of scale, in combination with sunk investment costs, can constitute a barrier to entry or expansion.⁹⁹¹ Entry on a large scale will often entail a high risk (that sunk investment costs may not be recovered) because it will generally be

⁹⁸⁹ CC3 (Revised), [Guidelines for market investigations: Their role, procedures, assessment and remedies](#), paragraph 205.

⁹⁹⁰ CC3 (Revised), paragraph 207.

⁹⁹¹ CC3 (Revised), paragraph 212.

successful only if the firm can expand the total market significantly, or substantially replace one or more existing firms.⁹⁹²

- 5.6 In this section, we analyse the size of cloud providers' investments in fixed asset infrastructure and the extent to which some or all of these fixed costs are unavoidable 'sunk' costs that are difficult to recover upon exit.
- 5.7 We focus on large, unavoidable, 'sunk' costs because they could deter new entry into the markets, as new entrants would need to make significant upfront investments in infrastructure without any certainty of return. If their plans to enter the markets failed, they would be unable to recoup the initial investment, further increasing the risk of entry for them.
- 5.8 AWS said that one of the 'six advantages of cloud computing' for its customers of using cloud services is that they can replace fixed, capital IT expenditure (CapEx) with variable, operating expenditure (OpEx). This is because cloud providers invest in the fixed asset infrastructure required to deliver computing resources and customers only pay when they consume these resources.⁹⁹³
- 5.9 The corollary of customers moving from CapEx to OpEx is that the cloud providers must make the CapEx investments in the fixed asset infrastructure for their customers. As the CapEx is in specific assets, with limited resale value, the cloud providers' CapEx therefore represent unavoidable 'sunk' costs.

Stakeholder views on investment

- 5.10 Microsoft submitted that significant capital investment is required to participate in the cloud markets, and economies of scale matter. It said that AI is attracting a lot of investment not only from cloud providers but also from private equity firms and venture capitalists. This investment can fund competitors' expansion of data centres and entry into the cloud markets, albeit at a more niche level.⁹⁹⁴
- 5.11 AWS submitted that:
- (a) investments by the larger cloud providers are pro-competitive, not anti-competitive and that the market would not be more competitive if some providers stopped investing;⁹⁹⁵
 - (b) smaller cloud service providers can both invest and secure funding, which allows them to grow alongside the larger providers. AWS listed several companies, which include Voltage Park, Lambda Labs, Aethir, Hive, Wasabi,

⁹⁹² CC3 (Revised), paragraph 213.

⁹⁹³ Six advantages of cloud computing - Overview of Amazon Web Services, accessed on 6 November 2024.

⁹⁹⁴ Summary of hearing with Microsoft, 5 March 2025, paragraph 23.

⁹⁹⁵ AWS' Economic response to the CMA's Competitive landscape working paper (prepared by CRA) dated 2 August 2024, paragraph 116.

Alibaba, Cloudsky, Skytap, Digital Ocean and OVHcloud as evidence of competition in the cloud services markets;⁹⁹⁶

- (c) Alibaba, Huawei, Tencent and others have been investing on a global scale, this putting competitive pressure on UK cloud providers;⁹⁹⁷
- (d) new entrants can defer fixed costs to third party data centre providers and by doing so incur costs that are proportionate to their size;⁹⁹⁸ and
- (e) its capital expenditure in AI is very substantial but it does not consider itself to be a strong rival, given it only invests a small fraction of the total capital going into AI. AWS also gave examples of major investments by venture capitalists, private equity firms, sovereign wealth funds, state-backed investors and other cloud providers into cloud services and data centres.⁹⁹⁹

5.12 Google submitted that cloud was fundamentally a scale out industry where providers could start small and grow. New entrants could come into the market, for example, by leasing data centres and selling capacity that way.¹⁰⁰⁰

Our assessment

5.13 Cloud providers primarily use their fixed asset infrastructure to deliver IaaS.

5.14 PaaS requires IaaS, but providers that only offer PaaS can use IaaS provided by a third-party, so there is no requirement for PaaS providers to own their own infrastructure. Hence, investment in infrastructure is a requirement for entry and expansion into IaaS, but not PaaS.

5.15 However, competition in PaaS using third party IaaS likely involves sourcing IaaS from firms that are also competing in the supply of PaaS. This limits the scope for PaaS providers to be fully effective competitors as:

- (a) IaaS providers are more inclined to facilitate PaaS provision that is complementary to, rather than a substitute for, their own PaaS; and
- (b) losing a sale to a PaaS provider will often involve the IaaS provider continuing to earn upstream margin.

⁹⁹⁶ AWS' Economic response to the CMA's Competitive landscape working paper (prepared by CRA) dated 2 August 2024, paragraph 89.

⁹⁹⁷ [REDACTED].

⁹⁹⁸ AWS' Economic response to the CMA's Competitive landscape working paper (prepared by CRA) dated 2 August 2024, paragraph 90.

⁹⁹⁹ Summary of hearing with Amazon Web Services, 12 March 2025, paragraph 11. [REDACTED].

¹⁰⁰⁰ Summary of hearing with Google, 17 March 2025, paragraph 19.

- 5.16 Our assessment should be viewed in the context of the barriers to entry and expansion that relate to fixed asset infrastructure and economies of scale mainly manifesting in IaaS, but also indirectly extending into PaaS.
- 5.17 Cloud providers group their fixed assets under three broad headings:¹⁰⁰¹
- (a) Data centre assets: these include the shell of the data centre, costs associated with building the data centre, the infrastructure for energy and water and the land on which the data centre is built. Data centres can be owned outright by cloud providers or they can be owned by third party data centre providers. If the data centre is owned by a third-party, cloud providers can co-locate in a third-party's data centre, which involves the third party data centre provider offering a managed service or providing space on a leased basis.
 - (b) Network assets: cloud providers connect their data centres using network assets, which connect server racks to each other, to other data centres, to customers and to the internet.
 - (c) Servers and components: servers run the software to enable cloud related services. Servers and components include racks and the components on these, for example CPUs and GPUs.¹⁰⁰²
- 5.18 Cloud providers depreciate their investment in a fixed asset over an asset's useful economic life. This means that the carrying value of an asset, as reflected in a cloud provider's annual financial statements, is less than the upfront cost of the asset.¹⁰⁰³ This also indicates that if a cloud provider had to sell its assets, for example if it exited the IaaS market, the recoverable amount would be less than the initial investment, as the value of the asset decreases over time.
- 5.19 The inability of a cloud provider to recover its initial investment on exit is likely to be further exacerbated by the highly specialised nature of the assets and the costs a cloud provider would incur disassembling and selling the assets.

Investment

- 5.20 We have reviewed the carrying value of the global fixed assets of AWS, Microsoft, Oracle Cloud and IBM Cloud as at 31 December 2024. The evidence shows that AWS and Microsoft have invested significantly more in global infrastructure than smaller cloud providers.¹⁰⁰⁴ For example:

¹⁰⁰¹ Responses to the CMA's information requests [X].

¹⁰⁰² These include central processing units (CPUs), graphics processing units (GPUs), dynamic random access memory (DRAM), Motherboards, solid state drives (SSDs) and hard disk drives (HDDs).

¹⁰⁰³ This is standard accounting practice.

¹⁰⁰⁴ See detail in appendix J.

- (a) AWS had global cloud fixed assets with a carrying value of [REDACTED] as at 31 December 2024;¹⁰⁰⁵
- (b) Microsoft had global cloud fixed assets with a carrying value of [REDACTED] as at 31 December 2024;¹⁰⁰⁶
- (c) Oracle had global cloud fixed assets with a carrying value of [REDACTED] as at 31 December 2024;¹⁰⁰⁷ and
- (d) IBM had global cloud fixed assets with a carrying value of [REDACTED] as at 31 December 2024.¹⁰⁰⁸

- 5.21 Google was unable to separate the assets of Google Cloud Platform from those of the larger Alphabet Group. Google also noted that it is unable to estimate the proportion of usage of Alphabet's assets attributable to Google Cloud or to GCP.
- 5.22 Our analysis of the carrying value of the global fixed assets of Alphabet Group shows that as at 31 December 2024 Alphabet had data centre, server and network assets had a carrying value of [REDACTED] billion, which is [REDACTED] to the carrying value of the global fixed assets of AWS and [REDACTED] Microsoft.¹⁰⁰⁹
- 5.23 Google also said that it manages its assets as a fleet and allocates assets to different businesses based on changing need.¹⁰¹⁰ This indicates that the types and capability of assets available to GCP are comparable to Microsoft and AWS.
- 5.24 Google's ability to allocate assets to its different businesses based on changing need is consistent with Microsoft who said on its 2025 Q2 earnings call that its long-lived infrastructure investments are fungible, enabling it to remain agile to meet customer demand globally across Microsoft Cloud including AI workloads.¹⁰¹¹ This shows that the three main cloud providers are able to benefit from the flexibility in how they allocate assets across their cloud and other businesses compared to smaller cloud providers who purchase assets to fulfil more specific needs and run a greater risk of asset under-utilisation and/or obsolescence.

¹⁰⁰⁵ AWS' response to the CMA's information request [REDACTED].

¹⁰⁰⁶ Microsoft's response to the CMA's information request [REDACTED].

¹⁰⁰⁷ Oracle's response to the CMA's information request under [REDACTED].

¹⁰⁰⁸ IBM's response to the CMA's information request under [REDACTED].

¹⁰⁰⁹ CMA analysis. See detail in appendix J.

Planned investments

- 5.25 Cloud providers and others such as private equity funds and sovereign wealth funds are planning to make further significant investments in infrastructure. These investments focus on infrastructure to support AI.
- 5.26 Public statements and internal documents from Microsoft and AWS indicate that they expect to increase their investments in global cloud fixed assets in the next couple of years, with the focus of the investment being on infrastructure to support AI capability:¹⁰¹²
- (a) Microsoft said that it plans to invest approximately \$80 billion to build out AI-enabled data centres to train AI models and deploy AI and cloud-based applications around the world in FY25 (ie the year to 30 June 2025).¹⁰¹³
 - (b) [REDACTED].¹⁰¹⁴
 - (c) More recently, AWS said that it expects to spend over \$100 billion globally on infrastructure in 2025, with the ‘vast majority’ of this spend going towards AI capability for AWS.¹⁰¹⁵
- 5.27 Google and Oracle are also planning to increase their investments in global cloud infrastructure in the next few years to support their AI capabilities.
- (a) Google is planning to invest approximately \$75 billion in global infrastructure in 2025 (primarily for servers, followed by data centres and network) to support its AI efforts across Google Cloud and all of its other businesses.¹⁰¹⁶
 - (b) An internal document from Google states that Google’s total Cloud Capex is forecast to be between [REDACTED] and [REDACTED] in each of years from 2025 to 2029, with “AI Machines Capex” accounting for between [REDACTED] and [REDACTED] each year.¹⁰¹⁷ The same internal document notes that [REDACTED]¹⁰¹⁸ [REDACTED].¹⁰¹⁹

¹⁰¹² [Amazon share pop on earnings beat, cloud growth](#), accessed on 6 November 2024; [Microsoft FY24 Third Quarter Earnings Conference Call](#), accessed on 6 November 2024; [REDACTED] response to the CMA’s information request [REDACTED].

¹⁰¹³ [The Golden Opportunity for American AI - Microsoft On the Issues](#); accessed on 8 April 2025.

¹⁰¹⁴ [REDACTED] response to the CMA’s information request [REDACTED].

¹⁰¹⁵ [Amazon doubles down on AI with a massive \\$100B spending plan for 2025 | TechCrunch](#); accessed on 8 April 2025.

¹⁰¹⁶ [GOOG Exhibit 99.1 Q4 2024](#); accessed on 8 April 2025.

¹⁰¹⁷ Google’s response to the CMA’s information request [REDACTED].

¹⁰¹⁸ Google’s response to the CMA’s information request [REDACTED].

¹⁰¹⁹ Google’s response to the CMA’s information request [REDACTED].

- (c) Oracle has announced several multi-billion-dollar investments in infrastructure in the UK,¹⁰²⁰ Malaysia,¹⁰²¹ and Spain,¹⁰²² to deliver AI and cloud computing services.

5.28 Private equity firms and sovereign wealth funds have also announced investments in the supply of data centres and energy infrastructure globally - key inputs for cloud services, for example:

- (a) BlackRock, Global Infrastructure Partners, Microsoft and the United Arab Emirates' MGX investment fund have announced a \$100 billion investment in new and expanded data centres and power infrastructure to deliver AI capabilities.¹⁰²³
- (b) Softbank Group, OpenAI, Oracle and the United Arab Emirates' MGX investment fund have collectively agreed to commit \$100 billion for the Stargate AI infrastructure.¹⁰²⁴
- (c) Blackstone said that it has over \$70 billion worth of data centre assets, with another \$100 billion in the pipeline, including facilities under construction.¹⁰²⁵ This includes a planned \$10 billion investment in a new AI data centre in Blyth, Northumberland.¹⁰²⁶ Blackstone links the growing demand for data centres and energy infrastructure to the power consumption and data intensity needed to deliver AI.¹⁰²⁷

5.29 The evidence cited above indicates that investments by private equity firms and sovereign wealth funds are focused on increasing the supply of data centres to support AI. We have not seen any evidence to suggest that private equity funds and sovereign wealth funds are making comparable levels of investment outside of infrastructure to support AI or that their investments represent competition for the three main cloud providers outside of AI.

5.30 In addition to their investments in data centre and server assets, Microsoft, AWS and Google have also made acquisitions to improve their network

¹⁰²⁰ [Oracle to Invest U.S. \\$5 Billion in Cloud Infrastructure in the UK | Oracle United Kingdom](#); accessed on 8 April 2025.

¹⁰²¹ [Oracle to Invest More Than US\\$6.5 Billion in AI and Cloud Computing in Malaysia | Oracle United Kingdom](#); accessed on 8 April 2025.

¹⁰²² [Oracle to Invest More Than \\$1 Billion in AI and Cloud Computing in Spain | Oracle United Kingdom](#); accessed on 8 April 2025.

¹⁰²³ [BlackRock, Inc. - BlackRock, Global Infrastructure Partners, Microsoft, and MGX Launch New AI Partnership to Invest in Data Centers and Supporting Power Infrastructure](#); accessed on 9 April 2025.

¹⁰²⁴ [SoftBank, OpenAI, Oracle and MGX commit to \\$100B for Stargate AI infrastructure | S&P Global](#); accessed on 8 April 2025.

¹⁰²⁵ [The Convergence of Data Centers and Power: A Generational Investment Opportunity | The Connection - Blackstone](#); accessed on 9 April 2025.

¹⁰²⁶ [Blackstone £10bn investment in Blyth AI data centre confirmed - BBC News](#); accessed on 23 April 2025.

¹⁰²⁷ [The Convergence of Data Centers and Power: A Generational Investment Opportunity | The Connection - Blackstone](#); accessed on 23 April 2025.

infrastructure.¹⁰²⁸ This is further evidence of significant capital investment, which could deter entry and act as a barrier to expansion by smaller providers.

Co-locating vs owning data centres

- 5.31 AWS and Google said that it is possible to enter the IaaS market without making the same levels of investment and that many cloud providers can start small and scale up as their business grows. They mentioned OVHcloud, eCloud VPC, Outscale, Scaleway and Brainboard as examples.¹⁰²⁹
- (a) AWS said that smaller cloud providers are able to compete despite the existence of scale advantages, noting that there is no need for a cloud provider to build its own global network of data centres and that infrastructure can be owned, leased or outsourced.¹⁰³⁰
- (b) Google said that small and mid-scale cloud providers can enter the IaaS market by using capital-efficient strategies such as leasing and co-location and scale out their capacity as their business grows.¹⁰³¹
- 5.32 An internal document from Microsoft said that it [REDACTED].¹⁰³² The same document also states that [REDACTED].¹⁰³³
- 5.33 In the UK, we have seen that cloud providers, including larger ones, have historically co-located or leased most, if not all, of their data centre capacity. However, recently Microsoft, AWS and Google have made significant commitments to build new data centres.
- (a) AWS said that most of its currently operating data centre sites in the UK are co-located.¹⁰³⁴ It recently committed to invest £8 billion between 2024 and 2028 to build and operate data centres in the UK.¹⁰³⁵
- (b) Microsoft builds and leases data centres in the UK and globally. In addition to its existing sites in the UK it has made public its plans to build new data centres in the UK.¹⁰³⁶

¹⁰²⁸ [AWS Acquires a Fiber Pair on MAREA Cable System on IRU Basis](#), accessed on 11 October 2024; [Why did Microsoft just buy a fiber optic cable company?](#) accessed on 11 October 2024; We also have some evidence of Google making investments in patents to improve its network infrastructure: [Google Acquires Fiber Optic Networking Patents](#), accessed on 11 October 2024.

¹⁰²⁹ Google's response to the CMA's information request [REDACTED]; AWS' response to the CMA's information request [REDACTED].

¹⁰³⁰ AWS' response to the CMA's information request [REDACTED].

¹⁰³¹ Google's response to the CMA's information request [REDACTED].

¹⁰³² Microsoft response to the CMA's information request [REDACTED].

¹⁰³³ Microsoft response to the CMA's information request [REDACTED].

¹⁰³⁴ AWS' response to the CMA's information request [REDACTED].

¹⁰³⁵ [AWS UK – Driving Digital Acceleration](#), accessed on 7 October 2024.

¹⁰³⁶ In 2023, Microsoft began construction a £1 billion data centre in North Acton (London) and in February 2024 Microsoft announced plans to build a data centre on the site of a former power plant in Eggborough, North Yorkshire [Eggborough datacentre project updates - Microsoft Local](#), accessed on 29 October 2024.

- (c) Google said that all of its data centres in the UK are currently owned and operated by third parties.¹⁰³⁷ It recently announced plans to invest \$1 billion to build its first UK data centre.¹⁰³⁸
- (d) Oracle said that it does not typically build or operate data centres but it leases them (globally).¹⁰³⁹

5.34 The evidence we have seen shows that new entrants (and existing providers) can use co-location to enter the IaaS market and to expand into new regions.

5.35 However, even if a new entrant enters the IaaS market, the scale of investment by Microsoft, AWS and to a lesser extent Google, means that any new entrant would not be able to compete directly with these providers in the IaaS market, as they would not have:

- (a) the geographic reach;
- (b) the same number of availability zones, which has implication for availability and resilience;
- (c) the network assets that allow for low latency, high security data transfers; or
- (d) the ability to shift demand between data centres to gain higher average rates of utilisation.

5.36 In addition, any new entrant would still:

- (a) incur upfront and irrecoverable costs when signing leases and entering into supply commitments; and
- (b) need to purchase the components to fit out and operate the data centre.

5.37 This shows that entry, even on a small scale, requires investment and involves considerable risk.

5.38 The combined impact of these factors is that, where any new entry does occur, it tends to focus on a 'niche' area in the IaaS market, rather than providing direct competition against the full-service offers of the largest cloud providers. This approach limits the competitive pressures exerted on these larger providers. For example, the companies that AWS noted in its submission to us as entering or obtaining investment in the IaaS market are firms offering GPU capacity to deliver Gen-AI (eg Voltage Park), clouds limited by geography (eg Cloudsky) or storage only services (eg Wasabi).

¹⁰³⁷ Google's response to the CMA's information request [🔗].

¹⁰³⁸ [Google starts building £790m site in Hertfordshire - BBC News](#), accessed on 29 October 2024.

¹⁰³⁹ Oracle's response to the CMA's information request [🔗].

Our conclusions on entry and expansion in the supply of IaaS

- 5.39 Entry and expansion in the supply of IaaS requires significant capital investment in fixed assets. These investments are mainly sunk costs that would not be recovered in full on exit.
- 5.40 The size of the investments made by the largest cloud providers are very large relative to those of smaller providers: the investments planned by each of Microsoft and AWS in the current financial year are over five times greater than the entire cloud-related fixed asset base of Oracle and IBM combined. In order to achieve a cloud asset base that would allow a rival to gain a competitively significant market share in IaaS provision, a rival provider would have to make even larger annual investments than Microsoft and AWS.
- 5.41 Microsoft and AWS have public plans to make the largest investments by cloud providers in cloud and AI related assets in the next few years, with Google announcing investments of a similar size. This will act to further widen the gap between Microsoft, AWS and Google and the smaller providers, which will raise the barriers to entry and expansion even higher.
- 5.42 Private equity firms and sovereign wealth funds have recently announced significant investments to deliver energy infrastructure and data centres. These investments are focused on increasing the supply of data centres to support AI. We have not seen any evidence to suggest that private equity funds and sovereign wealth funds are making significant investments outside of infrastructure to support AI.
- 5.43 While it is feasible in many markets for smaller rivals to constrain larger ones without matching their base of assets, the assets we have discussed in this section – the data centres, network assets, servers and components – are some of the core productive assets that allow cloud providers to operate in IaaS. This means that having access to a base of assets that is comparable to those of Microsoft and AWS is likely to be correlated with a firm's ability to compete with them, particularly in IaaS.
- 5.44 Large scale market entry will entail a risk that significant sunk investment costs may not be recovered, because it will generally be successful only if the entrant can expand the total market significantly, or win substantial business from an existing firm.¹⁰⁴⁰
- 5.45 On this basis, the size of Microsoft's, AWS' and to a lesser extent Google's investments and holdings of assets represent substantial barriers to entry for any

¹⁰⁴⁰ CC3 (Revised), paragraph 213.

provider that does not have the capital available to make investments of a similar magnitude.

- 5.46 The size of up-front investment can be reduced by co-locating or leasing data centres, allowing new entrants to scale their business as they grow their customer base.
- 5.47 However, entry on a small scale is unlikely to lead to a strong competitive constraint on the largest providers as they would not have the geographical reach, resilience, network or ability to shift demand compared to the largest providers. Also, data centres only make up part of the value of the three main cloud providers' fixed assets and a new entrant would still need to invest in the servers, components and network equipment for a co-located or leased data centre. Therefore, substantial levels of investment are still required.
- 5.48 The largest cloud providers are making significant further investment in their cloud infrastructure to meet growing demand for AI services. This increases the capital investment required by a new entrant, should they choose to offer customers accelerated compute capacity.
- 5.49 While there is some evidence of entry and expansion into niche areas within the IaaS market, this does not provide a comprehensive challenge across the full range of services offered by the largest cloud providers, and so any resulting competitive pressures would be limited.

Economies of scale

- 5.50 Economies of scale arise where average costs fall as the level of output rises over a range of output volume. As noted above, economies of scale, in combination with sunk investment costs, can constitute a barrier to entry or expansion.¹⁰⁴¹
- 5.51 In this section, we set out our assessment of whether the three main cloud providers benefit from economies of scale. The evidence that we have gathered shows that economies of scale are achieved through:
- (a) purchasing efficiencies, including bulk purchasing servers, components and network equipment;
 - (b) operating efficiencies, including in relation to energy requirements, data centre capacity and utilisation; and
 - (c) investment in research and development.

¹⁰⁴¹ CC3 (Revised), paragraph 212.

- 5.52 Economies of scale and, in particular, purchasing and operating efficiencies most clearly arise in relation to IaaS. However, they may also apply to PaaS in some circumstances, for example in relation to research and development.

Stakeholder views on economies of scale

- 5.53 Below we set out stakeholder views that we have received on economies of scale.

- (a) Microsoft submitted that significant capital investment is required to participate in the cloud market, and economies of scale matter.¹⁰⁴²
- (b) AWS submitted that some cost items, such as staff training and maintenance, scale with network size, but other costs such as security and management of the network increase more than proportionately to the network size as larger infrastructures are more exposed to outages and harder to monitor or manage. Cloud service providers with smaller networks incur these costs in a smaller measure, such that their overall unit cost can be lower than those of larger ones. This means that economies of scale are absent, or at least limited.¹⁰⁴³
- (c) A customer submitted that cloud services, and in particular the data storage and compute elements of the service stack, have many of the characteristics we see in natural or utility monopolies: early mover advantage, high fixed costs, economies of scale and high barriers to entry and expansion. It also said that because of the factors above, it is unsurprising that a small number of super-providers have come to dominate these markets.¹⁰⁴⁴

- 5.54 We also set out public statements made by stakeholders and responses to requests for information to cloud providers on economies of scale which are indicative of the scope for scale economies:

- (a) On its 2025 Q2 earnings call, Microsoft said that its 'data centres, networks, racks, and silicon are all coming together as a complete system to drive new efficiencies to power both the cloud workloads of today and the next-gen AI workloads'.¹⁰⁴⁵
- (b) An internal document from Microsoft identified [REDACTED].¹⁰⁴⁶
- (c) On its 2024 Q4 earnings call, Alphabet (Google's parent company) said that Google's 'leading infrastructure is also among the world's most efficient' and

¹⁰⁴² [Summary of hearing with Microsoft](#), 5 March 2025, paragraph 23.

¹⁰⁴³ [AWS' economic response to the CMA's competitive landscape working paper](#), dated 23 May 2024, paragraphs 92 and 93.

¹⁰⁴⁴ [REDACTED] submission to the CMA [REDACTED].

¹⁰⁴⁵ [Microsoft FY25 Second Quarter Earnings Conference Call](#); accessed on 9 April 2025.

¹⁰⁴⁶ Microsoft response to the CMA's information request [REDACTED].

that 'Google data centres deliver nearly four-times more computing power per unit of electricity compared to just five years ago'.¹⁰⁴⁷

- (d) On the same call, Alphabet also said that Google's data centres are 'industry-leading in terms of costs and power efficiency at scale'.¹⁰⁴⁸

Purchasing efficiencies

- 5.55 If larger cloud providers are able to achieve bigger purchasing discounts than smaller cloud providers, this would allow these larger ones to realise economies of scale by operating at a lower cost.
- 5.56 We asked cloud providers about their experience of purchasing servers, components and network equipment and they confirmed that there are higher discounts available for bulk purchasing:
- (a) AWS said that it (as any other business with large purchasing requirements) benefits from certain economies of scale/efficiencies derived from its global cloud infrastructure in relation to certain inputs, notably IT hardware (eg servers) with associated savings that are passed on to customers;¹⁰⁴⁹
 - (b) Microsoft said that purchasing hardware such as processors, servers, cooling units and network infrastructure in bulk allows it to achieve lower cost per MW of computing resources;¹⁰⁵⁰ and
 - (c) IBM said that procurement is shared across business units and divisions, allowing for purchase discounts and more efficient supply chain management.¹⁰⁵¹
- 5.57 This suggests that cloud providers can negotiate volume discounts on capital spending and that the volume discounts are likely to be proportionate to scale (ie larger providers will typically achieve lower purchasing prices).
- 5.58 Given that the cloud-related fixed asset bases of Microsoft and AWS are significantly larger than the smaller cloud providers and that Microsoft and AWS' cloud services business form part of much larger groups of companies (Amazon and Microsoft) that contain complementary businesses, AWS and Microsoft are likely to benefit most from volume discounts.

¹⁰⁴⁷ 2024 Q4 Earnings Call - Alphabet Investor Relations; accessed on 9 April 2025.

¹⁰⁴⁸ 2024 Q4 Earnings Call - Alphabet Investor Relations; accessed on 9 April 2025.

¹⁰⁴⁹ AWS' response to CMA's information request [38].

¹⁰⁵⁰ Microsoft's response to the CMA's information request [38].

¹⁰⁵¹ IBM's response to the CMA's information request [38].

- 5.59 Google cloud is also likely to derive benefit from volume discounts on capital spending arising from having large complementary businesses in the wider Alphabet Group.

Operating efficiencies

- 5.60 Below, we considered the extent to which larger cloud providers are able to achieve operating efficiencies, which lead to comparatively lower costs, by operating larger networks of data centres.
- 5.61 Microsoft has noted publicly that its cloud business benefits from three economies of scale that relate to operating efficiencies:
- (a) Data centres that deploy computational resources at significantly lower cost per unit than smaller ones.
 - (b) Data centres that coordinate and aggregate diverse customer, geographic and application demand patterns, improving the utilisation of computing, storage, and network resources.
 - (c) Multi-tenancy locations that lower application maintenance labour costs.¹⁰⁵²
- 5.62 We have broadly aligned our analysis with the first two operating efficiencies that Microsoft describes, as we consider these are where the largest cloud providers generate more economies of scale. We consider:
- (a) energy requirements;
 - (b) data centre capacities; and
 - (c) utilisation rates.

Energy requirements

- 5.63 We understand that energy is the largest variable cost incurred by a cloud provider when operating data centres, to such an extent that data centre capacity is typically measured in megawatts, ie the amount of energy that it consumes. Any efficiencies that larger cloud providers achieve in their energy consumption will give them an advantage compared to smaller cloud providers.

¹⁰⁵² [Microsoft 2023 Annual Report](#), accessed on 20 September 2024.

- 5.64 We have analysed the weighted average Power Usage Effectiveness (PUE)¹⁰⁵³ of Microsoft Azure, AWS, Google Cloud, Oracle Cloud and IBM Cloud. Our analysis shows that:¹⁰⁵⁴
- (a) all five cloud providers have average PUE ratings for their UK data centres that are lower (ie more effective) than industry average data centre PUE ratings;¹⁰⁵⁵ and
 - (b) Microsoft, AWS and IBM have [REDACTED].¹⁰⁵⁶
- 5.65 We also considered evidence on energy efficiency from cloud providers and data centre providers:
- (a) Google said that the supply of electricity comprises the largest portion of the operational costs of running a data centre, noting that energy supply is required for running and cooling the IT equipment. However, there are other day-to-day running costs involved including rent, maintenance, equipment, depreciation and labour costs.¹⁰⁵⁷
 - (b) Google also said that larger data centres will have greater energy overheads than smaller sites, but larger data centres also tend to be more energy efficient compared to smaller ones.¹⁰⁵⁸
 - (c) Microsoft said that large data centres achieve a lower cost per MW of computing resources through lower energy costs, through even distribution of workloads between servers and through optimised cooling solutions.¹⁰⁵⁹
 - (d) a global data centre provider, said that larger data centres achieve economies of scale through more efficient use of energy.¹⁰⁶⁰
 - (e) it also said access to energy can be a barrier to the provision of new data centre capacity, as uninterrupted and reliable sources of energy are not always available, particularly in locations with high demand due to a high density of networks.¹⁰⁶¹
- 5.66 Smaller cloud providers are able to benefit from some energy efficiencies by co-locating in larger data centres, for example, they may be able to use shared ventilation, air conditioning and cooling systems. However, smaller providers are

¹⁰⁵³ PUE is a calculation used by data centre providers to measure data centre efficiency by comparing the total facility energy to the IT equipment energy usage. [What is Power Usage Effectiveness \(PUE\)?](#) accessed on 20 September 2024.

¹⁰⁵⁴ See detail in appendix J.

¹⁰⁵⁵ See detail in appendix J.

¹⁰⁵⁶ [REDACTED].

¹⁰⁵⁷ Google response to the CMA's information request [REDACTED].

¹⁰⁵⁸ Google response to the CMA's information request [REDACTED].

¹⁰⁵⁹ Microsoft response to the CMA's information request [REDACTED].

¹⁰⁶⁰ Note of meeting with [REDACTED].

¹⁰⁶¹ Note of meeting with [REDACTED].

unlikely to benefit from other energy efficiencies to the same extent as larger providers. For example, when compared to smaller providers, Microsoft, AWS and to a lesser extent Google, have the ability to:

- (a) negotiate bigger volume discounts through their supply contracts;
- (b) shift demand across and within their large data centre networks to optimise energy efficiency across their networks; and
- (c) vertically integrate their energy supply chains – AWS, Microsoft and Google recently acquired stakes and/or entered into long-term supply commitments for nuclear energy to meet the power requirements for some of their data centres in the US.¹⁰⁶²

Data centre capacities

- 5.67 We have assessed global cloud data centre capacity in megawatts for the largest five cloud providers for the past few years and forecasts for future years.¹⁰⁶³ We have also considered the number and size of the data centres of six cloud providers: Microsoft, AWS, Google, Oracle, IBM and OVHcloud.¹⁰⁶⁴
- 5.68 Variation across providers in terms of their overall capacity, or in terms of the size of their data centres, may give rise to variation in the scope for firms to benefit from scale advantages such as energy efficiency, maintenance costs, ability to coordinate and aggregate demand across data centres. Physical proximity can often affect latency for customers, which might incentivise having more locations.¹⁰⁶⁵
- 5.69 Our analysis of capacity shows that [redacted] forecast the greatest absolute increases in capacity over the next few years. This is consistent with their forecast increases in investment in fixed assets, most notably investments to meet demand for accelerated compute services.¹⁰⁶⁶ [redacted].
- 5.70 There are differences in composition of the global data centre networks between the different cloud providers. However, in general, the three main cloud providers tend to have more larger data centres and smaller providers tend to have mainly smaller data centres.¹⁰⁶⁷

¹⁰⁶² [Amazon enlists nuclear small modular reactors in push for net carbon-zero](#), accessed on 29 October 2024; [Constellation to Launch Crane Clean Energy Center, Restoring Jobs and Carbon-Free Power to The Grid](#), accessed on 29 October 2024; [Google signs advanced nuclear clean energy agreement with Kairos Power](#), accessed on 29 October 2024.

¹⁰⁶³ See details in appendix J. Sources: Data from cloud providers.

¹⁰⁶⁴ See analysis in appendix J. Sources: Data from cloud providers.

¹⁰⁶⁵ [What to consider when selecting a region for your workloads](#), accessed on 7 October 2024.

¹⁰⁶⁶ Responses to the CMA's information requests [redacted].

¹⁰⁶⁷ See details in appendix J.

Utilisation

- 5.71 Cloud providers aim for high levels of utilisation of their data centres in order to maximise cost efficiencies.¹⁰⁶⁸
- 5.72 We gathered data from cloud providers to understand their levels of utilisation; however, not all the data that we obtained allowed for comparisons between companies.
- 5.73 We have analysed AWS' utilisation rates for its global EC2 (compute) physical server utilisation. Our analysis shows that AWS' utilisation rates have increased over the past six years (the period for which we have data) from an average utilisation of [%] in 2018 to an average utilisation of [%] in 2023.¹⁰⁶⁹ AWS' utilisation rate for 2023 of [%] is [%] than that of IBM, who noted that its current average level of both daily and monthly utilisation is approximately [%] and that it managed its capacity so as not to exceed [%] utilisation.¹⁰⁷⁰
- 5.74 AWS' increasing utilisation rates, [%], are consistent with AWS being better able to shift demand between its data centres and achieve greater operating efficiencies when compared to smaller rivals.
- 5.75 We have also analysed the utilisation rates that cloud providers consider to be 'high',¹⁰⁷¹ and found that:
- (a) provider views on what constitutes high rates of utilisation varies quite widely; and
 - (b) the utilisation rates that AWS and Microsoft consider to be high are higher than the other cloud providers.¹⁰⁷²
- 5.76 The evidence that we have received indicates that [%] utilisation rates compared to smaller providers and that the high utilisation rates that [%].¹⁰⁷³

Investment in research and development

- 5.77 Large incumbent firms may benefit from significant economies of scale in the innovative process. This is because large-scale firms that undertake large amounts of research and development (R&D) may be able to employ more specialised resources; they will face smaller average total costs because they can

¹⁰⁶⁸ [%] response to the CMA's information request [%].

¹⁰⁶⁹ We only have data for 10 months in 2023 (January to October). The average utilisation rates by year are 2018: [%], 2019: [%], 2020: [%], 2021: [%], 2022: [%], 2023 (10 months): [%] AWS' response to the CMA's information request [%].

¹⁰⁷⁰ IBM's response to the CMA's information request [%]

¹⁰⁷¹ See details in appendix J.

¹⁰⁷² See details in appendix J, Sources: Data from cloud providers.

¹⁰⁷³ See details in appendix J. Microsoft's response to the CMA's information request [%].

average the fixed costs of their innovative effort over a greater level of output; and they may be able to support a larger portfolio of R&D efforts, increasing the likelihood that this will develop an improved product or process likely to be applicable to at least one of its businesses.¹⁰⁷⁴

- 5.78 Spend on R&D can improve quality and increase efficiency. Customers stand to benefit from better quality, more innovative cloud services and/or from lower prices (assuming any benefits and/or efficiencies that are realised are passed on, eg through lower pricing). However, when spend on R&D is significant, it can be difficult for smaller firms to achieve and so could contribute to barriers to entry and expansion.
- 5.79 We have considered the following evidence on cloud providers' spend on R&D:
- (a) Microsoft said that it spent \$27 billion on R&D in financial year 2023 and that other cloud providers made similar investments in research and development. It also said that high levels of R&D spend are required to attract new customers and limit the number of customers switching to rivals.¹⁰⁷⁵
 - (b) IBM also said that it incurred significant annual expenditure on R&D exceeding \$[REDACTED] billion across its global technology businesses and that investment results in innovation benefitting products both in and out of cloud.¹⁰⁷⁶
 - (c) Google said that its internal R&D entails a degree of fixed and sunk costs. It also said that it is not necessary to make significant R&D investments or incur costs to begin to offer basic services (compute, storage and network).¹⁰⁷⁷
 - (d) AWS said that it works with its partners to custom-develop server hardware that it deploys into its data centres. This approach helps it to reduce the time to market for new innovative products and enables cost reductions for these inputs.¹⁰⁷⁸ [REDACTED].¹⁰⁷⁹
- 5.80 In summary, the size of the investments made by the largest cloud providers in R&D are very large in absolute terms and we have some evidence that they contribute to cost reductions for inputs. In order to compete directly with the largest cloud providers, a rival would also need to consider significant investments in R&D to achieve similar outcomes. The levels of investment in R&D needed would

¹⁰⁷⁴ CC3 (Revised), paragraph 183.

¹⁰⁷⁵ Microsoft's response to the CMA's information request [REDACTED].

¹⁰⁷⁶ IBM's response to the CMA's information request [REDACTED].

¹⁰⁷⁷ Google's response to the CMA's information request [REDACTED].

¹⁰⁷⁸ AWS' response to the CMA's information request [REDACTED].

¹⁰⁷⁹ AWS' response to the CMA's information request [REDACTED]. [REDACTED].

reduce if the rival was competing in a niche area with a more limited product portfolio.

Our conclusions on economies of scale in the provision of cloud services

- 5.81 There are material economies of scale in the provision of cloud services, in particular for IaaS. These include that larger cloud providers generally:
- (a) achieved purchasing efficiencies from bulk discounts on necessary equipment;
 - (b) achieved lower operating costs through a combination of a greater number and larger average size of their data centres, which generally results in lower average energy requirements and higher utilisation rates; and
 - (c) made material investments in R&D which they can then spread over a wider business.
- 5.82 The significance of the individual factors varies, but in aggregate and when viewed in conjunction with the significant sunk cost investments required to enter and expand in the IaaS market, we found that there are material economies of scale in the provision of cloud services that act as barriers to entry and expansion.

Size of cloud providers' product portfolio

- 5.83 We have considered whether having a large portfolio of cloud services across IaaS and PaaS may give a cloud provider advantages over its rivals.
- 5.84 Cloud providers typically offer a range of first-party cloud services through their platforms and some also offer third party services provided by ISVs through a marketplace accessible by customers on the provider's platform. Marketplaces are used by eligible ISVs to offer their own services to the customers of those providers.
- 5.85 In our assessment of cloud providers' product portfolios, we looked at both first- and third party services, as well as IaaS and PaaS.
- 5.86 We broke down product portfolios by discussing:
- (a) the importance of range of services;
 - (b) economies of scope; and
 - (c) network effects.

Importance of range of services

- 5.87 If customers value having access to a large portfolio of cloud services, then providers with a large portfolio may be more attractive to customers relative to smaller providers, such that smaller providers may find it harder to compete for customers.
- 5.88 We gathered customer views on the importance of range of first- and third-party products.
- 5.89 We asked large customers to rate the importance of a list of factors their organisation considers when choosing their main public cloud. In answering this question, customers rated the importance of the range of cloud infrastructure services alongside other factors.
- 5.90 Range of cloud infrastructure services was identified as one of the most important factors by customers we contacted when choosing their main public cloud provider. This was alongside other factors such as price (including discounts or cloud credits), service quality and data sovereignty requirements and number and location of data centres.¹⁰⁸⁰
- 5.91 Customers rating the range of services offered as either important or very important gave various reasons for this:
- (a) Some customers said having a broad range of services was important in allowing them to build their own end products effectively and with flexibility, and meet their business needs.¹⁰⁸¹
 - (b) Some other customers said that a provider offering a large range of services could be a good signal that the provider is introducing new and improved services.¹⁰⁸²
- 5.92 Some customers said that the range of cloud infrastructure services is less important because most providers offer the same services.¹⁰⁸³
- 5.93 The Jigsaw report shows that one of the key reasons why Microsoft, AWS and Google are the main providers customers would consider is that they offer a wide range of solutions that cover many needs. The Jigsaw report also shows that other cloud providers, such as IBM and Oracle, are perceived as being for secondary use or only for certain use cases.¹⁰⁸⁴

¹⁰⁸⁰ CMA analysis of responses to the CMA's information requests [§].

¹⁰⁸¹ Responses to the CMA's information requests [§].

¹⁰⁸² Responses to the CMA's information requests [§].

¹⁰⁸³ Responses to the CMA's information requests [§].

¹⁰⁸⁴ [Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw \(2024\)](#), paragraphs 3.4.4 and 3.4.8, page 27.

Our assessment of the importance of range of services

- 5.94 We found that the range of services is an important factor for customers when choosing which cloud provider(s) to use as their main cloud provider.

Economies of scope

- 5.95 Economies of scope arise when producing two (or more) services is less costly for a single firm than for two (or more) firms each to produce the services separately. Where economies of scope are significant, an entrant, if it is to be successful, might have to produce a range of services from the outset, adding to the costs of entry.¹⁰⁸⁵ Economies of scope might be relevant if, for example, R&D and operations spend can be spread over a wider range of services.
- 5.96 We asked cloud providers to explain whether they could benefit and have in the past benefitted from any efficiencies as a result of increasing their range of services.
- 5.97 Cloud providers said that these efficiencies could arise and some also said they have benefitted from them. The views of cloud providers varied as to the strength of these efficiencies. Some also said that expanding the range of services can sometimes lead to inefficiencies.
- 5.98 In relation to the efficiencies that could arise:
- (a) Microsoft said that it has benefitted from efficiencies through expanding its number of cloud offerings. It said that these efficiencies have benefitted its business both in the development and operation of services. On the development side, it said that it has been able to build new services using existing expertise and technology across different layers of the cloud stack. On the operations side, it said that it has been able to optimise the utilisation and allocation of its resources.¹⁰⁸⁶
 - (b) Oracle said that cloud providers benefited from increasing their range of services. It said that cloud providers also use their own services, just like customers do. Therefore, cloud providers also directly benefit from new services and features. It said that any new service likely builds upon the fundamental services of compute, storage, network, identity, and security. It also said that it has benefited from increasing the range of services.¹⁰⁸⁷
 - (c) AWS said that an increase in the number of services offered does not necessarily increase overall efficiency for a cloud provider.¹⁰⁸⁸ It also said it

¹⁰⁸⁵ CC3 (Revised), paragraph 214.

¹⁰⁸⁶ Microsoft's response to the CMA's information request [X].

¹⁰⁸⁷ Oracle's response to the CMA's information request [X].

¹⁰⁸⁸ AWS' response to the CMA's information request [X].

benefitted from simplifying operations and sharing learnings and technologies across teams where there are opportunities to do so that would benefit customers. It said that some of its products share commonalities across codebases, so if one of the service teams discovered a vulnerability or bug in the code, they would quickly share the learnings with the other team to remedy the issue for customers. It also said that its product teams can use the same services for commonly used functionality like logging, billing, and authentication, meaning they did not need to build those tools from the ground up to launch a new service and could pass on benefits to customers in the form of a consistent user experience.¹⁰⁸⁹

- (d) IBM said that cloud providers could benefit from efficiencies as they increased the number of customer-facing cloud services by developing processes and systems to support those services. It said that having centralised teams that build and manage base runtimes, pipelines, automation and architecture allowed the service teams to focus on the higher-order functions / customer-facing aspects of their service. In general, centralisation of resources allows cloud providers to better scale and drive down costs.¹⁰⁹⁰
- (e) Google said that any efficiencies that may be derived from developing a broader range of services were, in its experience, limited (albeit there may be exceptions).¹⁰⁹¹

5.99 In relation to the inefficiencies that could arise:

- (a) AWS said that many of the services it offers are discrete and distinct, and the learnings and functionalities that applied to one may not necessarily translate or be useful to others. It also said that increasing the number of service offerings introduced inefficiencies due to increased complexity. It said that customers generally expected that its services will work together and offer a consistent experience. Therefore, it said, as the number of its service offerings grows, its service teams must design, maintain, and adapt their services to satisfy an increasing number of parameters tied to its other services.¹⁰⁹²
- (b) Microsoft said that having a large service portfolio required a broad and deep set of operational resources and expertise to manage, debug, and secure a cloud comprised of so many different offerings. On the other hand, it said,

¹⁰⁸⁹ AWS' response to the CMA's information request [REDACTED].

¹⁰⁹⁰ IBM's response to the CMA's information request [REDACTED].

¹⁰⁹¹ Google's response to the CMA's information request [REDACTED].

¹⁰⁹² AWS' response to the CMA's information request [REDACTED].

entities that focused on just one or a limited set of cloud solutions also had their own efficiencies that can be generated by such specialisation.¹⁰⁹³

Our assessment

- 5.100 There could be some economies of scope in supplying a range of cloud services, but this would not be the case for all cloud providers and for all cloud services. In some cases, increasing the portfolio of cloud services might instead lead to inefficiencies.

Network effects

- 5.101 If customers value having access to a large portfolio of first and third-party services (via ISVs) through the cloud infrastructure providers they use, providers with a large pool of ISVs being hosted on their platform and/or listed on its marketplace would have an advantage over other providers. This could result in providers without a large pool of ISVs finding it harder to compete for customers.
- 5.102 This might be compounded if there are indirect network effects between cloud infrastructure providers, ISVs and customers. Indirect network effects may arise as follows:
- (a) The more customers that a cloud infrastructure provider has, the more attractive it becomes to ISVs because it provides them with access to a larger customer base and the more likely ISVs are to use that cloud provider's platform. ISVs may run their services on the cloud infrastructure of the provider and also list their services on its marketplace.
 - (b) The more ISVs available on a cloud infrastructure provider's platform, the more attractive the provider becomes to customers (eg because they can access more ISV services on that platform) and the more likely customers are to use that provider's platform.
- 5.103 To assess the strength of any network effects, we collected evidence from ISVs and customers.

Evidence from ISVs

- 5.104 We have found that ISVs are more attracted to cloud providers with more customers:
- (a) One ISV said that a cloud provider having a large user base is quite important. It explained that given limited resources, putting resources into

¹⁰⁹³ Microsoft's response to the CMA's information request [38].

preparing to work with or on a public cloud has to have sales opportunities downstream attached to it. It further explained that one way of helping guarantee future revenue is to make sure there is a large user base on that cloud and that many of the European clouds have sufficient user bases to be of interest to it.¹⁰⁹⁴

- (b) Another ISV said that its customers typically want to use its services on the cloud provider and in region(s) where their data resides. It explained that, because of this, the size of the cloud provider's customer base is relevant when it decides which public cloud providers to support.¹⁰⁹⁵
- (c) Similarly, some ISVs highlighted the size of user base of the different cloud providers, eg as proxied by their market shares, as an important consideration when deciding which cloud providers to run on and/or which to run on first.¹⁰⁹⁶
- (d) Some ISVs also noted that there are significant costs associated with running their services to additional public cloud providers beyond the largest ones.¹⁰⁹⁷

Evidence from customers

- 5.105 To assess whether customers value cloud providers with more ISVs, we asked large customers about range of services offered by ISVs as a factor in their choice of main public cloud. We also asked about the range of cloud services.¹⁰⁹⁸
- 5.106 Range of cloud infrastructure services was identified as one of the most important factors by customers when choosing their main public cloud provider. This was alongside other factors such as price (including discounts or cloud credits), service quality and data sovereignty requirements and number and location of data centres.
- 5.107 Range of services offered by ISVs was identified as one of the least important factors relative to the other selection criteria. Examples of reasons given for the lack of importance of this factor were:
- (a) some customers said that most ISV services were accessible via all cloud providers and therefore were not very important when deciding which main cloud provider to use;¹⁰⁹⁹ and

¹⁰⁹⁴ Note of meeting with [X].

¹⁰⁹⁵ Note of meeting with [X].

¹⁰⁹⁶ [X] response to Ofcom information request [X] Ofcom note of meeting with Confluent [X].

¹⁰⁹⁷ Responses to Ofcom's information requests [X].

¹⁰⁹⁸ There may be some overlap between these two criteria if customers do not always distinguish between services they receive from their main public cloud provider and those they receive from ISVs on that public cloud.

¹⁰⁹⁹ Responses to the CMA's information requests [X].

- (b) other customers said that the range of ISVs is not a significant differentiator when choosing a main provider and first party services are a more important factor.¹¹⁰⁰

5.108 A few customers said that the range of cloud infrastructure services offered by ISVs was becoming more important and could be important to their medium and long-term strategy.¹¹⁰¹

Our assessment

5.109 Range of services is important to customers, and ISVs appear to value a broad customer base.

5.110 Currently, customers perceive that equivalent ISV services are widely available, although this is mostly across the main cloud providers. Therefore, having a wide range of ISV services is not acting as a competitive differentiator, at least between the main cloud providers. However, any network effects arising from the range of cloud services may contribute to the barriers faced by smaller providers.

Other potential barriers to entry and expansion

Cloud credits

5.111 Cloud credits are a form of discounting by cloud providers. They work as follows:

- (a) Cloud providers offer customer credits free of charge as an incentive to switch clouds or use their cloud services more.
- (b) Customers can redeem cloud credits against spend on cloud services.
- (c) The amount of credits that a cloud provider offers can vary between customers, though providers tend to set limits on the amount of credits that they will offer by type of customer eg a start-up, an SME or an AI start-up.
- (d) Customers are aware of the amount of cloud credits they will receive in advance of using a cloud service as they are agreed upfront.
- (e) Cloud credits are limited by time and/or by amount. If they are limited by both time and by amount and customers do not use them within the timeframe offered, they lapse.

5.112 Cloud credits are offered by a range of cloud providers:

¹¹⁰⁰ Responses to the CMA's information requests [X].

¹¹⁰¹ Responses to the CMA's information requests [X].

- (a) AWS offers start-ups up to \$100,000 in AWS activate credits and up to an additional \$300,000 in AI specific credits.¹¹⁰²
- (b) Microsoft offers start-ups up to \$150,000 in Azure credits.¹¹⁰³
- (c) Google offers start-ups up to \$200,000 in cloud credits and AI start-ups up to \$350,000 in cloud credits.¹¹⁰⁴
- (d) Oracle gives start-ups the ability to purchase discounted credits to scale.¹¹⁰⁵
- (e) OVHcloud offers start-ups up to €100,000 in cloud credits.¹¹⁰⁶
- (f) some smaller cloud providers, such as Civo, also offer cloud credits.¹¹⁰⁷

5.113 As described above, many of the schemes from cloud providers are targeted at start-ups, although cloud providers also offer cloud credits to other types of customers.

Stakeholder views on cloud credits

5.114 Some smaller cloud providers, former UK Cloud employees and an industry body have raised concerns about the impact of cloud credits on customers' incentives and the ability of smaller cloud providers to compete.¹¹⁰⁸

5.115 Civo said that:

- (a) the three main cloud providers have the 'deepest pockets' which allows them to offer significant cloud credits to customers, which act as a loss leader.¹¹⁰⁹
- (b) cloud credits can be the deciding factor for customers when choosing a cloud provider. This is particularly relevant for start-ups and SMEs, as they are more sensitive to reducing their costs. Start-ups and SMEs tend to be the 'target market' for the smaller cloud providers.¹¹¹⁰
- (c) a large cloud provider offered a smaller cloud provider's customer three years of free compute to transfer its business and a different large cloud

¹¹⁰² [Get AWS Activate Credits - AWS Startups](#), accessed on 7 October 2024.

¹¹⁰³ [Unlocking Azure credits as your startup grows](#), accessed on 7 October 2024.

¹¹⁰⁴ [Google for Startups Cloud Program | Google Cloud](#), accessed on 7 October 2024.

¹¹⁰⁵ [Free cloud is just the beginning](#), accessed on 7 October 2024.

¹¹⁰⁶ [Startup Program](#), accessed on 7 October 2024.

¹¹⁰⁷ [Grow with the Civo Startup Program](#), accessed on 7 October 2024.

¹¹⁰⁸ Note of meeting with [redacted] submission to the CMA [redacted]. [Former UKCloud employees response to issues statement](#). We also received a submission from the Open Cloud Coalition, who also considered credits offered by the largest cloud providers to be anti-competitive and to act as a barrier to entry and expansion. [Open Cloud Coalition submission to the CMA dated 18 December 2024](#), page 5. Open Cloud Coalition members include Google Cloud.

¹¹⁰⁹ Note of meeting with [redacted].

¹¹¹⁰ Note of meeting with [redacted].

provider offered a customer of another small cloud provider \$500k to transfer its business.¹¹¹¹

- (d) smaller providers are unable to match the cloud credits offered by the large cloud providers and this means that smaller providers cannot compete with the large providers for start-up and SME customers.¹¹¹²
- (e) customers are unlikely to migrate back to the smaller providers after they have utilised the cloud credits offered by large providers, as they are 'locked-in' to the three main cloud providers and it is costly and time consuming for customers to switch.¹¹¹³
- (f) excessive cloud credits align with the concept of predatory pricing, creating barriers to entry where ultimately, customers that have become locked in will face "take it or leave" pricing with little opportunity to negotiate.¹¹¹⁴

5.116 Another small cloud provider said:

- (a) Microsoft and AWS offer \$150,000 and \$100,000 in cloud credits to start-ups and this effectively means that smaller providers can no longer compete with the larger providers for this business.¹¹¹⁵

5.117 Blackbox hosting said:

- (a) Microsoft and AWS have 'extremely deep pockets', making the cost of offering a free credit programme insignificant – whereas smaller private cloud companies cannot afford to offer free credits.¹¹¹⁶
- (b) AWS partners with venture capital funds, who then require all their portfolio companies to switch to AWS in exchange for credits and discounts. This approach effectively locks other cloud providers out of the market for start-ups.¹¹¹⁷
- (c) free credits and egress fees together are preventing and distorting competition, locking out smaller private cloud providers. Free credits worsen the anti-competitive effect of egress fees.¹¹¹⁸

5.118 OVHcloud said:

¹¹¹¹ Note of meeting with [REDACTED].

¹¹¹² Note of meeting with [REDACTED].

¹¹¹³ Note of meeting with [REDACTED].

¹¹¹⁴ [Civo's response to our provisional decision](#) dated 17 February 2025.

¹¹¹⁵ [REDACTED] submission to the CMA [REDACTED].

¹¹¹⁶ [REDACTED] submission to the CMA [REDACTED].

¹¹¹⁷ [REDACTED] submission to the CMA [REDACTED].

¹¹¹⁸ [REDACTED] submission to the CMA [REDACTED].

- (a) cloud credits present a risk to reinforce the market's concentration around the three main cloud providers as they play a significant role in attracting early-stage startups to the services of one provider, and the startup will face strong barriers to switch, or multi-cloud.¹¹¹⁹
- (b) the three main cloud providers are the only providers able to provide cloud credits in massive amounts and to numerous companies thanks to their superior financial means. This dynamic is now being extended to AI startups and this creates a serious risk of the AI ecosystem is being steered towards the services of certain cloud providers due to them offering cloud credits and not because they are the best fit for customer needs.¹¹²⁰

5.119 The former UK Cloud employees said:¹¹²¹

- (a) the largest providers' primary objective is to bring in new business, especially start-ups and they offer inducements such as cloud credits to achieve this.
- (b) the amount of cloud credits that the largest providers are willing to offer has increased. AWS now offers \$500,000 of cloud credits to start-ups wanting to build generative AI products and it is common for similar amounts of cloud credits to be offered to enterprise customers.
- (c) smaller cloud providers cannot compete with these levels of discounts and are put at a disadvantage and this disadvantage will deepen and grow, as the smaller cloud providers cannot build their revenues to fund the infrastructure needed to compete.
- (d) start-ups may be grateful for the cloud credits, but face lock-in and dependence on the largest providers.

5.120 The Open Cloud Coalition said:

- (a) the primary issue is how cloud credits function as an acquisition tool to attract early-stage businesses, before transitioning them into committed spend agreements that make switching prohibitively expensive.
- (b) while start-ups may initially benefit from free credits, the reality is that these credits serve as a hook to draw customers into an ecosystem. Once embedded – through proprietary tools, APIs, and preferential integrations – switching becomes difficult, especially when combined with other discounting mechanisms like committed spend agreements.

¹¹¹⁹ OVHcloud response to provisional findings report dated 21 February 2025, paragraph 11.

¹¹²⁰ OVHcloud response to provisional findings report dated 21 February 2025, paragraph 11.

¹¹²¹ Former UKCloud employees response to issues statement.

- (c) cloud credits primarily target start-ups and scale-ups – the future high-growth businesses of the UK economy. By capturing them early, dominant providers entrench their market position for the long term.
- (d) while the Jigsaw report suggests that start-ups plan to switch providers, this does not equate to genuine competition. In practice, start-ups move between dominant providers when incentivised by successive rounds of cloud credits.¹¹²²

5.121 The Startup Coalition and customers in the Jigsaw report highlighted the benefits of cloud credits, with an emphasis on the benefits for start-up customers.

5.122 The Startup Coalition,¹¹²³ also said that start-ups are very protective of any scheme or incentive that enables them to access cloud services more cheaply and that start-ups would have major concerns with any attempt to limit free credits regimes.¹¹²⁴

5.123 The Jigsaw report set out customer views on cloud credits:

- (a) credits were primarily mentioned by customers in the start-up phase of their lifecycle. These credits can be defined either in financial terms (eg \$100,000 credit) or as a fixed time period (eg free for first two years).¹¹²⁵
- (b) start-up businesses often make use of multiple cloud providers in their early years to take advantage of credits from each, as this keeps cost low, allows experimentation, and means they can draw on a range of provider input and advice.¹¹²⁶
- (c) the barrier to starting to use the public cloud is low. Participants mentioned start-up and migration credits, pay-as-you-go services and easy basic set-up for many public cloud providers. These were especially important for start-ups who cannot afford the initial capital expenditure needed for their own data centre or the staff to build their own bespoke IT services.¹¹²⁷
- (d) some start-ups used services offered by certain cloud providers because they were free or heavily discounted via start-up credits, even though they did not intend to continue to use these suppliers once credits ended.¹¹²⁸
- (e) participants from start-up businesses described the cloud providers as quite active in trying to win their business. One business spoke about using one

¹¹²² Open Cloud Coalition's response to the provisional decision, page 3.

¹¹²³ Google is a supporter of the Startup Coalition and offers cloud credits. Source: <https://startupcoalition.io/join-us/> accessed on 7 October 2024.

¹¹²⁴ Startup coalition's response to our Issues Statement: [Startup Coalition response to the CMA's Issues Statement](#)

¹¹²⁵ Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 6.1.17.

¹¹²⁶ Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 3.2.5.

¹¹²⁷ Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 3.3.3.

¹¹²⁸ Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 3.5.5.

cloud provider for 6 months (IBM) and then switching to another (Google) because the latter provider offered them \$300,000 in free credits over two years if they moved everything to Google.¹¹²⁹

- (f) in addition to credits for start-ups, there is also evidence that providers offer credits for migrating more of a company's on-premises workloads into the cloud. These credits are offered to mature organisations who have not yet fully migrated everything onto the cloud and can be offered in addition to a committed spend agreement.¹¹³⁰ Cloud providers offer cloud credits to assist enterprise customers migrating to their cloud.¹¹³¹

Our assessment

- 5.124 Cloud credits are one of many discounting strategies used by both large and small cloud providers to gain and retain customers.
- 5.125 Unlike other discounting strategies used by cloud providers, such as reserved instances and savings plans, where a customer commits to a consistent amount of usage,¹¹³² and committed spend agreements, where a customer commits to a pre-specified level of spend,¹¹³³ customers are not required to make a commitment in exchange for receiving cloud credits. The absence of commitment means that a customer is not obliged to use the cloud provider who gave them the cloud credits once those credits have expired.
- 5.126 The value of cloud credits is typically smaller than other forms of discount offered by cloud providers¹¹³⁴ and cloud credits are targeted at start-up customers in the hope/expectation that start-up customers will stay and grow their business with the cloud provider once the cloud credits have been utilised or have expired.
- 5.127 Customers, and specifically start-ups, view cloud credits positively. The Jigsaw report notes that start-ups benefit from cloud credits by minimising costs, experimenting in the cloud, and drawing on a range of cloud provider input and advice.¹¹³⁵
- 5.128 The Jigsaw report also found that a customer had utilised their cloud credits from one provider before moving their business to another cloud provider who also offered cloud credits, thereby allowing them to benefit from further discounts.¹¹³⁶

¹¹²⁹ [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), paragraph 6.1.18.

¹¹³⁰ [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), paragraph 6.1.19.

¹¹³¹ [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), paragraph 6.1.19.

¹¹³² [Compute Savings Plans – Amazon Web Services](#); accessed on 14 May 2025.

¹¹³³ See chapter on CSAs for more detail.

¹¹³⁴ Ofcom found that cloud credits are ten to 15 times smaller than committed spend discounts and approximately four times smaller than reserved instances and saving plans. Reserved instances and savings plans are volume discounts offered to customers in return for a customer committing to spend a specified amount on a cloud service (eg virtual machines) over a stated period (eg 1 year). [Cloud services market study final report](#), paragraph 5.1.86.

¹¹³⁵ [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), paragraph 3.2.5.

¹¹³⁶ [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), paragraph 6.1.18.

¹¹³⁷ Similarly, the Jigsaw report noted that some start-up customers planned to migrate away from their current cloud provider once they had used their cloud credits.¹¹³⁸ This anecdotal evidence is also consistent with our analysis of prevalence of switching that indicates that smaller customers are the most likely to switch, such that any lock-in effect would be lower for start-ups who benefit from cloud credits.¹¹³⁹

5.129 We recognise that cloud credits could have a negative impact on smaller cloud providers that are unable to match the significant amount of cloud credits offered by the three main cloud providers and that this could impact their ability to compete with them, particularly for start-up customers.

5.130 However, given that cloud credits are offered by a range of cloud providers, customers are not obliged to use the cloud provider who gave them the cloud credits once those credits have expired, and credits are considered positively by customers, particularly start-ups, we do not consider that cloud credits constitute a feature leading to harm to competition in the provision of cloud services in the UK.

5.131 Although we have not identified competition concerns, this does not mean that harm to competition might not materialise in the future. In particular, we consider that cloud credits could harm competition if:

- (a) the value of cloud credits offered by the three main cloud providers increases significantly; and/or
- (b) the three main cloud providers require a commitment in exchange for cloud credits.

5.132 To the extent that other barriers to switching and multi-cloud such as egress fees are factors that some smaller cloud providers consider are inhibiting their ability to compete for customers once the cloud credits have expired, we have taken these concerns into account in our analysis of those features.

Reputational barriers

5.133 One small cloud provider said that there are reputational barriers to entry and expansion as customers perceive the three main cloud providers to be superior and are seen as a safe choice for a Chief Information Officer when choosing a cloud provider.¹¹⁴⁰

¹¹³⁷ The Open Cloud Coalition also acknowledged that some start-up customers switch providers once they have utilised their credits with their first provider to access more credits offered by another provider.

¹¹³⁸ [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), paragraph 3.5.5.

¹¹³⁹ Our analysis on the prevalence of switching as set out in the section Market Outcomes.

¹¹⁴⁰ Note of meeting with [X].

- 5.134 Some customers also commented that they consider the three main cloud providers to be more credible and capable than smaller cloud providers.¹¹⁴¹
- 5.135 The Jigsaw report stated: ‘The main providers are seen as Microsoft, AWS and Google among participants. For some, this is the main or only consideration set in terms of who might even be on a shortlist of providers in the event of a review of the market or a switch, though most were not aiming to make any changes. They each have an excellent reputation, are seen to deliver a reliable service and offer a wide range of solutions that cover many needs’.¹¹⁴²
- 5.136 It also said that AWS was commonly seen by customers as a ‘first mover’. There was a sense that other public cloud providers were catching up but there was still a reputational advantage.¹¹⁴³

Our assessment

- 5.137 Our view is that there are some reputational barriers to entry and expansion in the cloud services markets but it is unclear how important reputational barriers are relative to other barriers to entry and expansion.

Regulatory barriers

- 5.138 In this section we consider barriers to entry and expansion that could arise from regulations and from public sector procurement.
- 5.139 AWS, Google and IBM have all identified a number of regulations they have to comply with in order to run their operations including data security (for example the Network and Information Systems Regulations) and data privacy (for example UK General Data Protection Regulation).¹¹⁴⁴
- 5.140 However, no cloud provider has flagged these as being particular barriers to entry or expansion. For example, Google said it does not believe that there are significant legal or regulatory barriers which restrict new entrants from developing and offering services.¹¹⁴⁵
- 5.141 Generally, cloud providers were also unable to provide accurately the costs of compliance and how this might differ globally and in the UK.

¹¹⁴¹ Notes of calls with [REDACTED]. [REDACTED] response to CMA’s information request [REDACTED].

¹¹⁴² [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), paragraph 1.3.6.

¹¹⁴³ [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), paragraph 3.4.10.

¹¹⁴⁴ AWS’ response to the CMA’s information request [REDACTED]; Google’s response to the CMA’s information request [REDACTED]; IBM’s response to the CMA’s information request [REDACTED].

¹¹⁴⁵ Google’s response to the CMA’s information request [REDACTED].

- 5.142 Cloud providers¹¹⁴⁶ commented on their overall regulatory burden, including European Union legislation (most notably the Data Act),¹¹⁴⁷ and the impact this had on their business. But they focussed on the broader consequences of this for their operations, for example diverting resource from delivery and innovation as opposed to representing barriers to entry or expansion.
- 5.143 In order to serve financial services customers, cloud providers in the UK have some specific regulatory obligations:
- (a) Financial services firms themselves have certain regulations and these place duties on any cloud providers that provide products or services to financial services firms.¹¹⁴⁸
 - (b) Providers recognised the tailored offerings they were required to offer financial services firms.¹¹⁴⁹
- 5.144 The UK regulatory landscape is changing as sectors and regulators respond to the increased adoption and importance of cloud services across the economy. For example:
- (a) The UK government has recently classed UK data centres as Critical National Infrastructure.¹¹⁵⁰
 - (b) In financial services, the Bank of England, Prudential Regulation Authority and Financial Conduct Authority (referred to as ‘the Regulators’) have published requirements, designed to strengthen the resilience of Critical Third Party (CTP) Services provided to the financial sector.¹¹⁵¹ HM Treasury has the powers to decide whether a third-party is designated as critical, based on whether a failure, or disruption to, its services could threaten the stability of, or confidence in, the UK financial system.¹¹⁵²
 - (c) The Regulators’ cost benefit analysis recognises that where a CTP’s services to the financial sector is already highly concentrated, the requirements could further entrench the market power of third parties designated as critical. However, they assess that the risk of this is small when considering the mitigations put in place by the Regulators, namely: rules prohibiting CTP’s

¹¹⁴⁶ Responses to the CMA’s information requests [§].

¹¹⁴⁷ [Data Act - Shaping Europe’s digital future](#), accessed on 21 March 2024.

¹¹⁴⁸ For example, [SS2/21 Outsourcing and third party risk management - Bank of England](#) in the UK and the [Digital Operational Resilience Act \(DORA\) - European Union](#) in the EU, both accessed on 02 April 2024.

¹¹⁴⁹ Responses to the CMA’s information requests issued [§].

¹¹⁵⁰ [Data centres to be given massive boost and protections from cyber criminals and IT blackouts](#), accessed on 7 October 2024

¹¹⁵¹ [PS24/16: Operational resilience: Critical third parties to the UK financial sector - FCA](#), accessed on 19 November 2024.

¹¹⁵² Nothing in this report should be read as implying or signalling that cloud service providers will be designated as CTPs. As noted above, this is ultimately a decision for HMT. However, the regulators’ cost-benefit analysis in their final CTP policy is of relevance and use to any potential future CTPs whose services to the financial sector are highly concentrated.

using designation as a mark of regulatory approval; and proportionate requirements, such that they are unlikely to result in an increase in costs that will significantly impact the structure of a third-party market.¹¹⁵³

- (d) The Information Commissioner's Office oversees compliance with UK data protection law,¹¹⁵⁴ and will consult on updated draft guidance on cloud computing in early 2026.¹¹⁵⁵

Our assessment

- 5.145 While cloud providers may face some costs associated with customers' regulatory compliance, in particular in sectors such as financial services, we have not identified any material regulatory barriers to entry and expansion in cloud services in the UK.

Public sector procurement

- 5.146 We have considered how public sector customers purchase cloud services in order to assess whether this customer segment has any characteristics that affect how competition works in cloud services by raising barriers to entry and expansion for new providers and smaller providers who are seeking to expand.¹¹⁵⁶
- 5.147 We have received submissions and evidence raising issues around public sector procurement policy and implementation.¹¹⁵⁷
- 5.148 The public sector is an important customer group for cloud providers and we estimate that it may represent 5% of UK IaaS and PaaS revenues.¹¹⁵⁸ Public sector customers are likely to grow their use of cloud services with the majority of workloads still to be migrated from on-premises and with UK government policy encouraging further migration of IT services to the cloud.¹¹⁵⁹
- 5.149 UK government procurement policy recognises the risk of raising barriers to entry and expansion potentially leading to market concentration and customer lock-in with regard to cloud services. It has a number of policies which seek to address

¹¹⁵³ CP26/23 - Operational resilience: Critical third parties to the UK financial sector - Bank of England, accessed on 14 February 2024.

¹¹⁵⁴ Cost benefit analysis, accessed on 7 November 2024.

¹¹⁵⁵ Cloud computing guidance for organisations, accessed on 22 October 2024.

¹¹⁵⁶ Our plans for new and updated guidance - ICO.

¹¹⁵⁷ More detail is included in appendix K.

¹¹⁵⁸ Former UKCloud employees', Prolinx's, OVHcloud's and Oracle's responses to the Issues Statement dated 17 October 2023. Available at: [Cloud services market investigation](#); [redacted] response to the CMA's information request [redacted]. Open Cloud Coalition submission to the CMA dated 18 December 2024. Open Cloud Coalition members include Google Cloud. [Open Cloud Coalition's response to the CMA's PDR dated 28 January 2025](#); [Former employee of UKCloud's response to the CMA's PDR dated 28 January 2025](#); [Prolinx's response to the CMA's PDR dated 28 January 2025](#); [Civo's response to the CMA's PDR dated 28 January 2025](#).

¹¹⁵⁹ [redacted] submission to the CMA [redacted].

¹¹⁵⁹ Note of meeting with [redacted].

this.¹¹⁶⁰ In particular, its latest policy Cloud Compute 2 requires competitive tendering.¹¹⁶¹ In fact, most cloud providers we contacted said public sector customers already tender competitively.¹¹⁶²

- 5.150 We have seen some inconsistency between the central policy and individual strategies. In particular, we found two instances of policies appearing to specify the use of ‘hyperscalers’.¹¹⁶³
- 5.151 AWS and Microsoft appear to be the largest providers to the public sector. Their share of supply to public sector customers appears to be at least the same as their overall position in the cloud services markets and their joint share may be even higher in the public sector.
- 5.152 Their leading positions amongst public sector customers are likely driven by the barriers to entry and expansion, and barriers to multi-cloud and switching we have identified elsewhere in this report as giving rise to their unilateral market power, not through public procurement practices, and therefore are addressed in our assessment of those features.
- 5.153 There have been a number of recent developments in government policy, including the state of digital government review¹¹⁶⁴ and work on the Digital Commercial Centre of Excellence.¹¹⁶⁵ As a result of these recent developments, and our provisional findings, [REDACTED].¹¹⁶⁶
- 5.154 We suggest that UK government should continue to improve data collection on public sector procurement of cloud services so that outcomes can be better assessed and the relative positions of cloud providers better understood.
- 5.155 We also suggest that the UK government should promote and share best practices, and address inconsistencies in public bodies’ cloud procurement strategies where necessary as public sector use of cloud services continues to grow.

Conclusions on barriers to entry and expansion

- 5.156 There are substantial barriers to entry and expansion in the provision of cloud services, in particular for IaaS.
- 5.157 Market entry and expansion in the supply of IaaS requires significant capital investment in fixed assets, which for many asset types would be largely

¹¹⁶⁰ One Government Cloud Strategy and Cloud First policy.

¹¹⁶¹ Note of meeting with [REDACTED].

¹¹⁶² Responses to Ofcom’s information requests [REDACTED].

¹¹⁶³ [NHS Cloud Strategy adoption plan - NHS England, Digital Cloud Strategic Roadmap for Defence.](#)

¹¹⁶⁴ [State of digital government review - GOV.UK](#), accessed on 10 April 2025

¹¹⁶⁵ [A blueprint for modern digital government \(HTML\) - GOV.UK](#), accessed on 10 April 2025.

¹¹⁶⁶ Note of meeting with [REDACTED].

irrecoverable upon exit. This combines with economies of scale, whereby the larger providers have comparatively lower average costs.

- 5.158 Unless a new entrant (or company seeking to expand) is willing to make investments of a similar magnitude to those of the largest suppliers, it is likely to face higher average costs to provide an equivalent level of service and so may struggle to compete effectively across a sufficiently wide range of cloud services. This is likely to be disincentivising IaaS market entry and expansion.
- 5.159 Furthermore, given the scale of investment and expansion that the three main cloud providers have made in IaaS, any new entrant (or company seeking to expand) would need to invest substantially more than the large existing providers in order to build its market share. The levels of investment that Microsoft, AWS and Google are expecting to make in IaaS and in AI infrastructure in the coming years may raise these barriers even higher.
- 5.160 A range of investment firms have recently announced significant investments in energy infrastructure and data centres. These investments are focused on increasing the supply of data centres to support AI.
- 5.161 Investment in infrastructure is a requirement for entry and expansion into IaaS, but not PaaS. However, competition in PaaS using third party IaaS likely involves sourcing IaaS from firms that are also competing in the supply of PaaS. This limits the scope for PaaS providers to compete effectively.
- 5.162 While investment by cloud providers may have pro-competitive effects and benefit customers, this does not preclude them also having the effect of deterring market entry.
- 5.163 The wide product portfolios of the larger cloud providers also contribute to the barriers to entry and expansion in both IaaS and PaaS markets because range of services is an important consideration for customers selecting a cloud supplier. ISVs also value access to the wider customer bases of the larger cloud providers and prioritise developing services compatible with cloud services offered by these cloud providers. In addition, customers place some importance on a provider's reputation and this contributes to overall barriers to entry and expansion in cloud services.
- 5.164 Some stakeholders raised concerns about the impact of cloud credits on customers' incentives and the ability of smaller cloud providers to compete. We recognise that cloud credits could have a negative impact on smaller cloud providers that are unable to match the credits offered by the three main cloud providers. We have assessed how cloud credits work and we do not consider that they are a feature which is harming competition in cloud services at present, although this harm could materialise in future if their value grew or scope changed.

- 5.165 We have considered whether procurement of cloud services by public sector customers affects competition in cloud services markets. Microsoft and AWS appear to be the largest providers to the public sector and this is consistent with their overall position in cloud services markets. Public sector procurement policy aims to maintain competition in the sector, including by requiring competitive tendering of contracts, and greater competition in cloud services would create greater choice for public sector customers.
- 5.166 We suggest that UK government should continue to monitor the outcomes of the outcomes of public procurement and drive best practice in the application of procurement frameworks.

6. Barriers to switching and multi-cloud

- Customers face both technical and commercial barriers when seeking to multi-cloud or switch their provider, and many think that the costs of switching or using multi-cloud outweigh the benefits.
- Providers' incentives to facilitate switching and multi-cloud are mixed, as there may be some incentives to win more business, but there are also disincentives if it becomes easier for customers to switch away.
- Cloud services have developed in different ways, with different functionalities and features. Accordingly, whether or not intentional, there are technical barriers that stem from the differentiation of features and interfaces in core and ancillary cloud services. This means that customers cannot easily compare, substitute, or integrate services from different providers. Further disincentives include higher latency, a lack of transferable skills across clouds and insufficient transparency on how to mitigate or overcome technical barriers.
- There are some ways for customers to partially mitigate the impact of technical barriers in specific use cases, although the patchwork of mitigations does not amount to a comprehensive solution to the barriers identified. Customers recognised some positive aspects of existing technical mitigations, but also thought that cloud providers could do more to improve them or offer additional mitigations.
- In terms of commercial barriers, egress fees are incurred when a customer transfers data out of a cloud provider's cloud. Egress fees are not typically the most important factor in customer decisions relating to switching and multi-cloud but some customers identify them as a challenge, and in some instances, a disincentive when considering switching and multi-cloud. The significance of egress fees as a barrier varies among customers, with the highest fees as a proportion of total cloud spend more likely to be paid by smaller customers and customers egressing larger amounts of stored data.
- Microsoft, AWS and Google have introduced free switching programmes, but we found that these have substantial limitations and restrictions, including the exclusion of ongoing multi-cloud egress fees from their scope. The introduction of these programmes does not impact our analysis of whether egress fees act as a barrier to switching and/or multi-cloud use.
- We have not seen sufficient evidence that investments and innovations cited by cloud providers as being funded by egress fees are dependent on or are solely driven by these fees, nor that they result in passed-on cost savings to customers or deter customers from inefficient egress usage.
- Technical barriers and egress fees are both barriers to switching and/or multi-cloud. Together, these barriers restrict customers from exercising effective choices about their

cloud providers and affect customers' ability to either switch provider or use multiple providers.

Introduction

- 6.1 This chapter presents our assessment of barriers affecting the ease with which customers can switch cloud provider and/or use multiple clouds.
- 6.2 In order to drive effective competition in a market, customers need to be both willing and able to access information about the various offers available, assess these offers to identify the good or service that provides the best value for them, and act on this assessment by switching to purchasing the good or service from their preferred supplier.¹¹⁶⁷ Customers' ability to either switch providers or use multiple providers (ie multi-cloud) therefore represents an important source of competitive pressure on providers.
- 6.3 In some cases, barriers to switching suppliers or barriers to using multiple suppliers may restrict customers from exercising effective choice. Switching from one provider to another, so as to respond to attractive offers, may be made difficult for customers by the costs of doing so.¹¹⁶⁸
- 6.4 When customers face significant impediments to switching or using multiple providers, sellers may be able to set prices or levels of quality, range and functionality with only limited regard to competition and, in particular, less regard to competition than they would in markets with lesser impediments to switching or using multiple providers.
- 6.5 Firms can enjoy some market power if customers cannot easily or effectively switch, because of the difficulty or cost of switching to better deals. If, for example, one provider raises its price for a particular service above the level of other providers (or fails to cut its price when rivals cut theirs), many of the provider's customers may switch to rivals if they are able to do so at little cost. In contrast, if the cost of switching or the cost of using multiple providers is higher than the benefits the customers stand to gain by switching away some or all of their demand, the provider may be able to raise or maintain its higher price without incurring lost sales – ie the provider has some degree of market power.
- 6.6 Providers may be able to charge higher prices (or offer lower quality, innovation or range of functionality) than they otherwise would even where some customers do switch or use multiple providers. This would be the case when the costs of

¹¹⁶⁷ CC3 (Revised), Guidelines for market investigations: Their role, procedures, assessment and remedies, paragraph 296.

¹¹⁶⁸ CC3 (Revised), paragraphs 297 and 316.

switching or multi-cloud are too high for some customers or for certain workloads to be moved to other providers.

- 6.7 We have looked closely at the costs of switching and multi-cloud because they directly affect customers' ability and incentives to choose the provider(s) that best fits their needs. As we set out elsewhere in this report, the supply of cloud services is concentrated, levels of switching are very low and multi-cloud is generally observed only amongst large or very large customers; there are also substantial barriers to entry in the provision of cloud services. These factors contribute to market power being enjoyed by a small number of providers.¹¹⁶⁹ In this context, customers' ability to switch and multi-cloud is particularly important.
- 6.8 The CMA may consider a wide range of potential barriers to switching or using multiple providers, including inconvenience, administrative obstacles, the presence of network effects locking customers into existing standards and large one-off costs or investments required to switch.¹¹⁷⁰
- 6.9 We have focused on technical and commercial barriers to switching and multi-cloud, in particular:
- (a) technical barriers related to the way cloud infrastructure services are designed, including differentiation of features and differentiation of interfaces;
 - (b) other technical barriers, including latency,¹¹⁷¹ skills and transparency; and
 - (c) egress fees related to switching and/or multi-cloud, ie fees customers have to pay to transfer data out of their provider's cloud and into another.
- 6.10 Before assessing these barriers, we considered the providers' submissions on their overall incentives to facilitate switching and multi-cloud.

Providers' incentives to facilitate switching and multi-cloud

Cloud provider views

- 6.11 Microsoft, AWS and Google all told us that they have incentives to enable switching and/or customers' use of multi-cloud.¹¹⁷²
- (a) AWS said that it and other cloud providers have, and will continue to have, strong commercial incentives to reduce technical limitations across differentiated systems and support interoperability due to the customers'

¹¹⁶⁹ CC3 (Revised), paragraph 189.

¹¹⁷⁰ CC3 (Revised), paragraph 318.

¹¹⁷¹ Latency measures the time it takes to transfer data between public clouds.

¹¹⁷² Providers also made more specific submissions in relation to their incentives to mitigate technical barriers to switch and multi-cloud. We discuss them in the section below on Technical barriers.

consideration of interoperability with existing IT solutions when moving to the cloud. AWS said IT services providers compete on a workload-by-workload basis and will not win these workloads if it does not offer sufficient interoperability.¹¹⁷³

- (b) it said there can be no presumed relationship between a firm's size and its incentives to support or degrade interoperability. By way of example, AWS said that Oracle's latest versions of its database software are restricted to run on its proprietary Oracle Cloud Infrastructure.¹¹⁷⁴ AWS said that its incentives are different to Microsoft and Oracle, whose software is internal to those companies (ie built and maintained by those companies), and who have an interest in maintaining control of their software through restrictive licence and support terms.¹¹⁷⁵
- (c) AWS said that the PDR acknowledges multi-cloud is 'not uncommon' amongst large customers, and that large customers are key drivers of competition, which it said supports the conclusion that CSPs are incentivised to support multi-cloud.¹¹⁷⁶
- (d) AWS said that it is committed to interoperability, and gave examples of how it offers Kubernetes services, supports open source database software (eg PostgreSQL and MariaDB), and File Server.¹¹⁷⁷ It also said that it works hard to ensure compatibility with Linux, and that its IAM services are compatible with widely used third-party technologies such as SAML and OIDC certificates, which it implemented as a result of customer demand.¹¹⁷⁸ AWS also said that it responds to customer demand, as seen with its introduction of vector database capabilities (following start-ups like Pinecone and the open-source community that supports PostgreSQL) by building and launching such capabilities to RDS, among other database services.¹¹⁷⁹
- (e) similarly, AWS said that it makes the Amazon S3 APIs publicly available under open-source licences and publishes documentation on them, so customers and competitors can build applications that use Amazon S3 or create software that is compatible, and that S3 compatible solutions are used extensively throughout industry.¹¹⁸⁰
- (f) AWS also said that at its inception, it had 100% share of supply (on the narrowest market definition) and did not adopt a strategy of reducing

¹¹⁷³ AWS' response to the CMA's PDR dated 28 January 2025, paragraphs 23 and 39; AWS' response to the CMA's updated issues statement and working papers, paragraphs 13 and 15.

¹¹⁷⁴ AWS' response to the CMA's PDR dated 18 February 2025, paragraph 39a.

¹¹⁷⁵ AWS' submission to the CMA [REDACTED].

¹¹⁷⁶ AWS' response to the CMA's PDR dated 18 February 2025, paragraph 39. [REDACTED].

¹¹⁷⁷ AWS' submission to the CMA [REDACTED].

¹¹⁷⁸ AWS' submission to the CMA [REDACTED].

¹¹⁷⁹ AWS' submission to the CMA [REDACTED].

¹¹⁸⁰ Summary of hearing with Amazon Web Services paragraph 27. AWS' submission to the CMA [REDACTED].

interoperability, and that if customers were locked in, there would be evidence of AWS or other providers taking advantage of that, but there is not. AWS said that this shows that AWS is not incentivised to reduce interoperability.¹¹⁸¹

- (g) AWS also submitted that generative AI and demand from FM developers will continue to incentivise IT providers to support interoperability. It said that FM developers are sophisticated buyers of IT services that want the best-in-class inputs from different IT providers and, to the extent necessary, the ability to move between them. To win FM-related workloads, providers will need to ensure their services work together with the services of other providers that a customer might choose to run their workloads. AWS said FM developers have control and flexibility over how they build their models, which may include offerings from more than one provider. As a result, IT providers will be incentivised to support interoperability and portability.¹¹⁸²
- (h) AWS said that it is misguided to assert that, in future, when a lot of the migration from on premises to cloud has occurred, the incentives to interoperate may weaken or change. It said that it is wrong to assume that switching from on-premises to cloud is inherently easier than switching between clouds and that switching from on-premises to the cloud is permanent.¹¹⁸³
- (i) Microsoft submitted that its position as a challenger to AWS in the cloud services market means that it has always been incentivised to make it as easy as possible for customers to switch to Microsoft (in particular, from AWS) or to multi-cloud as customers focus on diversifying beyond AWS.¹¹⁸⁴
- (j) Microsoft said that competition for new workloads is not drying up and therefore cloud providers' incentives to interoperate are not reducing. Instead, Microsoft said that the source of new workloads is changing. This is because the customer application lifecycle continues after migration to the cloud, including changes to make the application cloud-native and to take advantage of advanced cloud capabilities (eg integration of AI functionalities).¹¹⁸⁵
- (k) Google said that it agrees that challenger cloud providers are incentivised to promote switching and multi-cloud.¹¹⁸⁶

¹¹⁸¹ AWS' submission to the CMA [§].

¹¹⁸² AWS response to CMA's information request [§]. [AWS' response to the CMA's PDR](#), paragraph 24.

¹¹⁸³ [AWS' response to the CMA's Technical barriers working paper dated 06 June 2024](#), paragraph 75.

¹¹⁸⁴ [Microsoft's response to the Issues Statement](#), 17 October 2023, paragraphs 26.

¹¹⁸⁵ [Microsoft's response to the Technical barriers working paper dated 06 June 2024](#), paragraphs 21-25.

¹¹⁸⁶ [Google's response to the CMA's PDR dated 28 January 2025](#), paragraph 37. Google's response to CMA's information request [§].

- 6.12 We have received submissions from some providers on the incentives of AWS and Microsoft specifically to support features which improve portability and interoperability:
- (a) Google said that players with market power, like Microsoft, may in some cases have different incentives to other smaller providers, which can impact how they approach technical design decisions for their products and services. It said that Microsoft and AWS have significant market power in cloud and Microsoft in particular has significant market power in on-premises software which includes Active Directory, as part of Windows Server. Google said Microsoft is therefore not only incentivised to design its products in a way to maintain its market power in cloud, but also to maintain and extend the market power it holds on-premises and across its software ecosystem, including Office 365.¹¹⁸⁷
 - (b) [X] said that not all cloud providers have the same incentive to ensure that their services are portable and interoperable. The CMA should distinguish situations where (i) a provider has customers currently operating on two clouds, without underlying strategy, from (ii) cases where a provider promotes an intentional multi-cloud architecture. While in both cases the cloud provider may provide tools to facilitate integration, the incentive is different. The former is primarily driven by the desire to ensure customer satisfaction in order to retain and potentially increase the customer spend, but not an incentive to promote interoperability.¹¹⁸⁸

Our assessment

- 6.13 When cloud providers consider whether to reduce barriers to switching and multi-cloud, they face potential benefits and costs, as set out in our guidelines.¹¹⁸⁹ Providers may have an incentive to reduce these barriers if they are seen by customers to be relatively easy to switch to and from or multi-cloud with, and this will lead customers to be more willing to choose them as a provider in order to limit any long-term lock-in effects.
- 6.14 However, when cloud providers reduce the barriers to switching and multi-cloud, customers' ability to switch and multi-cloud increases, and so does their bargaining power. This increases the competitive pressure on providers to reduce prices or incur costs to raise quality.¹¹⁹⁰

¹¹⁸⁷ [Google's response to the Technical barriers working paper dated 06 June 2024](#), paragraph 13; [Google's response to the Egress fees working paper dated 23 May 2024](#), paragraph 57.

¹¹⁸⁸ [X] submission to the CMA [X].

¹¹⁸⁹ [CC3 \(Revised\)](#), paragraph 317.

¹¹⁹⁰ [CC3 \(Revised\)](#), paragraph 296.

- 6.15 The size of these benefits and costs to cloud providers may be affected by various factors:
- (a) all else being equal, the benefit of reducing barriers to switching and multi-cloud to win customers may be more important to smaller providers as they have more to gain and less to lose due to their relatively smaller customer base. In contrast, the costs of reducing switching barriers for existing customers may be more significant for larger providers.
 - (b) smaller providers are less likely to be a customer's main cloud supplier than larger providers and are more likely to therefore rely on customers being willing to multi-cloud, and this will tend to contribute to the incentive to interoperate with other clouds.
 - (c) technologies that facilitate migration from traditional IT may also be used to facilitate switching or multi-cloud between public cloud providers. For example, a cloud provider said that container portability could lead to more customers operating in multiple clouds and that it would 'double down' on containers in spite of the risks because it will quicken the migration of traditional workloads towards it.¹¹⁹¹ This adds complexity to the incentives of the providers:
 - (i) on one hand, providers may be incentivised to develop technologies that facilitate switching or multi-cloud if they also accelerate the rate of migration from other IT models, therefore allowing them to grow the market.
 - (ii) on the other hand, these technologies might not be relevant if they only facilitate the lift-and-shift process of migration from other IT models and do not facilitate the switching of modernised workloads developed on the cloud. We explore this further below in relation to technical barriers.
- 6.16 Microsoft and AWS have submitted examples of where they have responded to specific customer requests for features related to interoperability and portability.¹¹⁹²
- 6.17 However, we observe that a proportion of such customer requests was not granted and therefore it appears that cloud providers weigh up their incentives to implement such features, as well as other factors on a case-by-case basis. These other factors include the required technical effort, the existence of sufficient workarounds, and the magnitudes of customer demand and industry trends.¹¹⁹³

¹¹⁹¹ [redacted] response to Ofcom's information request [redacted].

¹¹⁹² AWS' response to the CMA's information request [redacted]. Microsoft's response to the CMA's information request [redacted].

¹¹⁹³ For example, see responses to the CMA's information requests [redacted].

- 6.18 AWS also submitted that there are various rational reasons why such requests might not be granted, for example if customers request an existing capability without realising they have not configured their workload in line with best practices. Therefore, the request may not be granted because the customer is provided with an explanation on how they can use the capability in question.¹¹⁹⁴
- 6.19 Internal documents indicate that Microsoft, AWS and Google consider several factors when setting egress fees, including competitors' pricing and the pricing of their other networking services. This indicated that there are strategic considerations in the setting of egress fee prices that go beyond customers' preferences for switching and multi-cloud.
- 6.20 We therefore consider that this evidence does not necessarily demonstrate a broader incentive for Microsoft and AWS to facilitate interoperability and portability or reduce any barrier to switching and/or multi-cloud.
- 6.21 The submissions above also show that the incentive to meet customer needs has resulted in the development of features that positively impact interoperability and/or portability in particular instances. We discuss the details of these features in more detail in the Mitigations section of this chapter.
- 6.22 However, in relation to the incentive to meet customer needs by providing such features, we note that:
- (a) In general, making it easier for customers to switch may not confer a competitive advantage on a cloud provider or result in increased sales if rivals replicate these changes. Providers may recognise this, which, in conjunction with the fact that their efforts to reduce switching costs may increase customers' bargaining power and reduce their profitability, would reduce their incentive to proactively reduce switching costs.
 - (b) Similarly, a customer's incentive to switch when its cloud provider does not provide sufficient interoperability or portability is reduced when the alternatives also provide similarly limited levels of interoperability or portability, and where the customer has a wide range of factors to consider when choosing a provider, potentially including other factors where there is greater differentiation between providers.
- 6.23 The responses from cloud providers show that there are positive incentives, in theory, to facilitate interoperability and portability. But as shown above, there are also disincentives.
- 6.24 Overall, we consider that cloud providers' incentives to facilitate switching and multi-cloud are not clear-cut and are at best mixed, and that this applies both to

¹¹⁹⁴ AWS' submission to the CMA [X].

workloads migrating to the cloud from other IT models and to cloud-native workloads.

- 6.25 We have therefore assessed the extent to which switching barriers exist and the significance of those switching barriers on customers' behaviour and competition.
- 6.26 The rest of the chapter considers potential barriers to switching and multi-cloud related to:
- (a) technical barriers; and
 - (b) egress fees.


Technical barriers

Introduction

- 6.27 This section presents our analysis of the impact of technical barriers on public cloud customers' ability and incentive to switch between clouds and/or use multiple clouds.
- 6.28 There are a number of sources of technical barriers and they can be grouped into those that relate to the way cloud services are designed and those that do not. Over time, new sources of technical barriers may emerge or become apparent.
- 6.29 Design-related technical barriers include:
- (a) differentiation of features: differences in the functionality of similar cloud services hosted on different public clouds; and
 - (b) differentiation of interfaces: differences in the interfaces (eg protocols or APIs) of similar cloud services hosted on different public clouds.
- 6.30 Other technical barriers include:
- (a) latency: the time it takes to transfer data between public clouds. This is a relevant factor when considering integrating between multiple public clouds, but also when customers need to move data across regions and/or availability zones;
 - (b) skills: the difference in technical skills needed to work with different public clouds; and
 - (c) transparency: the availability and discoverability of information about potential technical challenges and how they can be mitigated or overcome.

- 6.31 When customers seek to switch between clouds or simultaneously use clouds from different providers, they will typically incur additional costs (including the additional expense, time, use of human and other resources, or engineering effort) in an effort to overcome or mitigate one or more of these barriers.
- 6.32 Broadly, customers may incur two types of additional costs when trying to use multiple clouds for their workloads:
- (a) Operational costs: customers may incur additional costs when operating workloads on more than one public cloud, independent of the level of integration between the workloads. For example, a customer may expend additional engineering effort to align its use of ancillary services and tools across clouds, or to reconcile different billing systems.
 - (b) Integration costs: customers can incur additional costs when enabling workloads on multiple public clouds to communicate. This includes setting up connections over a network using service APIs, as well as any ongoing management and operation of any integrations.
- 6.33 Customers also incur various costs when switching workloads that exist on one public cloud (origin cloud) to another (target cloud):
- (a) redesigning the workload such that it can be run on the target cloud;
 - (b) setting up and operationalising services on the target cloud;
 - (c) moving relevant data from the origin cloud to the target cloud; and
 - (d) testing the new workload on the target cloud before switching it off on the origin cloud. This step may also incur temporary multi-cloud costs.
- 6.34 Customers will face a different mixture of these costs depending upon whether they are placing a new workload on a new cloud or switching an existing workload, and depending upon the level of integration between this workload and other workloads in the cloud. This is set out in the figure below.

Figure 6.1: The costs associated with different customer activities



Source: CMA

6.35 We set out below our conceptual framework for assessing the impact of technical barriers on customers' ability to switch and multi-cloud and our analysis of:

- (a) the existence of technical barriers to switching and multi-cloud and the impact of these barriers on customer behaviour;
- (b) the technical barriers to switching and multi-cloud in relation to core services,¹¹⁹⁵ ancillary services and tools;¹¹⁹⁶
- (c) other technical factors; and

¹¹⁹⁵ We have distinguished between what we have named 'core services' and 'ancillary services and tools'. Core services are the main IaaS and PaaS services that contribute to the key objectives of the customer's workload. Ancillary services and tools provide functions that support the operation and management of core services.

¹¹⁹⁶ We use the term 'ancillary services' to refer to managed ancillary services and, in this context, 'tools' to refer to ancillary software that customers deploy and manage on their own.

- (d) technical mitigations that may reduce the technical barriers to switching and multi-cloud.

Conceptual framework

- 6.36 We have considered whether, and to what extent, technical barriers prevent or restrict the ability of customers to switch and use multi-cloud, particularly integrated multi-cloud architecture.
- 6.37 We recognise that cloud customers are heterogeneous with different needs and architectures, and that means that some technical barriers may impact the behaviour of some customers more than others. For example:
 - (a) The largest enterprises may have more complex cloud architectures that are more difficult to port to another cloud, but they may also have more in-house expertise or ability to hire external consultants for mitigating any technical barriers in doing that. By contrast, smaller customers may have simpler cloud architectures that are easier to switch, but also less expertise and fewer resources.
 - (b) Customers who have been using the cloud for longer may have developed more practical experience that they can use when mitigating technical barriers, but they may also have integrated more deeply into their hosting cloud(s) such that it is more difficult to switch and multi-cloud. By contrast, customers who have just migrated to the cloud might find it more difficult to understand the technical requirements for switching and multi-cloud, even while their cloud architectures may be less entrenched.
 - (c) Customers who migrated their workloads from on-premises may also experience technical barriers differently to customers who built their workloads on the cloud in the first place, due to factors such as whether they use legacy software and the number of custom integrations they may need to recreate during a switch.
- 6.38 There are also a range of factors affecting the level of switching and multi-cloud costs, which relate to the types of workload that customers have on public clouds and therefore tend to be quite customer-specific. These can include: the number of applications in a workload; the architecture and IT stack fit (similarity between origin and target cloud, including compatibility of services); data sources, types and volumes; data sensitivity and compliance requirements (ie higher compliance requirements increase complexity); continuity requirements (ie sensitivity of workload to interruption); and availability of required skills and resources.¹¹⁹⁷

¹¹⁹⁷ Responses to the CMA's information requests [X].

- 6.39 As a result of the heterogeneity described above, the particular technical barriers experienced by customers may vary significantly from one customer to the next. A particular technical barrier to one customer's operations may not be perceived as a barrier at all by another customer because of the different cloud services that it uses or the way in which it uses them.¹¹⁹⁸ Therefore, we have sought views from customers both on the general technical ease of switching and multi-cloud and on particular barriers.
- 6.40 Even if technical barriers only affect a proportion of customers, or are only perceived by them when considering switching or using a multi-cloud strategy,¹¹⁹⁹ they could still have a significant effect on competition overall, especially when considered cumulatively. Therefore, we have also considered the cumulative impact of several potential costs or points of friction that may arise for customers.
- 6.41 Cloud providers told us that technical barriers are inherent or intrinsic, arising naturally from innovation and competition: Microsoft and AWS submitted that some factors are intrinsic to cloud services and that the intrinsic complexity means a world in which customers switch more regularly, or design multi-cloud architectures more frequently is unrealistic.¹²⁰⁰
- 6.42 Microsoft also said that the PDR presents a hypothetical scenario of frictionless switching and multi-cloud that does not exist, and in any event would not be what customers want.¹²⁰¹
- 6.43 Our view is that some technical barriers may be inherent or intrinsic,¹²⁰² such that no mitigations can fully address the barrier itself, and different barriers (and different individual frictions between particular services within each of the identified categories of barriers) may be intrinsic to varying degrees. For example, some of the barriers in relation to latency (see below) may be more intrinsic, whereas those related to the current lack of information in the market about which cloud functionalities are proprietary, the potential technical challenges with switching and multi-cloud and how they can be mitigated or overcome are clearly not intrinsic (see below).

¹¹⁹⁸ For example, as set out below, one customer who said that a lack of staff skilled with GCP meant that it was not feasible for it to switch, while another said that skills generally transferable between clouds and were not a factor it considers when choosing a main public cloud.

¹¹⁹⁹ As set out below, the Jigsaw report found that for many customers in the research sample, switching or multi-cloud is not an active consideration, so concerns about specific technical barriers or instances of interoperability were not front of mind and the concerns expressed were often of a more general nature. [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), paragraph 4.1.3.

¹²⁰⁰ [AWS response to the CMAs working papers and updated issues statement 25 June 2024](#), paragraph 63. [Microsoft's response to the Technical Barriers working paper](#), sections 3.2 and 3.3. See further, sections 1.1 and 1.2 of that same submission. Oracle and IBM also submitted that multi-cloud, in particular, was a very complex area that warrants a nuanced approach [IBM response to the Technical Barriers working paper](#), [Oracle response to the CMAs updated issues statement and working papers](#), page 4

¹²⁰¹ [Microsoft's response to the CMA's PDR](#) dated 28 January 2025, para 131.

¹²⁰² We use the term 'intrinsic' rather than 'inherent' to match our guidance. See for instance [CC3 \(Revised\)](#), paragraph 320.

- 6.44 We consider that an assessment of whether each individual technical barrier is intrinsic would not be helpful because cloud services are complex and numerous, and assessing how intrinsic each barrier is for every service or use case would be prohibitively resource-intensive and would fail to capture the fact that the extent to which a particular barrier is intrinsic may change over time.
- 6.45 In addition, it is unfeasible and unnecessary to ascertain the origin of any feature that might be claimed to be intrinsic. For example, if a particular technical barrier exists due wholly or partly to a cloud provider's past design choices we would not consider it to be intrinsic; it would only be intrinsic if it is an unavoidable aspect of the design of cloud services.
- 6.46 Where certain features of cloud services may be intrinsic to the extent that the relevant market cannot realistically be envisaged without them, we can nevertheless consider what effect these have in shaping competition. This is the approach we have taken in our assessment.
- 6.47 We also do not consider it appropriate or necessary to ascertain the motive behind these design choices,¹²⁰³ nor to attribute responsibility for their existence or make findings on the artificiality of a barrier, as our focus is instead on whether technical barriers to switching and multi-cloud affect competition.
- 6.48 Instead, we have assessed the related question of whether the mitigations currently available to customers can address any technical barriers, given their potential costs and benefits, and whether the effect of any technical barriers that may be intrinsic can nevertheless be removed or reduced.
- 6.49 Microsoft and AWS submitted that the PDR had not identified specific services, products or features that create barriers to switching and multi-cloud.¹²⁰⁴
- 6.50 In view of the complexity and number of cloud services, we consider it reasonable to focus our assessment on categories of barriers (and evidence on whether customers experienced the barriers within those categories), rather than identifying and listing out individual frictions between specific pairs/groups of services.
- 6.51 This does not mean that the technical barriers are not real or purely hypothetical. Whilst not all customers named specific services or providers, the challenges they described were often specific. For example, they mentioned difficulties relating to PaaS services such as databases and storage services, differences in interfaces,

¹²⁰³ See further [CC3 \(Revised\)](#), paragraph 159, noting that a feature of the market arising from the conduct of a market participant can even be unintentional.

¹²⁰⁴ [Microsoft's response to the CMA's PDR](#) dated 28 January 2025, paragraphs 130 and 131. [AWS' response to the CMA's PDR](#) dated 18 February 2025, paragraph 38.

differences in identity and access management services and access authentication. These difficulties (and others) are described in more detail below.

Technical barriers to switching and multi-cloud

- 6.52 We have focused on customers' experiences of the technical aspects of attempting to switch between clouds and/or use multiple clouds, as this provides the clearest view of the nature and extent of any barriers. Evidence from other stakeholders, such as cloud providers, ISVs and industry bodies is also included where relevant.

Customer views on switching and multi-cloud

- 6.53 The Jigsaw report found that most customers pointed to technical barriers when asked about their likelihood or willingness to switch clouds, or when asked about their consideration of a multi-cloud approach to their cloud infrastructure.¹²⁰⁵ Many customers cited general concerns about the perceived difficulty of switching or multi-cloud, such as the amount of coding time and resources that would have to be invested in this effort, and the opportunity costs and disruptions to the delivery of digital infrastructure a switching or multi-cloud exercise would mean to the rest of the business.¹²⁰⁶
- 6.54 The Jigsaw report found that customers sometimes struggled to identify specific technical challenges that would have to be overcome to achieve a smooth switching experience and/or an effective multi-cloud infrastructure. It found that this is because, for many customers in the research sample, switching or multi-cloud is not an active consideration, so concerns about specific technical barriers or instances of interoperability were not front of mind and the concerns expressed were often of a more general nature.¹²⁰⁷
- 6.55 The most common response from customers, when asked about technical barriers to switching or multi-cloud, was to mention the amount of recoding such a move would require, ie the need of their software engineers to rewrite a large amount of code to recreate applications for a different cloud environment or to integrate applications across different cloud environments.¹²⁰⁸
- 6.56 The Jigsaw report also found that there were other factors which were considered by customers in their weighing of the costs versus the benefits of switching or using multiple public clouds, which were grouped into two categories: security concerns and disruption to the business' IT service.¹²⁰⁹ Jigsaw also found that

¹²⁰⁵ [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), paragraph 4.1.2.

¹²⁰⁶ [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), paragraph 4.1.4.

¹²⁰⁷ [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), paragraph 4.1.3.

¹²⁰⁸ [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), paragraph 4.1.8.

¹²⁰⁹ [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), paragraph 4.2.1.

there are two key factors that make it more difficult to overcome the technical barriers: a skills gap and a resource gap.¹²¹⁰

- 6.57 A report from a market research firm states that a multi-cloud architecture substantially increases complexity and creates operational challenges: 'Using a multicloud overlay frequently adds significant expense, may reduce functionality and agility, and is usually a major long-term risk-in addition to being a new point of lock-in.'¹²¹¹
- 6.58 The Jigsaw report also found that most participants indicated that currently the technical barriers to switching and multi-cloud weigh heavier than the potential benefits.¹²¹² Some multi-cloud users and switchers in the sample demonstrate these barriers can be overcome in practice, with enough time, effort and resources. However, there seems to be an insufficient incentive for most participants to invest this time, effort and resources to switch clouds or implement a multi-cloud architecture.¹²¹³
- 6.59 Some customers told us that the technical effort to switch and use multiple clouds stems from the fundamental differences in how each of the public clouds have evolved over time.¹²¹⁴ They pointed to differences in approaches, APIs, technical implementations, tools, frameworks, methodologies and best practices and this is consistent with evidence received from some cloud providers discussed below.¹²¹⁵

Multi-cloud costs

- 6.60 In this section, we have considered further the potential costs relating to customers' integration of multiple clouds between applications, within applications and within workloads, as well as integration for the purpose of simplifying management of multiple clouds. We have also considered any operational costs relating to these architectures.

Multi-cloud operational costs

- 6.61 Customers told us that they face additional costs to operate more than one public cloud, regardless of whether they choose to integrate their use of them.¹²¹⁶ For example, one customer explained why one of its subsidiaries had moved from a multi-cloud approach to a single cloud approach. Its internal documents showed that its concerns included the associated increase in the breadth of its cloud estate that must be kept secured (ie the increase in its attack footprint); the increased

¹²¹⁰ [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), paragraph 4.3.

¹²¹¹ [redacted] response to the CMA's information request [redacted].

¹²¹² [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), paragraph 4.4.1.

¹²¹³ [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), paragraph 4.5.9.

¹²¹⁴ Responses to the CMA's information requests [redacted]; Notes of meetings with [redacted]; [redacted] submission to Ofcom [redacted].

¹²¹⁵ Responses to the CMA's information requests [redacted].

¹²¹⁶ Responses to the CMA's information requests [redacted]; Notes of meetings with [redacted].

difficulty in building 'high availability' architectures, decreased scalability and resilience; more patterns and standards to maintain; and increased effort for governance and audit.¹²¹⁷

- 6.62 This customer estimated that the direct cost of maintaining two clouds added an overhead of approximately 5% to its total cloud spend.¹²¹⁸ This customer also said that there are diminishing returns to using multiple clouds, and the fewer clouds, the easier to operate in terms of staff familiarity, connectivity, and security.¹²¹⁹

Multi-cloud integration costs

- 6.63 The evidence relating to the costs that customers incur when integrating multiple public clouds is mixed, consistent with our view that customers are heterogenous.
- 6.64 A supplier of professional services said that the barriers a customer would face in adopting multi-cloud would depend on its individual workloads and connectivity requirements.¹²²⁰
- 6.65 Some customers, ISVs and a supplier of professional services told us that there are challenges to integrating within applications and/or workloads across multiple clouds.¹²²¹ A subset of these customers said that although there are challenges, there are also some workarounds such as using third party tools or building custom solutions to connect services.¹²²²
- 6.66 Other customers told us that they experienced few barriers to integration across multiple public clouds.¹²²³ Reasons given for this included the availability of open APIs¹²²⁴ that make integration easier, workarounds provided by third parties, integrations that are enabled and/or documented by cloud providers and designing architectures such that the integration between clouds is minimal.
- 6.67 Some customers said that they reviewed the option to integrate public clouds but concluded that the benefits did not outweigh the costs of doing so for their current use cases.¹²²⁵ Other customers said that they viewed the benefits of integrating multiple clouds as being too low but didn't mention whether this was in comparison to the costs.¹²²⁶

¹²¹⁷ [REDACTED] response to the CMA's information request [REDACTED].

¹²¹⁸ [REDACTED] response to CMA's information request [REDACTED].

¹²¹⁹ Note of meeting with [REDACTED].

¹²²⁰ Note of meeting with [REDACTED].

¹²²¹ Responses to the CMA's information requests [REDACTED]; Notes of meetings with [REDACTED].

¹²²² Responses to the CMA's information requests [REDACTED]; Note of meeting with [REDACTED].

¹²²³ Responses to the CMA's information requests [REDACTED]; Notes of meetings with [REDACTED].

¹²²⁴ An open API is a free, publicly available application programming interface (API) that allows developers to access software and data.

¹²²⁵ Responses to the CMA's information requests [REDACTED]. Note of meeting with [REDACTED].

¹²²⁶ Responses to the CMA's information requests [REDACTED].

- 6.68 In contrast, some customers are integrating both between and within applications across multiple public clouds, suggesting they saw at least some benefit to doing so.¹²²⁷ For example, one of these customers said that during a trial running a ‘split workload’ across two clouds was a very valuable enabling ‘stepping stone’ in a migration between public clouds.¹²²⁸
- 6.69 Some customers said that multi-cloud is a part of their business strategy,¹²²⁹ and others said whilst they do not currently have a use-case for integrated multi-cloud, they may do so in the future.¹²³⁰ A supplier of professional services said that, of the multi-cloud projects it had supported to conclusion, it observed that customers can/do realise tangible benefits.¹²³¹
- 6.70 Evidence we have seen shows that customers tend to put related workloads on the same public cloud, in order to, for example, reduce operational complexity and prevent a reduction in resilience.¹²³² One customer, on the other hand, told us that it is flexible in placing new workloads due to the benefits of using multiple public clouds.¹²³³
- 6.71 A customer said that when utilising a multi-cloud strategy, consideration must be applied to risks associated with a business process traversing multiple cloud providers, to prevent the impact of a single cloud provider failure from introducing increased operational resilience risk to the business process.¹²³⁴ We understand this to mean that the addition of a second public cloud could double the likelihood that an application fails, if that application were designed in a way that a service outage of either of the public clouds it is dependent on would impact it.
- 6.72 However, another customer said it had been able to avoid this challenge by designing its application such that a user’s experience would not immediately be affected if either cloud stopped operating.¹²³⁵

Switching costs

- 6.73 In this section, we consider the potential switching costs for customers moving workloads from one public cloud to another.
- 6.74 Many customers told us that they anticipate or experience significant costs to switch public clouds.¹²³⁶ Customers described the costs as significant either in

¹²²⁷ Between applications: Responses to the CMA’s information requests [redacted]; Notes of meetings with [redacted]. Within applications: Responses to the CMA’s information requests [redacted]; Notes of meetings with [redacted].

¹²²⁸ Note of meeting with [redacted].

¹²²⁹ Responses to the CMA’s information requests to [redacted].

¹²³⁰ [redacted] response to the CMA’s information request [redacted]; Note of meeting with [redacted].

¹²³¹ [redacted] response to the CMA’s information request [redacted].

¹²³² Notes of meetings with [redacted]; response to the CMA’s information request [redacted].

¹²³³ Note of meeting with [redacted].

¹²³⁴ Note of meeting with [redacted].

¹²³⁵ [redacted] response to the CMA’s information request [redacted].

¹²³⁶ Responses to the CMA’s information requests [redacted]; Notes of meetings with [redacted].

absolute terms – eg a customer said ‘it would take 12 months and tie up approximately 1,000 employees’,¹²³⁷ or in relative terms – eg some customers described technical barriers as the main barrier to switching.¹²³⁸

- 6.75 The Jigsaw report found that switching cloud providers is seen as the equivalent of moving other kinds of infrastructure, such as ‘moving house’ or moving a business from one country to another. It is not something to undertake lightly or consider at all unless it leads to significant business benefits long term that override the perceived cost and risk of changing.¹²³⁹
- 6.76 Some of these customers indicated that these costs had stopped them from switching or considering switching.¹²⁴⁰ For example, a customer said that even if its cloud provider raised all its prices by 5%, this would not be enough of a driver to move everything to a competitor.¹²⁴¹ This customer also said that a switch would cost a similar amount to the initial migration.¹²⁴²
- 6.77 Some customers said that, given the similarity of the current offerings by cloud providers from their perspective, the value of switching is low in comparison to the costs.¹²⁴³ The Jigsaw report also found this.¹²⁴⁴
- 6.78 KPMG, a supplier of professional services, said that most of its clients find switching affordable, but that in general, the size of the resource required (10,000-15,000 hours) poses a big challenge for customers as they could allocate those resources elsewhere. It said that, for example, even if a customer was to save 5-10% by switching cloud provider, it could be better off allocating those resources to developing new products or otherwise building or profitability. It said that this is a reason they see customers choosing not to switch. However, it also said, that if a customer was moving at least 20% of a large estate, a large provider is likely to pay a large part upfront to receive the workload, and that many customers would not move if they did not have this incentive.¹²⁴⁵
- 6.79 Centerprise, another supplier of professional services, said that technical barriers to switching between cloud providers are significant. It said that proprietary technologies and platform-specific services can result in vendor lock-in which make it complex and costly to switch.¹²⁴⁶

¹²³⁷ Note of meeting with [REDACTED].

¹²³⁸ Responses to the CMA’s information requests [REDACTED]; Notes of meetings with [REDACTED].

¹²³⁹ [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), paragraph 1.3.13.

¹²⁴⁰ Notes of meetings with [REDACTED].

¹²⁴¹ Note of meeting with [REDACTED].

¹²⁴² [REDACTED] response to the CMA’s information request [REDACTED].

¹²⁴³ Responses to the CMA’s information requests [REDACTED].

¹²⁴⁴ [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), paragraph 4.5.8.

¹²⁴⁵ Note of meeting with [REDACTED].

¹²⁴⁶ Centerprise’s response to the CMA’s information request [REDACTED].

- 6.80 Open Cloud Coalition said that the PDR correctly highlights the fact that technical barriers (such as lack of standardisation and proprietary APIs) limit portability.¹²⁴⁷
- 6.81 The International Center for Law and Economics said that compatibility issues and switching costs are not zero, but are not insurmountable barriers in practice. It said that many cloud services adhere to common standards (eg Linux environments, containerisation, open-source databases) which allow customers to migrate or balance workloads flexibly.¹²⁴⁸
- 6.82 The International Center for Law and Economics said that there is truth to the notion that moving a complex application from one cloud platform to another is not trivial. It likened cloud infrastructure services to other types of technology, and said that no matter the system, changing providers is not as easy as flipping a light switch or dragging and dropping files on Google Drive.¹²⁴⁹
- 6.83 Rayo, a cloud technology consulting firm, said that ‘lock in’ is often the result of customers making valid choices to use native higher-level services to deliver better business outcomes. It said that it helps organisations identify switching costs, and in its view invariably it is better to use native capability and tooling to save money and improve resilience or security, while being ‘eyes wide open’ about how to migrate.¹²⁵⁰

Operational concerns arising from switching

- 6.84 We have also seen a range of operational concerns that arise when customers consider switching between public clouds.¹²⁵¹
- (a) for example, one customer said that, particularly in relation to IaaS, the barriers to switching are predominantly due to the complexity of moving complex inter-connected architectures that naturally have a high effort and risk of moving at scale.¹²⁵²
 - (b) similarly, a customer’s internal documents showed that factors contributing to the length of its planned exit strategy included the extensive planning, approvals and testing required to switch large, sensitive workloads.¹²⁵³
 - (c) additionally, a public sector customer said that to bring in another cloud vendor to its current portfolio of three, there would need to be a specific ‘unique selling point’ or service not offered by its incumbent suppliers, as it

¹²⁴⁷ [Open Cloud Coalition’s response to the CMA’s PDR](#) dated 28 January 2025, page 4.

¹²⁴⁸ [International Center for Law and Economics’ response to the CMA’s PDR](#) dated 28 January 2025, section II.

¹²⁴⁹ [International Center for Law and Economics’ response to the CMA’s PDR](#) dated 28 January 2025, section IV.

¹²⁵⁰ [Rayo’s response to the CMA’s PDR](#) dated 28 January 2025.

¹²⁵¹ Responses to the CMA’s information requests [REDACTED].

¹²⁵² [REDACTED] response to the CMA’s information request [REDACTED].

¹²⁵³ [REDACTED] response to the CMA’s information request [REDACTED].

would necessitate significant investment and a clear business case even before any actual migration takes place.¹²⁵⁴

6.85 These views are consistent with the findings of the Jigsaw report (see above).¹²⁵⁵

Cloud provider views on the costs of multi-cloud

6.86 Google and Oracle said that there are benefits arising from, and a need for, the ability to integrate multiple public clouds.

(a) Google said it believes there is real customer appetite for integrated multi-cloud strategies, and that integration between multiple clouds is more likely to be adopted for cloud-native¹²⁵⁶ workloads. It said that digital native customers who do not have a historical reliance on legacy on-premises software products, are well-positioned and more likely to adopt multi-cloud strategies. It said that in contrast, traditional enterprises across all sizes and sectors often find an integrated multi-cloud set-up more challenging.¹²⁵⁷ Google said integrated multi-cloud brings significant benefits to customers, including avoiding vendor lock-in, enhanced operational resilience, combining best-in-class cloud products to suit commercial needs, and the ability to carry out pricing arbitrage between two providers to reduce costs.¹²⁵⁸

(b) Oracle submitted that as businesses move to the cloud, they may not need or be able to afford a multi-cloud architecture at the outset. Over time they may then be too entrenched with their first provider to justify spending on multi-cloud (thereby increasing single vendor lock-in). Oracle said that it is critical to ensure multi-cloud is accessible and free from anticompetitive barriers at all stages of a company's IT modernisation timeline.¹²⁵⁹

6.87 However, Microsoft and AWS said that there are operational challenges that disincentivise customers from using a multi-cloud approach.

(a) AWS said that operational challenges include increased data latency, data governance issues, security and data privacy issues due to managing multiple IT environments.¹²⁶⁰ It said that these operational challenges are inherent to integrating multiple IT environments and are not caused by any issues specific to or restrictions imposed by cloud providers.¹²⁶¹ It said that

¹²⁵⁴ [X] response to the CMA's information request [X].

¹²⁵⁵ Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 4.2.1.

¹²⁵⁶ 'Cloud-native workloads' refers to workloads created on the cloud, not migrated from on-premises.

¹²⁵⁷ [X] response to the CMA's information request [X].

¹²⁵⁸ [Google's response to the Egress Fees Working paper](#), paragraph 44.

¹²⁵⁹ [Oracle's response to working papers and updated issues statement](#), pages 3 and 4.

¹²⁶⁰ [AWS' response to the CMA's PDR](#) dated 18 February 2025, paragraph 34.

¹²⁶¹ AWS' response to the CMA's information request [X].

customers are sometimes willing to incur additional costs, if it allows them to choose ‘best of breed’ services.¹²⁶²

- (b) Microsoft said that integrated multi-cloud gives customers the ability to use services from different cloud providers but is generally the least preferred approach by customers. This is because it increases the complexity of building, maintaining and securing applications and also creates multiple points of failure across different clouds. Therefore, customers will typically only choose this model where there is a particularly differentiated service that represents a unique value proposition for their needs.¹²⁶³ Microsoft also said that multi-cloud is less prominent among customers with low spend on cloud services, and this is because there are fixed overheads associated with adopting a multi-cloud architecture, such as additional management costs, reduced performance due to increased latency and increased risk of downtime.¹²⁶⁴ It said however that such customers do have credible alternatives to switch to, should they need.¹²⁶⁵

6.88 Microsoft said that integrated multi-cloud does not necessarily make switching easier, and customers will only adopt ‘Cross-Cloud Applications’ if they see value for their specific use case. It said that where customers see a clear benefit in integrating, they will do so, but not merely for the sake of bargaining power.¹²⁶⁶

6.89 Microsoft said that, despite the fact that customer research shows that the core services of Microsoft, AWS and Google are all roughly equivalent (as each innovates and the others catch up), the CMA created a hypothesis that there are ‘best of breed’ features that customers want to mix and match, whether in siloed or integrated multi-cloud applications.¹²⁶⁷

Cloud provider views on technical barriers to both switching and multi-cloud

6.90 Microsoft, AWS and Google acknowledged that there are technical barriers, but said that these are either inherent or inherent in large part.

- (a) AWS said that there will always be some inherent technical barriers,¹²⁶⁸ but that the introduction of new cloud services, and cloud providers’ efforts to support interoperability, have made switching between IT providers easier than ever before. AWS said that customers can, and do, switch and multi-cloud because of the effective support offered by AWS and other IT

¹²⁶² [AWS’ response to the CMA’s PDR](#) dated 18 February 2025, paragraph 37.

¹²⁶³ Microsoft’s response to the CMA’s information request [§<]. We note that increased risk of downtime and reduced performance due to latency both assume a model in which cloud services are highly dependent on each other for continuous service delivery, which may not be desirable or necessary for all customers.

¹²⁶⁴ [Microsoft response to the Technical Barriers working paper](#), paragraphs 50.

¹²⁶⁵ [Microsoft response to the Technical Barriers working paper](#), paragraphs 44.

¹²⁶⁶ [Microsoft response to the Technical Barriers working paper](#), paragraphs 49.

¹²⁶⁷ [Microsoft response to the Technical Barriers working paper](#), paragraph 27.

¹²⁶⁸ [§<].

providers, who are heavily commercially incentivised to ensure that customers have the ability to multi-cloud and switch if and when they need or want to.¹²⁶⁹ It also said that multi-cloud, or hybrid cloud, can deliver the 'desired outcomes' that might have otherwise motivated customers to switch cloud provider.¹²⁷⁰

- (b) AWS said that its customer feedback indicates that a switch between cloud providers would cost a similar (if not less) amount to the initial migration from on-premises IT infrastructure to a cloud provider. AWS' view is that this further supports that these costs are inherent to IT services, despite the mitigation efforts undertaken by cloud providers.¹²⁷¹
- (c) regarding multi-cloud costs, AWS said that it is not surprising that customers who multi-cloud incur some degree of higher costs than customers using a single cloud, due to the inherent costs of integrating multiple IT environments. However, AWS' view is that customers may be willing to incur these inherent costs if it allows them to enjoy the perceived benefits of multi-clouding.
- (d) Microsoft said that there is inherent complexity in designing and maintaining IT infrastructure which creates friction, both real and perceived, that customers face when switching and/or deploying multi-cloud.¹²⁷²
- (e) Google said that most technical differentiation is the natural result of innovation which it said is one of the hallmarks of effective competition in the cloud market and an important way for smaller cloud providers to differentiate themselves from the two market leaders.¹²⁷³

6.91 Oracle, IBM and Civo said that there are technical barriers to switching, and IBM also noted technical barriers to multi-cloud.

- (a) Oracle submitted that the CMA is correct to identify that currently customers must invest extra effort and resources to mitigate lock-in, though cloud providers can facilitate and ease that additional burden.¹²⁷⁴
- (b) IBM said that it agrees that customers face technical challenges related to switching [X]. IBM said that portability is key and that it allows customers access to the best services for their use cases and preserves cloud

¹²⁶⁹ [AWS' response to the Updated issues statement and working papers](#) dated 23 May 2024 and 06 June 2024, paragraphs 57-58; and [AWS' response to the Technical barriers working paper](#) dated 06 June 2024, paragraphs 4-5 and 107.

¹²⁷⁰ [AWS' response to the CMA's PDR](#) dated 28 January 2025, paragraphs 35 and 49.

¹²⁷¹ [AWS' response to the Updated issues statement and working papers](#) dated 23 May 2024 and 6 June 2024, paragraph 63.

¹²⁷² [Microsoft's response to the Technical barriers working paper](#) dated 06 June 2024, paragraph 1.1.

¹²⁷³ [Google's response to the Technical barriers working paper](#) dated 06 June 2024, paragraphs 4-5 and 9.

¹²⁷⁴ [Oracle's response to the Updated issues statement and working papers](#) dated 23 May 2024 and 06 June 2024, pages 3 and 4.

providers' incentives to innovate.¹²⁷⁵ IBM said that customers face technical barriers when using multiple clouds. IBM added that it is essential to enable integration between clouds for management and for integration between applications [X].¹²⁷⁶

- (c) Civo said that too frequently business managers are not aware of the dangers of implementing proprietary features. It said that the ease by which junior technical staff can access such features, without any governance or control, is encouraged by the 'hyperscalers' and is an effective form of lock-in.¹²⁷⁷

6.92 Cloudflare said that its networking services are designed to be fully interoperable which is the antithesis to the way the dominant cloud providers have traditionally operated. It said that customers face operational challenges such as a lack of interoperability or cloud ecosystems that intentionally preference particular tools.¹²⁷⁸ Cloudflare also said that its customers face significant technical barriers, and issues around interoperability and technical features are central contributors to vendor lock-in. In particular, it mentioned issues it sees relating to cryptographic certificates, which are set out in more detail in the Ancillary Services and Tools section of this chapter.¹²⁷⁹

6.93 OVHcloud said that technical barriers are often seen as the most important obstacle to switching or multi-cloud because they are resource-intensive to overcome.¹²⁸⁰

6.94 Some cloud providers said that technical barriers arising from differences between clouds are a natural result of innovation and competition, from which customers benefit.

- (a) AWS said that IT providers offering proprietary services based on innovative and new technologies is not anti-competitive, adding that, in its view, the ability to profit from innovation is what incentivises competitors to provide new products that best meet their customers' needs.¹²⁸¹ AWS said that customers switch between cloud providers and multi-cloud primarily to benefit from the heterogeneity that rapid innovation and technical differentiation brings about.¹²⁸²

¹²⁷⁵ [IBM's response to the Technical barriers working paper](#) dated 06 June 2024, page 1.

¹²⁷⁶ [IBM's response to the Technical barriers working paper](#) dated 06 June 2024, page 1.

¹²⁷⁷ Civo's response to the CMA's information request [X].

¹²⁷⁸ [Cloudflare's response to the CMA's PDR](#) dated 28 January 2025, page 2.

¹²⁷⁹ [Cloudflare's response to the CMA's PDR](#) dated 28 January 2025, page 6.

¹²⁸⁰ [OVHcloud's response to the CMA's PDR](#) dated 28 January 2025, page 4.

¹²⁸¹ AWS' response to CMA's information request [X], [AWS' response to the CMA's PDR](#) dated 18 February 2025, paragraph 31.

¹²⁸² [AWS' response to the CMA's PDR](#) dated 18 February 2025, paragraph 33b.

- (b) AWS gave the example of its choice to adopt three availability zones per region, which resulted in differences in underlying infrastructure and APIs to other cloud providers.¹²⁸³ It said that whilst this may lead to some technical burden when switching, its customers view its approach as a key reason for choosing AWS.¹²⁸⁴ Similarly, AWS gave another example that most cloud providers offer data storage services. All these services meet the same basic customer need, namely storing data. However, it said that customers are looking for cost-effective storage options and high availability and data durability, among other needs, so this is where cloud providers try to differentiate themselves by offering innovative features or strength in a particular type of storage.¹²⁸⁵
- (c) Microsoft said that cloud providers invest heavily in innovation to differentiate themselves, which brings inevitable complexity to customers' cloud architecture. It added that new cloud services may be inherently less interoperable or portable, if they are the result of technical innovation which is either not available on all clouds or as a result of parallel innovation.¹²⁸⁶ It said that the market is overcoming technical barriers where customers want it, and there will always be some technical barriers to interoperability, as there are in all areas of computing.¹²⁸⁷
- (d) Microsoft also said that differentiation can exist in the form of the cost, security features, scalability/agility, technology and performance, compliance features, sustainability and resilience of cloud infrastructure.¹²⁸⁸ Microsoft said that there may be benefits arising from the extent of differentiation in cloud providers' services, and that they may counterbalance the friction in switching and integrated multi-cloud.¹²⁸⁹ It said that customers weigh up this friction against innovation/technical complexity.¹²⁹⁰
- (e) Oracle said that architectural engineering fundamentally differentiates some of the cloud providers' offerings. It said that it has facilitated a multi-cloud strategy to help customers take advantage of each cloud provider's architectural innovations, even when that innovation creates fundamental differences in engineering. Oracle said that architectural innovation targets the entire stack of technology used to deliver cloud services and can result in better performance, lower costs, higher security and a smaller environmental

¹²⁸³ AWS' submission to the CMA [§]; [AWS' response to the Updated issues statement and working papers](#) dated 23 May 2024 and 6 June 2024, paragraph 64.

¹²⁸⁴ [AWS' response to the Updated issues statement and working papers](#) dated 23 May 2024 and 6 June 2024, paragraph 64.

¹²⁸⁵ [AWS' response to the CMA's PDR](#) dated 18 February 2025, paragraph 33b.

¹²⁸⁶ [Microsoft's response to the Issues Statement](#) dated 17 October 2023, paragraphs 23-25.

¹²⁸⁷ [Microsoft's response to the CMA's PDR](#) dated 28 January 2025, para 118.

¹²⁸⁸ [Microsoft's response to the Issues Statement](#) dated 17 October 2023, paragraphs 23-25.

¹²⁸⁹ [Microsoft's response to the Technical barriers working paper](#) dated 06 June 2024, paragraph 7 (a).

¹²⁹⁰ [Microsoft's response to the CMA's PDR](#) dated 28 January 2025, paragraph 129.

footprint for similar services. It also said that it is differentiated from other cloud providers in part because it optimises for speed and performance.¹²⁹¹

- (f) IBM said that it agrees that differences in features of core services may generate some technical costs but views these costs as ‘inherent to a competitive market and no more than a function of the differentiation between suppliers on the merits’.¹²⁹²

6.95 Other cloud providers expressed different views on technical differentiation.

- (a) Google said that to preserve competition and foster future innovation, it is critical for customers to have unconstrained ability to switch providers and adopt multi-cloud strategies.¹²⁹³
- (b) IBM said that, while innovation increases the quality of service for customers, it may also increase technical switching barriers if improvements are only available to first-party services of a cloud provider [with market power], but not to competing third party services.¹²⁹⁴
- (c) OVHcloud said that the justification for technical differentiation of cloud infrastructure services and cloud ancillary services through proprietary technologies is less clear.¹²⁹⁵

6.96 Some cloud providers said that the customer evidence relating to technical barriers was mixed, and agreed with us that this is because of the heterogeneous nature of cloud customers.

- (a) AWS said that the mixed and contradictory customer views that it said were expressed in the PDR illustrate that technical barriers are specific to each customer, as different customers value different aspects of cloud services, have different requirements, and face unique challenges depending on their use case.¹²⁹⁶
- (b) Google said that, whilst some customers have described potential challenges to switching, other than in respect of IAM services (where it has raised consistent concerns), their feedback is generally mixed. It said that the feedback appears to reflect a more general recognition of the resources and challenges involved in any major IT change/transformation project, rather

¹²⁹¹ [Oracle's response to the Issues Statement](#) dated 17 October 2023, pages 1-2.

¹²⁹² [IBM response to the Technical Barriers working paper](#), page 3.

¹²⁹³ [Google's response to the Issues Statement](#) dated 17 October 2023, paragraphs 3, 10-11 and 18-19.

¹²⁹⁴ [IBM's response to the Issues Statement](#) dated 17 October 2023, paragraph 3.2.

¹²⁹⁵ [OVHcloud's response to the Issues Statement](#) dated 17 October 2023, page 4.

¹²⁹⁶ [AWS' response to the Technical barriers working paper](#) dated 06 June 2024, paragraph 13.

than specific technical barriers to switching that are unique to the cloud market.¹²⁹⁷

- (c) IBM said that differing customer views are to be expected, since customers will have markedly different experiences with cloud depending on their levels of expertise, their industry and their use cases.¹²⁹⁸

- 6.97 Microsoft said that the evidence presented relates almost entirely to anecdotes, primarily from customers, and that these examples do not represent convincing evidence that customers' incentives to switch and multi-cloud are harmed by technical barriers or differentiation.¹²⁹⁹
- 6.98 Microsoft also said that it is not possible for Azure to implement lopsided portability, ie that it is not possible to design a service to be both seamless to switch into and hard to move out of.¹³⁰⁰
- 6.99 Similarly, AWS said that interoperability does not work in only one direction.¹³⁰¹ As an example, AWS said that its S3 APIs have been implemented by many IT services providers, who market these APIs to promote customer migration away from AWS to those providers' services.¹³⁰²

Our assessment of multi-cloud and switching costs

- 6.100 The way that cloud services have developed in the past, as well as elements of competition and innovation, has led to some differentiation in cloud services. Whilst these differences have associated benefits for customers, they also create technical barriers and we considered what effect these have on competition.
- 6.101 Furthermore, even if cloud offerings are considered by many customers to be roughly equivalent, this does not mean there cannot be 'best of breed' services for particular use cases or customer requirements.¹³⁰³ Nor does it rule out the existence of customer benefits that could be realised through a greater ability to switch and multi-cloud. At any stage in the evolution of a market, we expect to see a range of services that are broadly comparable and widely available, alongside

¹²⁹⁷ [Google's response to the Technical Barriers Working Paper](#) dated 06 June 2024, paragraphs 4-5, 9.

¹²⁹⁸ [IBM's response to the Technical barriers working paper](#) dated 06 June 2024, page 1.

¹²⁹⁹ [Microsoft's response to the PDR dated 28 January 2025](#), para 131.

¹³⁰⁰ [Microsoft's response to the Issues Statement](#) dated 17 October 2023, paragraphs 26-32.

¹³⁰¹ AWS submission to the CMA [redacted].

¹³⁰² [Summary of hearing with Amazon Web Services \(AWS\)](#), 12 March 2025, paragraph 27; AWS submission to the CMA [redacted].

¹³⁰³ As set out in the Industry Background, some customers view the cloud providers as having areas of specialism. The Jigsaw research found that while the main operating model for customers is a single cloud provider, or two to three cloud providers with siloed workloads, other models exist in which customers integrate services that are seen as superior (such as GCP BigQuery or Microsoft's IAM) alongside infrastructure from their main cloud provider. [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), paragraphs 1.3.11.

innovative new products and services which are only available from some suppliers.

- 6.102 In relation to Microsoft's views that the evidence is anecdotal, whilst the nature of our evidence is largely qualitative, we note that numerous sources of evidence cover a wide range of different customers and sectors. For example:
- (a) in addition to seeking views directly from customers, we also sought views from suppliers of professional services who work with a wide range of their own customers as well as industry bodies that represent broad segments of industry;
 - (b) the evidence we gathered directly from customers is consistent with the Jigsaw report and other market research, which together represent the experience of a large number of customers from a wide variety of business sectors; and
 - (c) we asked a range of questions including those that considered the present supply of cloud services as well as the ways in which it is continually evolving. Our qualitative approach enabled us to get a deeper understanding of stakeholder views.
- 6.103 Our view is that it will in many cases be easier for customers to switch to cloud services than away from them, regardless of whether the services are designed for 'lopsided portability'.
- 6.104 Customers typically make their initial migration to the cloud by lifting-and-shifting on-premises workloads to IaaS, which can be relatively simple. Theoretically, customers can also lift-and-shift their applications between clouds after they have initially migrated from on-premises, which again does not entail significant technical challenges. However, as customers upgrade their architectures to make use of a range of PaaS services, they generally then find it harder to switch and multi-cloud. Similarly, services could be designed such that they are easy to migrate to using a simple configuration, but more difficult to migrate away from once specific features or a more complex configuration is adopted.
- 6.105 In terms of multi-cloud, although some customers are able to multi-cloud (in some cases due to their willingness to adopt a workaround, or because their desired integration was minimal), technical barriers to multi-cloud do negatively affect other customers' decisions to use and integrate multiple public clouds.
- 6.106 Customers often think about using a multi-cloud architecture in cost-benefit terms, meaning that they weigh the costs – including technical costs of doing so – against perceived benefits. Customers expressed differing views on the extent to which there were benefits to multi-cloud but we note that customers, professional service

firms and cloud providers all recognised that, in particular circumstances, there are benefits to adopting a multi-cloud architecture.

- 6.107 Due to their varying needs and requirements, not all customers will see value in integrated multi-cloud approaches at any one time or for any specific use case. However, we still find the existence of technical barriers that impact customers' ability to multi-cloud a cause for concern because this necessarily limits customers' ability to exercise this choice.
- 6.108 In terms of switching, the costs to overcome or mitigate technical barriers affect customers' ability to switch between public clouds. In general, switching costs are seen by customers as significant.
- 6.109 Whilst customers are able to switch if they are willing to expend all the required resources to do so, customers think about doing so in cost-benefit terms, and some customers would not switch even if they were forgoing significant benefits in failing to do so.
- 6.110 Below we look in more detail at the technical barriers to switching and multi-cloud in relation to core services, ancillary services and other technical factors.

Core services

- 6.111 We have considered whether technical differentiation of the features and interfaces of core services is itself a technical barrier to switching and multi-cloud.
- 6.112 As described above, core services are the main IaaS and PaaS services that contribute to the key objectives of the workload.
- 6.113 Differentiation of features relate to the functionality of similar core services hosted on different public clouds, and differentiation of interfaces affects the interfaces (eg protocols or APIs) of comparable core services hosted on different public clouds by cloud providers and/or ISVs.

Cloud provider views

- 6.114 All cloud providers that commented on this topic, including Microsoft and AWS, said that there are differences in the features and interfaces of core services.¹³⁰⁴
- 6.115 AWS also said that these are specific to each IT provider and, although they might lead to some inherent technical burden when switching between IT services

¹³⁰⁴ [AWS' response to the Technical barriers working paper](#) dated 06 June 2024, paragraphs 16, 19 and 22; [Microsoft's response to the Issues Statement](#) dated 17 October 2023, paragraphs 23 and 24; [Google's response to the Technical barriers working paper](#) dated 06 June 2024, paragraphs 3-5; [IBM's response to the Issues Statement](#) dated 17 October 2023, paragraphs 1.6 and 1.7; [OVHcloud's response to the Issues Statement](#) dated 17 October 2023, pages 3-5; [Company A's response to the Issues Statement](#) dated 17 October 2023, paragraph 3.2; Notes of meetings with [X]; Responses to the CMA's information requests [X].

providers, customers choose their specific IT providers because of such differences.¹³⁰⁵

- 6.116 AWS said that PaaS is not necessarily more differentiated than IaaS and that while it is true that moving a VM from on-premises to the cloud or between clouds is easier than moving PaaS applications, the CMA overestimates the difficulty of moving PaaS applications from one cloud provider to another.¹³⁰⁶ As an example, AWS said that there are performance, reliability, and other characteristics of IaaS that vary significantly between cloud providers, depending on each providers' unique hardware, supply chains, infrastructure design choices, and operational practices.¹³⁰⁷
- 6.117 AWS said that feature and interface differentiation is inherent to IT services and reflective of a healthy level of competition, as it indicates high levels of innovation and customer choice. AWS said that, while such differentiation might lead to some inherent technical burden when switching between IT services providers, customers choose their specific IT providers because of such differences.¹³⁰⁸
- 6.118 AWS also said that it offers customers a range of services for a given 'use case', some of which are highly interoperable, and some of which contain proprietary technology. It said that this trade off is a feature of a highly competitive market, and customers are able to choose the solutions that work for them.¹³⁰⁹
- 6.119 AWS gave the example that customers can choose to set up entirely self-managed databases using AWS infrastructure and leveraging the standard protocols and open-source software that they prefer (eg, by running MySQL on EC2), or they can instead choose to use a managed database service like Amazon RDS. It said that the differentiated features of RDS do not mean that there is less interoperability, or that customers are locked in, because of the standard database protocols and languages used, the ability to use open-source database engines, the ability to move data from RDS to other databases and the ability to connect third party tools.¹³¹⁰
- 6.120 AWS also identified some other features which it said differentiate its services from its competitors', such as [REDACTED].¹³¹¹
- 6.121 AWS also said that the evidence on differentiation of features and interfaces is mixed and cannot support a conclusion that such differentiation results in any AEC. It said that the evidence includes several pieces of feedback which illustrate

¹³⁰⁵ [AWS' response to the Technical barriers working paper](#) dated 06 June 2024, paragraphs 16, 19 and 22.

¹³⁰⁶ [AWS' response to the Technical barriers working paper](#) dated 06 June 2024, paragraph 63.

¹³⁰⁷ AWS' submission to the CMA [REDACTED].

¹³⁰⁸ [AWS' response to the Technical barriers working paper](#) dated 06 June 2024, paragraphs 14 and 16.

¹³⁰⁹ [AWS' response to the CMA's PDR](#) dated 28 January 2025, paragraph 33(a).

¹³¹⁰ [AWS' response to the CMA's PDR](#) dated 28 January 2025, paragraph 32(c).

¹³¹¹ AWS' submission to the CMA [REDACTED].

that feature differentiation and interface differentiation do not raise any concerns.¹³¹²

Customer views

- 6.122 In general, customers said that there are differences in the features and interfaces of core services and this hinders their ability to switch and multi-cloud.¹³¹³
- 6.123 This was echoed by the Jigsaw report, which said that that the most common technical challenge to switching and multi-cloud cited by customers was the amount of recoding involved, driven by several factors, including differences in interfaces, features and capabilities across providers.¹³¹⁴ There were also some customers that said that differences in features or interfaces of core services did not affect their ability to switch or multi-cloud.¹³¹⁵
- 6.124 We heard that as the technical complexity of a customer's workloads increases, the difficulty of switching increases. Technical complexity is the result of a customer's chosen architecture, but relevant factors include the number of the workloads, the degree dependence on more differentiated services, and strict operational or compliance requirements.¹³¹⁶

Differentiation of features

- 6.125 In relation to differentiation of features, customers were concerned about the impact on their ability to switch between clouds, with some noting the need to plan, remap, rework and test workloads.¹³¹⁷
- 6.126 For example, one customer said that the differences in the services of different cloud providers can require rework or mapping workloads to different services, which increases the cost of switching.¹³¹⁸ Some customers also said that, while cloud providers offer similar features in some of their core services, there are still differences in how these features have been implemented.¹³¹⁹
- 6.127 Some customers also said that differences in features also increase the cost of using multiple clouds.¹³²⁰ For example, a customer said that there are circumstances where multiple providers offer identical solutions but that there are

¹³¹² [AWS' response to the Technical barriers working paper](#) dated 06 June 2024, paragraphs 17 and 20.

¹³¹³ Responses to the CMA's information requests [REDACTED]; Notes of meetings with [REDACTED].

¹³¹⁴ [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), paragraphs 4.1.9.

¹³¹⁵ Responses to the CMA's information requests [REDACTED]; Notes of meetings with [REDACTED].

¹³¹⁶ [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), paragraphs 4.1.13; Responses to the CMA's information requests [REDACTED]; Notes of meetings with [REDACTED].

¹³¹⁷ Responses to the CMA's information requests [REDACTED]; Notes of meetings with [REDACTED].

¹³¹⁸ [REDACTED] response to the CMA's information request [REDACTED].

¹³¹⁹ Responses to the CMA's information requests [REDACTED].

¹³²⁰ Responses to the CMA's information requests [REDACTED]; Notes of meetings with [REDACTED].

also situations where functional and non-functional differences exist which complicate the wider use of multiple clouds.¹³²¹

6.128 Customers were particularly concerned with the differentiation of features of PaaS services.¹³²² For example:

- (a) one customer said that the use of PaaS with proprietary features significantly increases the complexity of switching to a different cloud and would result in a loss of functionality in most cases.¹³²³
- (b) another customer said that functional differences in the way certain PaaS services operate, including their configuration and functionality cause vendor lock-in. It said that therefore, fully embracing the PaaS services of one provider can mean it is much harder to then move to a different provider.¹³²⁴

6.129 This is consistent with the Jigsaw report which found that customers have a choice between accepting lack of code portability (and hence greater vendor lock-in) or not taking advantage of vendor-specific PaaS solutions.¹³²⁵

6.130 Customers also mentioned specific types of PaaS services which are particularly challenging to switch away from, such as data warehouse and analytics services:¹³²⁶

- (a) for example, a customer said, as an example of engineering effort required to switch clouds, that GCP Big Query uses a particular type of SQL and so if a company wanted to migrate to AWS Redshift, it would have to conduct SQL conversion.¹³²⁷
- (b) similarly, an internal document from a customer contained estimates of 'one-off migration costs' associated with switching an application from AWS to GCP. This document stated that a principal driver of these costs was the need to re-engineer applications to use GCP services such as databases, message queues, and caches.¹³²⁸

6.131 This was echoed in the Jigsaw report, which stated that challenges migrating databases and storage services are among the most commonly raised by customers and can necessitate rewriting database queries or other code.¹³²⁹

¹³²¹ [X] response to the CMA's information request [X].

¹³²² Responses to the CMA's information requests [X]; Notes of meetings with [X].

¹³²³ [X] response to the CMA's information request [X].

¹³²⁴ [X] response to the CMA's information request [X].

¹³²⁵ [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), paragraphs 4.1.3.

¹³²⁶ Responses to the CMA's information requests [X]; Notes of meetings with [X].

¹³²⁷ Note of meeting with [X].

¹³²⁸ [X] response to the CMA's information request [X].

¹³²⁹ [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), paragraphs 4.1.14 and 4.5.2.

Differentiation of interfaces

- 6.132 In relation to differentiation of interfaces, customers were concerned about the impact on their ability to multi-cloud, with many noting the need to create and use workarounds.¹³³⁰ For example, a customer said that each cloud has different APIs and so it has to use third-party tools to integrate them, as it is not as easy to have systems talk across clouds.¹³³¹
- 6.133 Some customers also said that differences in interfaces also increase the cost of switching by necessitating work to recreate the integrations of a switched workload with other services or workloads.¹³³² For example, a customer said that different interfaces would mean Infrastructure as Code¹³³³ and other software managing applications must be re-written when switching clouds.¹³³⁴
- 6.134 This was echoed in the Jigsaw report, which stated that challenges due to inconsistent APIs and interfaces are among the most commonly raised by customers.¹³³⁵
- 6.135 Some customers said that whilst there were differences in APIs for similar services across clouds, integration using different but open APIs did not cause major challenges.¹³³⁶ For example, a customer said that it does not find differences in APIs to be an insurmountable challenge because it can use translation layers to convert between different clouds or have ‘over-the-top’ solutions that support multiple versions at once in a multi-cloud architecture.¹³³⁷

Other stakeholders

- 6.136 ISVs, suppliers of professional services and other stakeholders also said that there are differences in the features and of interfaces of core infrastructure services. Similar to customers, these other parties generally said that such differentiation hinders customers’ ability to switch and multi-cloud.¹³³⁸ For example, an ISV said that the main restrictions to porting to any cloud are down to differences in automation technique and topology between public clouds, as well as differences in management of resources and infrastructure in each cloud.¹³³⁹

¹³³⁰ Responses to the CMA’s information requests [§]; Note of meeting with [§].

¹³³¹ [§] response to the CMA’s information request [§].

¹³³² Responses to the CMA’s information requests [§]; Note of meeting with [§].

¹³³³ Infrastructure as Code (IaC) is the use of high-level descriptive coding language to automate and standardise the provisioning and deployment of IT infrastructure such as networks, virtual machines, load balancers, and connection topologies required by any application.

¹³³⁴ Note of meeting with [§].

¹³³⁵ [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), paragraphs 4.1.9 and 4.5.2.

¹³³⁶ Responses to the CMA’s information requests [§]; Notes of meetings with [§].

¹³³⁷ [§] response to the CMA’s information request [§].

¹³³⁸ [§] submission to the CMA [§]; Notes of meetings with [§].

¹³³⁹ Note of meeting with [§].

- 6.137 The International Center for Law and Economics said that switching ‘core’ cloud functions is far easier than in the pre-cloud era, and many industry standards and open-source tools exist to minimise lock-in. It said that the apparent ease of multi-homing suggests that other compatibility issues are not a major hinderance to competitive pressures, and that there is still robust competition.¹³⁴⁰

Conclusion on core services

- 6.138 In general, core services hosted on different clouds are technically differentiated in relation to both their features and interfaces. The differentiation of features manifests in the feature itself and/or its implementation, ie in *what* the feature is doing and/or *how* it is doing it. While such differentiation may bring benefits in terms of innovation and increased customer choice, it has associated costs which we set out below.
- 6.139 Based on evidence from customers and other parties, we have found that differentiation of features of core services hinders at least some customers’ ability to use multiple clouds.
- 6.140 The evidence more broadly showed that differentiation of features of core services, particularly PaaS services, leads to a lack of product substitutability and thereby hinders customers’ ability to switch clouds. Many customers, some other organisations and some smaller cloud providers detailed how the differences in features increase the cost of switching, including the need to plan, remap, rework and test workloads. This was also echoed in the Jigsaw report.
- 6.141 PaaS in particular is generally highly differentiated which creates a trade-off for some customers between taking advantage of perceived benefits of using PaaS, but limiting their ability to switch and multi-cloud, and using IaaS, but limiting their use of PaaS features. This is based on evidence from customers, other organisations and smaller cloud providers. It is also consistent with the findings of the Jigsaw report.
- 6.142 Differentiation of features in core services does not impact all customers in the same way and to the same extent, as customer needs are heterogeneous as described above.
- 6.143 In terms of the differentiation of interfaces of core services, these hinder customers’ ability to integrate multiple clouds. Many customers and other organisations said that the differences in APIs of core services were significant and detailed how this increased technical cost when integrating multiple clouds

¹³⁴⁰ [International Center for Law and Economics's response to the CMA's PDR](#) dated 28 January 2025, section IV.

and required creating or using workarounds. This was also echoed in the Jigsaw report.

- 6.144 In addition, based on evidence from some customers and other parties, we found that differentiation of interfaces of core services hinders at least some customers' ability to switch clouds. As with differentiation of features, differentiation of interfaces in core services does not impact all customers in the same way and to the same extent.
- 6.145 We recognise that there are some technical mitigations to the current differentiation of features and interfaces of core services but the evidence showed that these mitigations only overcome the technical barriers to a limited extent. We have discussed these below at Technical mitigations.

Ancillary services and tools

- 6.146 We have considered whether the design and implementation of ancillary services and related software tools separately leads to technical barriers to switching and multi-cloud.¹³⁴¹
- 6.147 Ancillary services and tools provide functions that support the operation and management of core services. They include IAM, billing, observability, and provisioning and orchestration.¹³⁴²
- 6.148 By 'differentiation of ancillary services and tools' we mean differences (in features and/or interfaces) of comparable ancillary services and tools hosted on different public clouds and provided by any supplier, including cloud providers and ISVs.
- 6.149 We begin by setting out evidence and our analysis relevant to differentiation of ancillary services and tools in general. We then set out evidence and our analysis in relation to IAM because we heard from stakeholders that IAM is especially important for customers who wish to switch or multi-cloud.

Differentiation of ancillary services and tools

- 6.150 In principle there are various ways in which differentiation of ancillary services and tools can impact customers' ability to switch and multi-cloud. For example, if a customer has a multi-cloud architecture and the ancillary services and tools hosted on the relevant clouds are incompatible, the customer may need to run them

¹³⁴¹ As noted above, we use the term 'ancillary services' in this section to refer to managed ancillary services and, in that context, tools' to refer to ancillary software that customers deploy and manage on their own.

¹³⁴² IAM is ancillary services and tools that allow customers to control who (person or application) can access what in the cloud. Billing is the mechanisms used to monitor, analyse, and charge for cloud services. Observability is the process of measuring, analysing and visualising the current state of a customer's cloud architectures based on the data it generates, such as logs, metrics, and traces. Provisioning and orchestration is the process of automating the tasks needed to manage deployment, connections and operations of workloads.

separately, without integration and automation, which can lead to increased multi-cloud operational costs. If a customer wants to switch and there is differentiation of ancillary services and tools between the origin and target clouds, the customer will need to make additional changes to its applications, data and associated tools so that they can work and perform well on the target cloud.¹³⁴³

Cloud provider views

- 6.151 AWS and Google said that generally ancillary services and tools are differentiated.¹³⁴⁴
- 6.152 AWS said that technical barriers related to differentiation of features and interfaces of ancillary services and tools are inherent to IT services.¹³⁴⁵
- 6.153 AWS said that our customer evidence on ancillary services illustrates why technical barriers are specific to each customer and simply a product of the inherent technical barriers that it works hard not to exacerbate but cannot ever fully mitigate. AWS noted two of its ancillary services (AWS Glue and Amazon CloudWatch) that it said support interoperability and multi-clouding.¹³⁴⁶ It also said that ‘monitoring is not necessarily a “dependency” that prevents containers from promoting portability’.¹³⁴⁷
- 6.154 IBM suggested that there are technical barriers associated with ancillary services and tools that hinder customers’ ability to switch or multi-cloud. In particular, IBM said that concerns may arise where there is a lack of portability of ‘Day2 operations tools’ which are used once an application is running, eg to monitor performance, ensure security, compliance, logging and metering. It said that a non-portable ‘Day2 operation tool’ has to be recreated when using another cloud provider, which will involve additional technical work (mainly recoding) and may require different skillsets.¹³⁴⁸
- 6.155 In addition, IBM said that it can be a lot of work for providers of ancillary services which enable portability to keep up with changes in underlying providers’ public clouds.¹³⁴⁹
- 6.156 A cloud provider said that when a customer uses proprietary ancillary services this can create lock in even if the customer uses otherwise portable services, such as

¹³⁴³ Differentiation of core services could also cause customers who multi-cloud or switch to expend resources redeveloping or reconfiguring the associated ancillary services and tools.

¹³⁴⁴ [AWS' response to the CMA's Technical barriers working paper](#) dated 06 June 2024, section III(B); [AWS' response to the Updated issues statement and working papers](#) dated 23 May 2024 and 6 June 2024, paragraph 65; and [Google's response to the Technical barriers working paper](#) dated 06 June 2024.

¹³⁴⁵ [AWS' response to the CMA's Technical barriers working paper](#) dated 06 June 2024, paragraph 31.

¹³⁴⁶ [AWS' response to the CMA's Technical barriers working paper](#) dated 06 June 2024, paragraphs 40 and 41.

¹³⁴⁷ [AWS' response to the CMA's Technical barriers working paper](#) dated 06 June 2024, paragraph 40 and 66(b).

¹³⁴⁸ [IBM's response to the Issues Statement, 17 October 2023](#), page 3; IBM's response to the CMA's information request [redacted].

¹³⁴⁹ Note of meeting with IBM [redacted].

IaaS services.¹³⁵⁰ It said that ensuring that access to cloud providers' services is protected and enhanced is essential to improving competition in the market.

[REDACTED]¹³⁵¹ [REDACTED].¹³⁵²

- 6.157 Cloudflare submitted that applications from the three main cloud providers are designed in such a way that limits customers' ability to use their own cryptographic certificates. It said that ultimately this limits customers' ability to use alternative providers of security services and tools.¹³⁵³ In particular, [REDACTED] which has implications for security as well as competition.¹³⁵⁴

Customer views

- 6.158 Many customers we contacted said that the differentiation of ancillary services and tools increases the complexity and time in switching between clouds and managing multiple clouds.¹³⁵⁵ For example:

- (a) one customer said that differences in APIs of ancillary services and tools require reworking of deployment pipelines, tooling or code.¹³⁵⁶
- (b) another customer told us that the methods and philosophy behind ancillary services and tools are often completely different, which means it is hard to have one approach to multi-cloud. It also told us that differentiation of ancillary services and tools strongly impacts its ability to switch, since infrastructure as code (IaC) and other software managing applications needs to be re-written or adapted to work with another cloud.¹³⁵⁷
- (c) another customer told us that each tooling 'jump' between suppliers of cloud services requires bridging a technical and conceptual gap, for example in relation to different security frameworks and different systems of access control.¹³⁵⁸ It also said that cloud billing is painful and difficult to work with, and a lack of standardisation on what different fields mean in its bills makes comparing bills difficult across cloud providers. This customer said it had to develop an expensive workaround, including an in-house tool to make billing data comparable across clouds and present it in a central dashboard.¹³⁵⁹

- 6.159 Another customer told us that the biggest barrier to switching is tooling and telemetry, which is usually tied to the specific cloud platform. It said that any

¹³⁵⁰ Note of meeting with [REDACTED]; [REDACTED] submission to the CMA [REDACTED].

¹³⁵¹ See for example [REDACTED].

¹³⁵² [REDACTED] submission to the CMA [REDACTED].

¹³⁵³ [Cloudflare's response to the CMA's PDR](#) dated 28 January 2025, page 6.

¹³⁵⁴ Cloudflare's submission to the CMA [REDACTED].

¹³⁵⁵ Responses to the CMA's information requests [REDACTED].

¹³⁵⁶ [REDACTED] response to the CMA's information request [REDACTED].

¹³⁵⁷ [REDACTED] response to the CMA's information request [REDACTED].

¹³⁵⁸ Note of meeting with [REDACTED].

¹³⁵⁹ Note of meeting with [REDACTED].

migration would not only need re-engineering of the application or platform, but all of the security and observability configuration.¹³⁶⁰ However, some other customers told us that differences between ancillary services and tools do not necessarily impact their ability to switch or multi-cloud.¹³⁶¹

Other stakeholder views

- 6.160 An ISV, a supplier of professional services and two other organisations also suggested that there are technical barriers associated with ancillary services and tools that hinder customers' ability to switch and multi-cloud.
- (a) an ISV said that the difficulties in switching and multi-cloud arise from ancillary services and tools,¹³⁶² as well as the resource and staffing required to make the necessary changes to conduct the switch or to manage more than one cloud. It said that mapping from one service to another is achievable, but it is the technical effort required to make this operational that is the main source of difficulty. It said that the developer environments that customers use are highly specialised to specific clouds and this makes it difficult to deploy code from these environments on different clouds.¹³⁶³
 - (b) a supplier of professional services said that setting up and running multi-cloud is reasonably simple, but running it efficiently is quite complex and difficult. It said that for running an integrated multi-cloud architecture efficiently, a customer needs ancillary services and tools that cover all layers of their multi-cloud architecture, including infrastructure, network, APIs, applications and customer experience. It said that, while the integration of such services is technically possible, it is also complex and requires specific skillsets and significant technical effort. It also said that ancillary services and tools differ significantly on how they perform their functions. For example, it said that some ancillary services and tools produce different data in a different format and with different frequency.¹³⁶⁴
 - (c) a non-profit organisation told us that there is a large ecosystem of tools needed for application development, which are often unique to individual clouds. It said that its concern is less about mobility in the cloud, but rather

¹³⁶⁰ [X] response to the CMA's information request [X].

¹³⁶¹ Responses to the CMA's information requests [X] response to the CMA's information request [X]. Notably, this customer also told us that IAM can pose challenges to integration as there are variances in what different cloud providers standardise on, but that most other ancillary services and tools provide clear API interfaces that are reasonably easy to integrate with.

¹³⁶² The ISV said that these include observability, custom tools the customer has written (or off the shelf tooling), specialised Continuous Integration/Continuous Deployment (CI/CD) pipelines and associated custom workflows.

¹³⁶³ Note of meeting with [X].

¹³⁶⁴ Ofcom note of meeting with [X].

the lock-in of developers to clouds, through the cost and time of retraining on these tools.¹³⁶⁵

- (d) Linux Foundation told us that that using proprietary services or add-on features can lead to barriers to switching and integrated multi-cloud.¹³⁶⁶

6.161 In addition, a market research firm report states that [REDACTED].¹³⁶⁷

6.162 ISVs and others told us that using cloud-agnostic ancillary services and tools in particular can help reduce technical barriers associated with switching and multi-cloud. For example:

- (a) Okta, an ISV, said that a cloud-agnostic IAM service can help simplify managing IAM-related logs within a multi-cloud architecture.¹³⁶⁸
- (b) another ISV said that cloud-agnostic provisioning tools can be used to simultaneously set up cloud services from multiple cloud providers.¹³⁶⁹
- (c) Linux Foundation told us that OpenTelemetry exemplifies the case in which value-added features unique to specific cloud providers become commoditised, and open-source standards naturally arise.¹³⁷⁰

Conclusion on ancillary services and tools

6.163 The evidence set out above shows that, in general, ancillary services and tools hosted on different clouds are technically differentiated both in relation to their features and interfaces.

6.164 Based on evidence from customers and cloud providers, we have found that the differentiation of ancillary services and tools hinders customers' ability to switch and multi-cloud.

6.165 We also consider that the differentiation of ancillary services and tools does not impact all customers in the same way and to the same extent, as customer needs are heterogeneous (as described above).

6.166 We recognise that there are some ways in which customers can avoid or mitigate the effects of this differentiation. For example, we heard from a customer that some cloud-agnostic ancillary services and tools can have a positive impact on customers' ability to switch and multi-cloud since the same ancillary service or tool

¹³⁶⁵ Note of meeting with [REDACTED].

¹³⁶⁶ Note of meeting with the Linux Foundation [REDACTED].

¹³⁶⁷ [REDACTED].

¹³⁶⁸ Okta's response to the CMA's information request [REDACTED].

¹³⁶⁹ Note of meeting with [REDACTED].

¹³⁷⁰ Note of meeting with the Linux Foundation [REDACTED].

can be used across different clouds.¹³⁷¹ However, our evidence on technical mitigations shows that these are not fully effective. We discuss these below at Technical mitigations.

Identity and Access Management (IAM)

- 6.167 IAM is part of the broader category of security services and tools which aim to secure customers' cloud architectures. IAM services and tools allow customers to control which services, data, or resources each user or application can access in the cloud. IAM services and tools perform two main functions: authorising access and authenticating identity.¹³⁷²
- 6.168 We heard from various stakeholders that IAM is an important consideration and/or poses technical challenges for customers who switch or multi-cloud, which is why we consider it separately here. For example:
- (a) Okta, said that customers who use the public cloud must use the cloud provider's own IAM service for certain functionalities, in particular the 'identity layer' that enables access into the cloud provider's cloud applications.¹³⁷³ Customers are able to choose other IAM services and tools for additional functionalities and/or for multi-cloud purposes; and
 - (b) the Jigsaw report noted that one of the most mentioned examples of integrated multi-cloud was using Microsoft's IAM services in combination with cloud infrastructure services from a different cloud provider.¹³⁷⁴
- 6.169 We also received submissions that relate specifically to Microsoft's IAM services and tools (and the integration between these and other Microsoft software), which we set out below.

Cloud provider views

- 6.170 AWS said that IAM services do not act as an artificial technical barrier to switching and multi-cloud. It said that naturally, IAM has to be specific to each provider, not because of artificial 'barriers' imposed by cloud providers but because each provider has its own underlying infrastructure and security architecture, which necessitates a specific IAM and cannot be homogenised. It said that differences in authorisation across different IT environments are inherent limitations that should not be framed as a cloud-specific issue.¹³⁷⁵ AWS said that it supports a number of industry-standard IAM protocols which allows customers to store their identities

¹³⁷¹ Note of meeting with [REDACTED].

¹³⁷² 'What is Cloud Identity and Access Management?', accessed on 7 July 2025.

¹³⁷³ Note of meeting with Okta [REDACTED].

¹³⁷⁴ Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 3.5.5

(b).

¹³⁷⁵ AWS' submission to the CMA [REDACTED].

with a third-party identity services provider, and then manage and use those identities with AWS services or external applications, but that not all applications or IAM services may meet the same standards, as different services meet different customer needs and support standards based on those needs.¹³⁷⁶ It said, for example, that 'IAM Roles Anywhere' allows AWS customers to authenticate external workloads and apply permissions to those workloads. It also identified other features of AWS services which support interoperability standards (such as Cognito Identity Pools).¹³⁷⁷

- 6.171 AWS shared a list of product feature requests made by customers in the last 12 months in relation to multi-cloud functionalities, and we note that some of these requests involved technical challenges with IAM.¹³⁷⁸
- 6.172 AWS said that it has invested heavily in security-related solutions that help customers switch between IT services or use IT services across different suppliers. It has developed Cedar, which is an open-source policy language and authorisation engine for fine-grained permissions management, and an associated cloud service known as Amazon Verified Permissions.¹³⁷⁹
- 6.173 We note that AWS Cedar has been adopted by at least one ISV, in a service that helps customers develop software applications.¹³⁸⁰
- 6.174 Google told us that IAM is an important consideration for customers' choice of cloud provider, and that customers who use public cloud typically use the chosen cloud provider's own IAM service/tool for the main functions of an IAM service.¹³⁸¹
- 6.175 IBM said that technical differences between core IAM services from different clouds can significantly contribute to the challenge of using IAM services across clouds. It said that cloud providers have the option of implementing standards and protocols related to how identity information is managed and processed which can result in differences, eg in the level of abstraction provided by IAM services or the granularity of access supported. It also said differences in the ecosystem and related configurations can also result in added complexity to the use of IAM services across clouds.¹³⁸²
- 6.176 Oracle gave IAM as an example of the technical difficulties which it or its customers can encounter when seeking to integrate its services with those of other public clouds. In particular, this cloud provider said that authorisation policies are cloud-specific and concepts used by one cloud provider may not have exact

¹³⁷⁶ AWS' submission to the CMA [§].

¹³⁷⁷ AWS' submissions to the CMA [§].

¹³⁷⁸ AWS' response to the CMA's information request [§].

¹³⁷⁹ [AWS' response to the CMA's updated issues statement and working papers](#) dated 23 May 2024 and 6 June 2024, paragraph 66(b).

¹³⁸⁰ 'Permit.io Cedar Implementation Q&A: Everything you need to know', accessed on 7 July 2025.

¹³⁸¹ [Google's response to the CMA's Technical barriers working paper](#) dated 06 June 2024, paragraph 12.

¹³⁸² [§] response to the CMA's information request [§].

translations in another public cloud. For example, it said that different cloud providers may represent users in different ways or as different/multiple entities. It also said these differences may contribute to difficulty for customers, but that depends on factors including how familiar or comfortable users are with one cloud provider's resources and authorisation policies versus another cloud provider's. It said that customers will likely find the options that cause them the least amount of work and/or rework to implement and/or maintain cross-cloud solutions.¹³⁸³

- 6.177 Conversely, Oracle said that there are no significant feature differences between IAM services amongst cloud providers and that, in general, the major platforms all seem to support common industry standards.¹³⁸⁴
- 6.178 Cloudflare told us that practices of the three main cloud providers mean that customers are required to rely on certain providers of IAM services.¹³⁸⁵
- 6.179 Civo told us that IAM is an area in which standardisation is important since it could simplify cross-cloud access control and security management, which would provide a unified approach to managing users, roles and permissions.¹³⁸⁶
- 6.180 Microsoft told us that IAM is an example of a service that customers wish to unify across their environments, because this simplifies the authentication of users across all of the customer's applications.¹³⁸⁷
- 6.181 Some cloud providers submitted specific comments on identity authentication as one of the main functions of IAM.¹³⁸⁸ For example, Microsoft told us that identity authentication is well-understood and commonly solved, and most enterprises using multiple cloud services deploy some kind of single sign-on.¹³⁸⁹
- 6.182 However, a cloud provider said that each cloud provider has its own nuances in the implementations of the interoperable frameworks (OAuth, SAML and OIDC) in terms of configuration and access mapping, which can impact consistency.¹³⁹⁰ We understand that these nuances can add to the technical costs involved for customers who set up this type of integration.

Customer views

- 6.183 Many customers said that IAM is an important consideration and/or poses technical challenges when switching between clouds or adopting a multi-cloud

¹³⁸³ Oracle's response to the CMA's information request [REDACTED].

¹³⁸⁴ Oracle's response to the CMA's information request [REDACTED].

¹³⁸⁵ [Cloudflare's response to the CMA's PDR](#) dated 28 January 2025, page 6.

¹³⁸⁶ Civo's response to the CMA's information request [REDACTED].

¹³⁸⁷ Microsoft's response to the CMA's information request [REDACTED].

¹³⁸⁸ [AWS' response to the CMA's Technical barriers working paper](#) dated 06 June 2024, paragraphs 32-36.

¹³⁸⁹ [Microsoft's submission to the CMA dated 8 August 2024, CMA Cloud Services MIR - Identity and Access Management Follow-up](#), page 2.

¹³⁹⁰ [REDACTED] response to the CMA's information request [REDACTED].

architecture.¹³⁹¹ Most of these customers said that the differentiation of IAM across clouds was either the main reason or a significant reason for these challenges.¹³⁹²

- 6.184 Two customers, told us that, for them, IAM may not pose technical challenges to switching or multi-cloud of associated workloads between clouds, since they would be able to do so without also switching their current IAM service and the cloud it is hosted on.¹³⁹³
- 6.185 Some customers provided specific comments on access authentication as one of the main functions of IAM. They said that differences in access management can lead to additional costs for customers who switch or multi-cloud.¹³⁹⁴ For example:
- (a) one customer said that it is experiencing increasing challenges in IAM, with individual cloud providers' PaaS offerings becoming tightly coupled with their IAM services. It said that while this simplifies managing business services operating on one cloud, it increases the effort to exit any provider because the access policies need to be translated and reproduced into another provider's context.¹³⁹⁵
 - (b) another customer said that access management is a challenge in IAM. It gave the example of IAM policy scripts which it said are particularly important for security. In particular, it said that IAM policy scripts are currently not standardised between public clouds (eg AWS and Azure) and require deep knowledge, which is not transferable between public clouds.¹³⁹⁶
- 6.186 This is consistent with the Jigsaw report, which found that authentication methods are seen to be different between providers, which it said suggests that the portability of proprietary IAM tools is particularly difficult.¹³⁹⁷
- 6.187 Some evidence from customers is specifically related to identity authentication as one of the main functions of IAM. From the evidence we obtained from customers¹³⁹⁸ and from the Jigsaw report, we understand that many customers who integrate multiple public clouds are able to integrate different IAM services and tools to some degree for identity authentication.¹³⁹⁹

¹³⁹¹ Responses to the CMA's information requests [§]; Notes of meetings with [§]; [§] submission to Ofcom [§].

¹³⁹² Responses to the CMA's information requests [§]; [§] submission to Ofcom [§].

¹³⁹³ Responses to the CMA's information requests [§].

¹³⁹⁴ Responses to the CMA's information requests [§]; Notes of meetings with [§].

¹³⁹⁵ [§] response to the CMA's information request [§].

¹³⁹⁶ Note of meeting with [§].

¹³⁹⁷ [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), paragraph 4.1.9

(b).

¹³⁹⁸ Responses to the CMA's information requests [§]; Note of meeting with [§]; [§] response to Ofcom's information request [§].

¹³⁹⁹ This type of integration allows a customer to use a single IAM service (such as Microsoft Entra ID or Okta) through which users (staff) sign in to different cloud services (so do not need to keep track of multiple usernames and passwords, which can be burdensome (see Note of meeting with [§])).

6.188 However, we also heard that there are still some limitations with this type of integration:

- (a) Some customers noted technical challenges associated with the integration of identity authorisation functionality across clouds or switching suppliers.¹⁴⁰⁰
- (b) A customer said that it is possible to integrate Okta into multiple clouds, but the integration is complex because the integration efforts mostly sit with the customer and entail continuous work as integrations need to be applied to any new applications, services, workloads. This customer said that as a result, from an integration perspective, it may be simpler to use just one cloud but that is not necessarily the right choice.¹⁴⁰¹

Other stakeholder views

6.189 Two ISVs said that there are technical differences between the IAM services of different cloud providers, and in particular how access policies work, and suggested that this can be a technical barrier to customers' ability to switch and multi-cloud:

- (a) An ISV said that access management in general lacks the same level of industry standards as identity management. However, it said that, in its view, AWS is working to solve this through the open-source Cedar Policy Language and that there may be other options.¹⁴⁰² It also said that if a customer wants to migrate from Entra ID, for example, to another IAM service that supports common standards (such as OIDC and OAuth) then there should be limited friction, since Entra ID offers the same features and follows these same standards. It added however that such a migration also involves migrating data, the specifics of which it is not familiar with.¹⁴⁰³
- (b) Okta told us that previous attempts to solve issues similar to differentiation of IAM access policies (such as the development of Extensible Access Control Markup Language, XACML) have not been successful.¹⁴⁰⁴

6.190 Research from a market research firm states that IAM infrastructure will differ between cloud platforms, so some provisions will have to be made to manage enterprise users and their access privileges on multiple cloud platforms. It states that although all cloud providers offer a range of security and IAM capabilities,

¹⁴⁰⁰ Responses to the CMA's information requests [§]; Notes of meetings with [§].

¹⁴⁰¹ Ofcom note of meeting with [§].

¹⁴⁰² [§] response to the CMA's information request [§].

¹⁴⁰³ Note of meeting with [§].

¹⁴⁰⁴ Note of meeting with Okta [§].

there is no common approach among the cloud providers, each has its own combination of features, and some providers are still working to close gaps.¹⁴⁰⁵

- 6.191 Specifically in relation to AWS Cedar, Okta said that for Cedar to significantly mitigate access management-related challenges for customers who switch or multi-cloud, it would need to be used as the basis of a new common IAM policy language or framework specifically for cloud services, which would then have to be adopted by multiple cloud providers in a consistent way. Okta also said that Cedar is just a policy language and is not the complete set of components needed for interoperability of authorisation policy. It said that without interoperability, there is no easy way to simply reuse policies across different clouds. Okta said industry working groups like AuthZEN are attempting to define some of the missing APIs needed for interoperability.¹⁴⁰⁶

Conclusion on IAM

- 6.192 We found that IAM is an important factor for customers' ability to switch and integrate multiple public clouds. Many customers, as well as some ISVs and cloud providers, noted that differences in access management policies in particular increase the costs of customers when they switch or multi-cloud.
- 6.193 The development and consistent adoption by cloud providers and ISVs of common IAM standards (such as SAML, OIDC and OAuth) has helped facilitate customers' ability to integrate multiple clouds (via single sign-on) and potentially to a lesser extent, switch between clouds. Some cloud providers have services (such as AWS Cedar or Amazon Verified Permissions) that help customers develop software relating to IAM.¹⁴⁰⁷
- 6.194 However, we found that such standards and services are not sufficient to mitigate the technical barriers related to IAM set out above (such as the differentiation of access management policies), and the design and implementation of IAM services and tools hinders customers' ability to switch and multi-cloud.

Microsoft's IAM services and tools

- 6.195 We have received submissions that specifically relate to Microsoft's IAM services and tools, including Microsoft's cloud IAM service (Entra ID), Microsoft's on-premises IAM software (Active Directory) and the integration between these and other Microsoft software.

¹⁴⁰⁵ [§] response to the CMA's information request [§].

¹⁴⁰⁶ Note of meeting with Okta [§]; Okta submission to the CMA [§]. AuthZEN is an industry working group.

¹⁴⁰⁷ We understand that Amazon Verified Permissions is a cloud service that customers can use to help them develop custom software applications, including both consumer applications such as photo sharing as well as internal software applications, such as an HR system.¹⁴⁰⁷ It is therefore unlikely to significantly help mitigate challenges for customers who switch between IAM services or integrate multiple public clouds.

- 6.196 Use of Active Directory may affect customers' ability to switch and multi-cloud, for example if they use a hybrid cloud architecture with both a public cloud IAM service and Active Directory on-premises. Okta said that most customers use a mix of cloud-based and on-premises software, and that enterprise customers need additional IAM services and tools to synchronise their cloud-based IAM service with their on-premises IAM software.¹⁴⁰⁸
- 6.197 We set out the substance of these submissions in our PDR and sought stakeholder views. In addition, we conducted further targeted evidence-gathering on this issue following the publication of the PDR to further understand these issues. As explained further in our assessment below, although it does appear to us that there are frictions related to Microsoft's IAM services and tools, the evidence we have considered does not show an impact on switching and multi-cloud arising from these frictions that is distinct from the wider issues we have found in relation to IAM.

Stakeholder submissions relating to Microsoft's IAM services and tools

- 6.198 We have received stakeholder submissions on the following three issues relating to Microsoft's IAM services and tools:
- (a) Microsoft Entra ID's interoperability with other providers' IAM services;
 - (b) technical integration between Entra ID and Active Directory; and
 - (c) technical integration between Entra ID and other Microsoft software and services using Entra ID.
- 6.199 We have heard the following from stakeholders about how aspects of identity authentication (including 'federation'¹⁴⁰⁹) works for Entra ID compared to other cloud providers' IAM services:
- (a) two cloud providers told us that there are technical limitations with Entra ID which limit customers' ability to synchronise IAM information between Entra ID and other cloud-based IAM services. They said that this creates friction for customers and consequently Entra ID often remains a customer's primary IAM service.¹⁴¹⁰

¹⁴⁰⁸ Okta's response to the CMA's information request [38].

¹⁴⁰⁹ The work needed to make identity authentication work across multiple cloud providers is commonly referred to as 'federation' (see [Microsoft's submission to the CMA dated 8 August 2024, CMA Cloud Services MIR - Identity and Access Management Follow-up](#)).

¹⁴¹⁰ Responses to the CMA's information requests [38]. One cloud provider explained the steps required to create a custom integration with Entra ID, which included i) using Microsoft's Graph API, which requires developers to learn the API and create the objects and attributes mappings; ii) building a Graph API-specific anti-entropy solution (to avoid inconsistencies in data being synchronised) and iii) create a mechanism for connecting to Graph API to make changes. The identity service provider would need to test and maintain these custom synchronisation implementations for Graph

- (b) an organisation told us that Microsoft has failed to support standard protocols in Entra ID for some customers in certain use cases, even though it supports standard protocols for customers keeping their identities in Entra ID. It suggested that Microsoft should ensure that Entra ID supports an identity industry standard (such as SCIM) for full interoperability, instead of requiring customers to use proprietary protocols such as the Microsoft Graph API.¹⁴¹¹
- (c) an ISV said that even though customers who use Entra ID can federate from a third-party IAM service, if they do not use Entra ID as their primary IAM provider then they do not have access to features such as Microsoft Autopilot, Windows Hello for Business and Conditional Access.¹⁴¹²

6.200 In relation to these concerns, Microsoft said that its IAM services are comparable to the services offered by other public cloud providers in terms of their core functionalities and features, but that there are technical and commercial differences between first-party IAM services offered by different cloud providers that reflect their development and use in different settings.

6.201 It said [REDACTED].¹⁴¹³

6.202 Microsoft added that it offers a high degree of parity for customers using third-party IAM services to access Azure services, compared to customers who use Microsoft's own IAM service, Entra ID.¹⁴¹⁴ Microsoft said that while integrating cloud services from multiple vendors necessarily has an overhead, the solutions to interoperating amongst multiple IAM services are well understood and enterprises operate effectively today.¹⁴¹⁵

6.203 We have also considered the following stakeholder submissions on whether there are any technical barriers which hinder the ability of customers who use Active Directory to use non-Microsoft IAM services and whether customers who use Active Directory to manage Windows Server-based workloads are consequently more likely to migrate those workloads to Azure:

- (a) Google told us that Microsoft has imposed artificial technical restrictions in relation to Active Directory which exacerbate the lock-in effects of its licensing practices and are the main technical barrier that customers face.¹⁴¹⁶ Google told us that Microsoft's IAM services and tools create additional barriers, beyond those associated with IAM more generally. It said that

API, in addition to their SCIM implementation. It added that since each cloud provider's integration is different, customers need to separately evaluate each integration, which causes delays and adds friction.

¹⁴¹¹ [REDACTED] submission to the CMA [REDACTED].

¹⁴¹² [REDACTED] response to the CMA's information request [REDACTED].

¹⁴¹³ Microsoft's response to the CMA's information request [REDACTED].

¹⁴¹⁴ Microsoft's response to the CMA's information request [REDACTED].

¹⁴¹⁵ [Microsoft's submission to the CMA dated 8 August 2024, CMA Cloud Services MIR - Identity and Access Management Follow-up.](#)

¹⁴¹⁶ [Google's response to the Technical barriers working paper](#) dated 06 June 2024, Section II.

Microsoft aims to foreclose third-party IAM services, and indirectly, rival cloud providers.¹⁴¹⁷

- (b) specifically, Google told us that customers who use Active Directory and migrate to the cloud are more likely to use Entra ID. It said that Entra ID is not a replacement for Active Directory and customers are in practice required to continue using both Active Directory and Entra ID in a hybrid cloud architecture. It also said that customers' dependency on Active Directory reflects their dependence on Windows Server. It also said that some functionality of Active Directory (such as managing Group Policies) is not available in Entra ID, which means that customers with legacy Microsoft workloads are effectively tied into using a combination of Active Directory and Entra ID.
- (c) a cloud provider said that, in its view, Microsoft does not support modern protocols such as SAML and OIDC in Active Directory and does not provide sufficient access to Active Directory's APIs or technical information, which it said makes it very challenging for users to integrate Active Directory with IAM services and tools other than Microsoft's own cloud IAM service, Entra ID.¹⁴¹⁸ It also said that non-Microsoft cloud providers must typically rely on offering a managed version of Active Directory, which for technical reasons¹⁴¹⁹ is inherently limited compared to Entra ID (specifically, limited compared to the equivalent feature of Entra ID known as Azure Domain Services). It said that as a result, most customers with Microsoft-related workloads continue using their on-premises Active Directory with Entra ID.¹⁴²⁰
- (d) an ISV told us that with Active Directory, Microsoft is de facto the only provider of on-premises IAM software and that even customers who shift towards a cloud-first IAM approach still retain some dependence on Active

¹⁴¹⁷ [Google's response the CMA's PDR](#) dated 28 January 2025, paragraph 40; Google's responses to the CMA's information requests [§]; [Google's response to the Technical barriers working paper dated 06 June 2024](#), paragraph 16.

¹⁴¹⁸ [§] response to the CMA's information request [§].

¹⁴¹⁹ As an example, this cloud provider said that Azure customers can authenticate to Azure-hosted Windows virtual machines using their Entra ID credentials over Remote Desktop Protocol (RDP). It said that conversely, non-Azure customers can authenticate to Windows-hosted virtual machines using the protocol specified in Microsoft's Open Specifications documentation (RDPEAR), but that this, on its own, would not enable customers to offer a single sign-on (SSO) experience, which is (according to this cloud provider) highly valued by customers as it enables seamless authentication across cloud services and their own remote desktop instances. This cloud provider said that offering SSO would require synchronising the customer's main cloud IAM service with managed Active Directory service, which is an additional technical hurdle. In addition, this cloud provider said that it is necessary to synchronise across the non-Azure cloud identity directory and Active Directory to achieve SSO and that third-party cloud IAM providers must rely on AD Directory Sync, which is based on the Lightweight Directory Access Protocol (LDAP) that is (according to this cloud provider) dated and is unsuitable for synchronising on-premises IAM tools with cloud IAM services (and is primarily suited for supporting IAM workloads hosted in an on-premises environment). It said that if Microsoft should be required to either make Active Directory compatible with SCIM or make its proprietary solution (Entra Connect) and the necessary information available to third-party IAM providers in a way that is comparable to how this solution is made available to Entra ID (which, according to this cloud provider, offers superior synchronisation capabilities with Active Directory). This cloud provider also said that Microsoft should make all necessary Windows authentication systems' API documentation available to all third-party IAM providers or make AD Domain Services a multi-cloud offering on 'FRAND' terms such that it can be hosted on third-party cloud providers' infrastructure.

¹⁴²⁰ [§] response to the CMA's information request [§].

Directory, as it is typically costly and disruptive to switch to an entirely cloud-based environment. It said that for technical reasons, customers who use a third-party IAM solution can face synchronisation delays and other workflow problems, which can ultimately incentivise customers to continue using Active Directory and Entra ID in their hybrid environments.¹⁴²¹

- (e) another cloud provider identified ways in which Entra ID offers features that cannot be matched by other cloud providers' IAM services. For example, it said that Microsoft's Entra ID Domain Services provides a managed Active Directory service to Azure customers, with a built-in capability to synchronise Entra ID users into Active Directory, including their passwords. This password-synchronising capability does not exist in Microsoft's Graph API, which means that customers cannot sync their Entra ID user passwords to other identity service providers or manage Active Directory outside of Microsoft's Entra ID Domain Services product.¹⁴²²
- (f) Google told us that Microsoft's Open Specifications documentation is not sufficient, and it is not true that other IAM providers can integrate with Active Directory. It said that Microsoft selectively strikes deals with some providers to give them special access.¹⁴²³
- (g) Okta said that Microsoft's Open Specifications documentation does not cover all of the necessary protocols and is often vague. It said that the provision of detailed documentation and technical specifications (and the required tools and APIs) to third-party providers of IAM services would promote further interoperability.¹⁴²⁴
- (h) some customers indicated that they do use both Active Directory and Entra ID in a hybrid cloud architecture.¹⁴²⁵ In some cases customers told us that they chose Entra ID based on their use of Active Directory. For example, one customer told us that Windows Server (which contains Active Directory) has built-in capabilities to join Entra ID in numerous ways and is designed specifically with this in mind.¹⁴²⁶

6.204 Regarding this issue, Microsoft said that as enterprises migrate from Active Directory to cloud-based IAM services, suppliers offer competing solutions to help with migration or hybrid deployments. It said that it publishes comprehensive documentation on Active Directory through its Open Specifications documentation,

¹⁴²¹ [redacted] response to the CMA's information request [redacted].

¹⁴²² [redacted] response to the CMA's information request [redacted].

¹⁴²³ [Google's response to the CMA's PDR](#) dated 28 January 2025, paragraph 41.

¹⁴²⁴ Okta's response to the CMA's information request [redacted].

¹⁴²⁵ Responses to the CMA's information requests [redacted].

¹⁴²⁶ [redacted] response to the CMA's information request [redacted].

which enables complete interoperability. It said that multiple IAM providers (such as Okta) support comprehensive integration with Active Directory.¹⁴²⁷

6.205 Finally, we have considered the following stakeholder submissions on whether customers who use other Microsoft software and services may choose Entra ID over competing IAM services and whether this may affect those customers' ability to switch between clouds and use multiple clouds in general and their ability to switch between different IAM services in particular.¹⁴²⁸ Cloud providers submitted that:

- (a) once customers adopt Entra ID they are then unlikely to switch away due to its 'integration' with other Microsoft products, and its lack of interoperability with third party identity solutions.¹⁴²⁹
- (b) Microsoft has hindered third parties' cloud identity management tool technical interoperability with Microsoft 365. It said that using competitive collaboration tools is thus more complex.¹⁴³⁰
- (c) the only way to use a third-party IAM service to connect to Microsoft Azure (in order to use Azure cloud services) is by syncing it with a corresponding Microsoft Entra ID. It said that in this way Microsoft forces customers to use its own cloud IAM service.¹⁴³¹
- (d) if Entra ID supported inbound SCIM, customers could use any compatible IAM service that supports SCIM as their primary identity service provider for their workforce identities, then connect their identities from their primary identity service provider via SCIM to Entra ID, and then manage and use those identities with Microsoft applications such as Microsoft 365. It said that instead, Microsoft only supports outbound SCIM, and by also automatically enabling Entra ID when enterprise customers adopt Microsoft 365, and given that there is a lack of support for inbound SCIM, the net effect is that customers are pressured to adopt Entra ID as their primary identity service provider.¹⁴³²

6.206 CISPE suggested that it is impossible to use a third-party IAM cloud service in order to properly manage Microsoft 365. It also said that Microsoft does not disclose the interoperability information, nor provide any APIs needed, to allow

¹⁴²⁷ Microsoft's submission to the CMA dated 8 August 2024, CMA Cloud Services MIR - Identity and Access Management Follow-up.

¹⁴²⁸ Google's response to the Technical barriers working paper dated 06 June 2024, paragraph 24; Submissions to the CMA [REDACTED]; Responses to the CMA's information requests to [REDACTED]; Note of meeting with [REDACTED].

¹⁴²⁹ Google's response to the Technical barriers working paper dated 06 June 2024, paragraph 24.

¹⁴³⁰ [REDACTED] submission to the CMA [REDACTED].

¹⁴³¹ [REDACTED] response to the CMA's information request [REDACTED].

¹⁴³² [REDACTED] response to the CMA's information request [REDACTED].

interoperability between Entra ID and third-party identity management products.¹⁴³³

- 6.207 A customer said that its use of Microsoft 365 means that it is required to use Entra ID. It said that this is an example where it is not possible to use alternatives.¹⁴³⁴ Another customer said that using Entra ID does not necessarily incentivise or force them to use other Azure cloud services.¹⁴³⁵
- 6.208 Microsoft said that, in general, each cloud provider determines how to manage authentication for its services and most providers rely on their own IAM services to accomplish this. It said that users of Google Workspace must authenticate to Google's identity system, users of Salesforce must authenticate to Salesforce and users of Office 365 must authenticate to Microsoft's identity system, Entra ID.¹⁴³⁶
- 6.209 Microsoft said that integrating cloud services from multiple suppliers necessarily has an overhead, but that the reality is that enterprises already deal with this scenario and the solutions to interoperating amongst multiple IAM services are well-understood and operate effectively today.¹⁴³⁷

Our assessment

- 6.210 We have considered whether there are technical challenges associated with Microsoft's IAM services and tools that are different from and/or additional to the frictions we have identified in relation to IAM more generally.
- 6.211 There appear to be some frictions specifically related to Microsoft's IAM services and tools, particularly where this relates to Microsoft Entra ID's interoperability with other providers' IAM services. For example, Microsoft's IAM services do not have built-in support for certain types of two-way synchronisation between Entra ID and other cloud-based IAM services. This is consistent with our findings on IAM more generally.
- 6.212 Stakeholders also submitted that there are additional frictions arising from the ways in which Entra ID may be better able to integrate with Active Directory and Microsoft software than other providers' IAM services and tools are able to. We note that Microsoft has supplied Open Specifications documentation to assist with software integration to mitigate these frictions, but that stakeholders consider that this could be improved to promote further interoperability.

¹⁴³³ CISPE's submission to the CMA [§].

¹⁴³⁴ [§] submission to the CMA [§]; Note of meeting with [§].

¹⁴³⁵ [§] response to Ofcom's information request [§].

¹⁴³⁶ [Microsoft's submission to the CMA dated 8 August 2024, CMA Cloud Services MIR - Identity and Access Management Follow-up](#), page 2.

¹⁴³⁷ [Microsoft's submission to the CMA dated 8 August 2024, CMA Cloud Services MIR - Identity and Access Management Follow-up](#), page 4.

6.213 We consider that the evidence we have gathered following consultation with stakeholders and in response to information requests raising questions about these issues supports our finding earlier in this chapter that all providers' IAM services and tools have an impact on switching and multi-cloud. However, this evidence is not sufficient to support a finding that Microsoft's conduct in relation to IAM has an impact on customers that is additional to the interoperability concerns we have identified across the market more broadly.

Other technical factors

6.214 We have considered whether certain other technical factors are also barriers to switching and multi-cloud, specifically data latency, skills and transparency.

Latency

6.215 Latency refers to the time it takes for data to travel between any two points on a network. In relation to public cloud, this could be the time it takes for data to move within a single data centre, between the data centres of the same public cloud, between different public clouds or between the public cloud and other IT environments.

6.216 Latency is bound by the speed of light. This means that latency is generally lower when data is travelling over shorter distances. Latency is also determined by the speed of processing at any interim nodes on the path of data transfer. This means that latency is lower when data is travelling directly between the origin and target points of the network and higher when it must pass through various nodes that process the data between these points.

Cloud provider views

6.217 Some cloud providers told us that increased cost and higher latency can be an unavoidable reality when attempting to integrate between cloud platforms which may discourage customers from integrating a single workload across multiple clouds.¹⁴³⁸ For example, AWS said that when a single solution is spread between multiple cloud providers, information may need to flow many hundreds of miles across the internet to move between services. This increases latency and cost due to the additional time it takes to transfer data between cloud providers.¹⁴³⁹

6.218 Similarly, AWS said that some customers require ultra-low latency and sometimes only one cloud provider is geographically close enough to offer it, or the customer chooses to purchase hardware that it can place in its own premises, because no

¹⁴³⁸ [Microsoft's response to the Issues Statement](#) dated 17 October 2023, paragraph 34; [Google's response to the CMA's Technical barriers working paper](#) dated 06 June 2024, paragraph 26(b); Responses to the CMA's information requests [X].

¹⁴³⁹ AWS' response to the CMA's information request [X].

third-party infrastructure can meet its needs. In such circumstances, there is not much cloud providers can do, other than offering infrastructure that is even closer, which may not be feasible.¹⁴⁴⁰

- 6.219 Cloud providers highlighted existing mitigations for latency in their responses.¹⁴⁴¹ For example, Microsoft said that direct interconnects between cloud providers' data centres already exist and do not resolve problems for latency-sensitive applications.¹⁴⁴²
- 6.220 A cloud provider said that latency is primarily relevant for a limited subset of customers and that these customers have ways to mitigate latency concerns in a multi-cloud infrastructure by using existing technology, [REDACTED].¹⁴⁴³

Customer views

- 6.221 We heard from a few customers that latency requirements mean that applications perform more strongly when they are located on a single cloud platform, particularly when they are processing real-time data.¹⁴⁴⁴ This is because transferring data between separate clouds means it may have to traverse intermediate steps such as over the public internet.
- 6.222 The Jigsaw report found that customers which need to process or transmit data in real time see latency as a barrier to multi-cloud because it will likely decrease the speed of their applications and workloads and that it is therefore more logical to keep such workloads on a single cloud provider.¹⁴⁴⁵
- 6.223 However, some customers did not see latency as an impediment to operating a multi-cloud architecture.¹⁴⁴⁶ One customer said that it previously moved some of its latency sensitive workloads during a trial but carried on using the database stored in the previous cloud provider. The customer explained that the latency between the two cloud providers within the same region was low and did not cause any issues; it could even have maintained this architecture for an extended period.¹⁴⁴⁷

¹⁴⁴⁰ [AWS' response to the CMA's Technical barriers working paper](#) dated 06 June 2024, paragraph 45.

¹⁴⁴¹ [Microsoft's response to the CMA's Technical barriers working paper](#) dated 06 June 2024, paragraph 78; and [Google's response to the CMA's Technical barriers working paper](#) dated 06 June 2024, paragraph 26(b).

¹⁴⁴² [Microsoft's response to the CMA's Technical barriers working paper](#) dated 06 June 2024, paragraph 78.

¹⁴⁴³ [REDACTED] submission to the CMA [REDACTED].

¹⁴⁴⁴ Responses to the CMA's information requests [REDACTED].

¹⁴⁴⁵ Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 4.1.15.

¹⁴⁴⁶ Responses to the CMA's information requests [REDACTED].

¹⁴⁴⁷ Notes of meetings with [REDACTED].

Other stakeholder views

- 6.224 Two ISVs said that latency can be an obstacle to multi-cloud because it can impact the performance and functionality of customers' applications.¹⁴⁴⁸

Conclusion on latency

- 6.225 The evidence shows that latency is an issue for a subset of customers which require real-time or near real-time data transfers and in such cases, the effect of latency contributes to the cumulative technical cost of integrating multiple clouds.
- 6.226 There are mitigations available to customers which may reduce the effect of latency, but we consider that these are not fully effective. Direct connections between clouds via IXPs can reduce latency for data transfers as compared to transfers over the internet, especially when the data centres are located close to each other, but the remaining latency may still be too high for latency-sensitive workloads and latency will remain significant for cross-region workloads. Co-location of data centres can also mitigate the effects of latency, but this has so far had very limited take-up by cloud providers.

Skills

- 6.227 To effectively manage its cloud architecture, an organisation will need employees who are skilled and experienced in cloud engineering. This discipline is highly technical which means that organisations must employ skilled professionals and train their existing workforce to optimise their environments.

Cloud provider views

- 6.228 Cloud providers told us about the training they offer to customers.¹⁴⁴⁹
- (a) AWS said that one of its most popular training courses is one which informs customers on how to migrate workloads from other providers to AWS.¹⁴⁵⁰
 - (b) Microsoft submitted that, to the extent that there are skill gaps, professional services firms have emerged to provide support for developers.¹⁴⁵¹
 - (c) Google said that whilst there is a degree of technical differentiation between services because of innovation, cloud providers offer broadly the same core services and the principles underpinning these are common across the

¹⁴⁴⁸ Notes of meetings with [REDACTED].

¹⁴⁴⁹ [AWS' response to the CMA's Technical barriers working paper](#) dated 06 June 2024, paragraph 66(f); and [Microsoft's response to the CMA's Technical barriers working paper](#) dated 06 June 2024, paragraph 82.

¹⁴⁵⁰ AWS' response to the CMA's information [REDACTED].

¹⁴⁵¹ [Microsoft's response to the CMA's Technical barriers working paper](#) dated 06 June 2024, paragraph 86.

industry. It said that engineers certified on other clouds can easily become certified on GCP.¹⁴⁵²

- 6.229 However, a small cloud provider said that the three main cloud providers, through comprehensive training and certifications, deeply influence students and professionals, swaying their cloud preferences from an early stage. It said that aggressive outreach (by the three main cloud providers) can create a tech ecosystem where new talent is predominantly trained and biased towards a single cloud vendor, limiting multi-cloud knowledge and curbing future diversification.¹⁴⁵³
- 6.230 Similarly, another small cloud provider (BlackBox Hosting) said that the free training provided by AWS and Azure, which is specific to the public cloud, further disincentivises customers from switching.¹⁴⁵⁴

Customer views

- 6.231 Some customers told us that the skillsets required by their staff to operate on different cloud providers are distinct and that retraining staff significantly increases costs associated with switching providers or operating an integrated multi-cloud architecture.¹⁴⁵⁵
- 6.232 Some customers also said that they make decisions related to their cloud architecture based, at least in part, on the skillsets that are already present in their organisations.¹⁴⁵⁶ For example, one customer said that while GCP is a credible alternative cloud provider, it would be very difficult to switch because it has no skills and experience with it.¹⁴⁵⁷
- 6.233 We also heard from some customers that the differentiation in skillsets required for each cloud necessitates separate teams for each; the cost of this can be prohibitive to a multi-cloud architecture.¹⁴⁵⁸ These costs include the recruitment difficulties that can arise from having to hire new cloud engineers to operate on different clouds.¹⁴⁵⁹
- 6.234 In contrast, some customers said that they have found it relatively easy to retrain staff in other cloud providers' services, or that the skills they have are transferable between providers.¹⁴⁶⁰

¹⁴⁵² [Google's response to the CMA's Technical barriers working paper](#) dated 06 June 2024, paragraph 26(a).

¹⁴⁵³ [Cloud Provider 1's response to the Issues Statement](#) dated 17 October 2023, section 3.

¹⁴⁵⁴ [Black Box Hosting's response to the CMA's PDR](#) dated 28 January 2025, paragraph 18.

¹⁴⁵⁵ Responses to the CMA's information requests [redacted]; [redacted] response to Ofcom's information request [redacted].

¹⁴⁵⁶ Responses to the CMA's information requests [redacted]. [redacted]; Note of meeting with [redacted].

¹⁴⁵⁷ [redacted] response to the CMA's information request [redacted].

¹⁴⁵⁸ Responses to the CMA's information requests [redacted]; Notes of meetings with [redacted].

¹⁴⁵⁹ Responses to the CMA's information requests [redacted] submission to the CMA [redacted].

¹⁴⁶⁰ Responses to the CMA's information requests [redacted].

- 6.235 Some customers said that the training offered by cloud providers is an important consideration and can be an enabler of multi-cloud.¹⁴⁶¹
- 6.236 The Jigsaw report found that exacerbating factors to the technical challenges in the report are the major skills gap (both in customers' current teams and a difficult hiring environment) and the resource gap, ie the engineering resources that switching or multi-cloud use would take away from their core business. The research further highlighted that most cloud engineers are expert in only one cloud provider, which necessitates expensive training or hiring new teams if seeking to switch or multi-cloud.¹⁴⁶²

Other stakeholder views

- 6.237 An ISV said it did not find it difficult to train its staff to operate across providers. It added however, that software engineers would have to learn some concepts around fault isolation and operations when learning how to work with a new provider and that this was akin to learning a new syntax.¹⁴⁶³
- 6.238 Another ISV said that it is not possible to train a human engineer to be expertly versed in all the clouds, as there is too much minutiae in the management of the services. It said that there is no good solution and that in general it is complex to switch between any cloud.¹⁴⁶⁴
- 6.239 Some suppliers of professional services said that skills can be a problem when switching because companies are hesitant to train staff on multiple clouds due to cost and competition for staff skilled in multiple clouds.¹⁴⁶⁵ One (KPMG) said that it is the biggest migration challenge facing customers.¹⁴⁶⁶
- 6.240 KPMG also said that, whilst it does not take much effort for cloud engineers to become skilled with multiple providers, many deliberately choose not to and instead prefer to specialise on one cloud. It also added that there is a 'tribal' element to the preferences of engineers and how each cloud is perceived, which is encouraged by the providers.¹⁴⁶⁷
- 6.241 An industry body said that there is now a large ecosystem of tools which are needed for developing applications and these tools are often unique to individual

¹⁴⁶¹ Responses to the CMA's information requests [§].

¹⁴⁶² [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), paragraphs 4.3.4-4.3.14 and 4.5.4.

¹⁴⁶³ Note of meeting with [§].

¹⁴⁶⁴ Note of meeting with [§].

¹⁴⁶⁵ Notes of meetings with [§].

¹⁴⁶⁶ Note of meeting with KPMG [§].

¹⁴⁶⁷ Note of meeting with KPMG [§].

clouds. Its concern is less about mobility in the cloud, rather the lock-in of developers to clouds, through the cost and time of retraining on these tools.¹⁴⁶⁸

Conclusion on skills

- 6.242 We found that the differentiation between cloud services discussed above means that a lack of transferable skills can restrict customers' ability to switch and multi-cloud. Currently, the skills required to operate individual clouds are sufficiently distinct such that introducing a new provider necessitates the retraining or hiring of staff which increases customers' costs. Customers face challenges, including the opportunity cost, of both retraining their existing technical staff and recruiting new engineers to work on different clouds.
- 6.243 The impact of skills on customers depends on the nature of the workloads and services that specific engineers utilise. Customers running complex workloads or workloads using less cloud-agnostic software are more likely to find that their engineers have to build a skillset more focused on a specific provider and thus will be less able to deploy this skillset in another cloud.
- 6.244 We also found that, because providers offer broadly similar services, some engineers can find it easier than others to use alternative clouds and that training offered by cloud providers can aid this. However, customers provided evidence of the difficulties they face arising from a lack of transferable skills despite the training currently offered by cloud providers.
- 6.245 Therefore, whilst cloud providers do offer training to assist customers in upskilling on, or migrating to, their clouds and some customers have spoken of the benefits of this, this is not sufficient to overcome the effect of the lack of transferable skills on switching and multi-cloud.

Transparency

- 6.246 We have considered whether cloud providers make technical information to help customers overcome technical barriers when switching clouds or using multiple public clouds available and discoverable.

Cloud provider views

- 6.247 Cloud providers told us that they put significant effort into ensuring that customers are well-informed and that they supply sufficient documentation related to their services. For example:

¹⁴⁶⁸ Note of meeting with [REDACTED].

- (a) AWS said that it takes many active steps to inform and educate its customers, for example on the programming languages behind various tools that can be used to build on AWS, on the differences between its managed open-source services and the underlying open-source project,¹⁴⁶⁹ and on the use of containers to support seamless multi-cloud.¹⁴⁷⁰ AWS also said that it regularly publishes blog posts dedicated to such topics and provides free courses and guidance on how to move workloads to or from AWS through the AWS Migration Acceleration Program.¹⁴⁷¹
- (b) Microsoft said it educates customers how to switch or integrate public clouds. It said it publishes information and training on its website, including information to developers about the services available in Azure and how they can access their functionality. It also said it runs virtual and in-person training sessions.¹⁴⁷²
- (c) Google said that it does not believe lack of transparency to be a meaningful barrier. It said that it publishes extensive documentation for all its services.¹⁴⁷³ Google offers public training and guidance to switch between cloud services or pursue a multi-cloud strategy with Google.¹⁴⁷⁴

Customer views

- 6.248 Customers indicated that in some cases, there is an insufficient degree of transparency in the technical documentation available for cloud services, which can increase the costs or efforts associated with switching or multi-cloud.¹⁴⁷⁵ This includes instances of customers only becoming aware of issues after integration with another provider and customers not being aware of which features of particular services are proprietary to a specific provider.¹⁴⁷⁶ Another customer also said that some providers are more transparent than others and where transparency is lower, this can make integration more challenging.¹⁴⁷⁷
- 6.249 In contrast, some customers said that the available documentation is usually sufficient.¹⁴⁷⁸

¹⁴⁶⁹ [AWS' response to the CMA's Technical Barriers Working Paper](#) dated 06 June 2024, paragraph 53.

¹⁴⁷⁰ AWS' submission to the CMA [§].

¹⁴⁷¹ AWS' response to the CMA's information request [§]; AWS' submissions to CMA [§].

¹⁴⁷² Microsoft's response to the CMA's information request [§].

¹⁴⁷³ [Google's response to the CMA's Technical barriers working paper](#) dated 06 June 2024, paragraph 26(c).

¹⁴⁷⁴ Google's response to the CMA's information request [§].

¹⁴⁷⁵ Responses to the CMA's information requests [§].

¹⁴⁷⁶ Responses to the CMA's information requests [§].

¹⁴⁷⁷ [§] response to the CMA's information request [§].

¹⁴⁷⁸ Responses to the CMA's information requests [§].

Other stakeholder views

- 6.250 Accenture, a supplier of professional services, said that cloud providers generally offer documentation to help customers understand various aspects of their services, including lock-in risks and exit strategies. However, the depth and clarity of this documentation can vary. Furthermore, whilst providers generally offer high level guidance on data export, multi-cloud and migration, they lack detailed customer-specific exit plans.¹⁴⁷⁹
- 6.251 This is consistent with market research which found that subtle differences in capabilities or configuration which may exist below the threshold of awareness can be frustrating for customers when portability is attempted.¹⁴⁸⁰
- 6.252 Accenture also said that awareness among customers about the cost, effort and risks of switching and managing multi-cloud projects varies and is often insufficient.¹⁴⁸¹

Conclusion on transparency

- 6.253 We found that a lack of transparency in the technical documentation for cloud services acts as a barrier for some customers that want to switch or multi-cloud. Whilst providers make available a large amount of documentation related to their services, some customers find the information to be insufficient to inform them about the potential technical barriers they may face when switching and/or using multiple clouds.

Technical mitigations

- 6.254 In this section we discuss ways in which technical barriers to switching and multi-cloud faced by customers could be mitigated and their impact lessened. We have considered how far the market itself is responding to these barriers and have examined the efforts made by providers, ISVs and third parties, the open-source community and customers to develop mitigations that address the current technical barriers to switching and multi-cloud.
- 6.255 We have categorised customer-led technical mitigations into the following broad categories:
- (a) Infrastructure as Code (IaC) is the use of high-level descriptive coding language to automate and standardise the provisioning and deployment of IT infrastructure such as networks, virtual machines, load balancers and

¹⁴⁷⁹ Accenture's response to the CMA's information request [redacted].

¹⁴⁸⁰ [redacted] response to the CMA's information request [redacted].

¹⁴⁸¹ Accenture's response to the CMA's information request [redacted].

connection topologies required by any application. Customers can procure IaC tools from a third party or develop them internally.

- (b) Open-source software can be used by customers seeking to overcome the differentiation of features and interfaces between different cloud services. Customers typically have to procure and manage open-source software on their own. Customers may also contract a supplier of professional services to help with the procurement and management of open-source software.
- (c) Containers are a software tool that enable the packaging and isolation of applications with their entire runtime environment – all of the files necessary to run. Kubernetes is an open-source software for automating deployment, scaling and management of containerised applications.¹⁴⁸²
- (d) Cloud-agnostic services and tools aim to provide a consistent set of capabilities and interfaces across different clouds.
- (e) Adaptors are small, focused pieces of software that facilitate communication between two or more components that cannot directly interoperate with each other.

6.256 The EU Data Act includes provisions which allow the EU to introduce standards that require cloud providers to ensure compatibility with common specifications based on open interoperability specifications or harmonised standards for interoperability within 12 months of a standard being introduced. On 1 July 2025, the European Commission took an implementing decision on a standardisation request to the European standardisation organisations.¹⁴⁸³ The European Standards Organisations have officially agreed to develop standards to support the implementation of the EU Data Act and will work to publish a framework addressing cloud service provider switching mechanisms and cloud service interoperability to which organisations can refer when seeking to interpret the requirements of the EU Data Act.¹⁴⁸⁴

6.257 It is too early to assess the types of standards that may be required but we expect that more details will emerge following the entry into force of these provisions in September 2025. It is also not clear at the present time to what extent any standards implemented in response to the EU Data Act would impact cloud services in the UK. We discuss this further in our assessment of potential interventions in appendix W.

¹⁴⁸² Containers are an example of a cloud-agnostic architecture that has become popular with customers. These lightweight packages of application code house the dependencies required to run software services in the cloud, such as specific versions of the programming language runtimes and libraries.

¹⁴⁸³ [Register of Commission Documents - C\(2025\)4135](#).

¹⁴⁸⁴ [Data Act: Standardization Request Officially Accepted by CEN and CENELEC - CEN-CENELEC](#).

laC

Cloud provider views

- 6.258 AWS said that the fact that customers need to do some extra work to enjoy the benefits of abstraction layers, such as Terraform, does not mean that these tools do not effectively support efficient multi-cloud and switching. AWS said that learning Terraform to perform multi-cloud deployments is no different than learning one of the other laC tools needed to perform single cloud deployments. It said that, even with the time needed to learn laC tools like Terraform, these tools can ultimately save customers time and costs with respect to deploying and versioning their cloud infrastructure configuration changes. AWS also said that, even though laC tools like Terraform can become integrated into customers' development pipelines, adding steps to transition from one laC tool to another for certain use cases, they can ultimately save customers time and costs and promote and facilitate multi-clouding.¹⁴⁸⁵
- 6.259 [REDACTED].¹⁴⁸⁶
- 6.260 We also understand that some cloud providers have made investments to support the growth of laC services and tools [REDACTED].¹⁴⁸⁷
- 6.261 On the other hand, Civo said that an laC tool can abstract some but not all differences between clouds.¹⁴⁸⁸

Customer and other stakeholder views

- 6.262 We heard from a non-profit organisation that laC tools like Terraform have emerged to overcome the absence of standardised cloud provider APIs.¹⁴⁸⁹
- 6.263 However, we also heard from an ISV that provides laC tools that an laC tool can abstract some but not all differences between clouds. It added that some cloud services like storage are more easily abstracted than others like security and that laC tools only support the provisioning of services and not their runtime.¹⁴⁹⁰
- 6.264 Some customers and the Jigsaw report said that laC platforms enable multi-cloud deployments by standardising the deployment of infrastructure across cloud providers.¹⁴⁹¹ Additionally, some customers and ISVs and a provider of

¹⁴⁸⁵ AWS' submission to the CMA [REDACTED].

¹⁴⁸⁶ [REDACTED] response to the CMA's information request [REDACTED].

¹⁴⁸⁷ Responses to the CMA's information requests to [REDACTED]; Transcript of hearing with [REDACTED].

¹⁴⁸⁸ Note of meeting with Civo [REDACTED].

¹⁴⁸⁹ Note of meeting with [REDACTED].

¹⁴⁹⁰ Note of meeting with [REDACTED].

¹⁴⁹¹ Notes of meetings with [REDACTED]; Responses to the CMA's information requests [REDACTED]; [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), paragraph 4.4.7.

professional services said IaC tools improve a customer's ability to manage their multi-cloud architecture.¹⁴⁹²

6.265 However, the Jigsaw report and other market research stated that IaC tools can introduce their own lock-in.¹⁴⁹³ Additionally, some customers, an ISV and the Jigsaw report said that IaC tools can lead to higher costs, such as cost of acquiring a cloud-agnostic IaC tool, hiring and training staff and increased operational rework.¹⁴⁹⁴

6.266 We heard from an ISV that provides IaC tools that IaC tools cannot support seamless switching but can support multi-cloud with 80% of the differences abstracted.¹⁴⁹⁵ Additionally, it said that the IaC vendor lock-in can be overcome,¹⁴⁹⁶ but competing with IaC offered by cloud providers is harder since those tools are offered for free to customers.¹⁴⁹⁷

Open-source software

Cloud provider views

6.267 Cloud providers support open-source software in multiple ways, including by (i) offering managed (PaaS) services based on open-source software, (ii) enabling customers to manage open-source software themselves using IaaS services and (iii) using open-source software or standards within their services. Cloud providers may also invest in projects developing open-source software.

6.268 In general, cloud providers said that they support open-source software:

- (a) Microsoft and AWS said they support open-source software that facilitates switching and provided the examples of Linux and Kubernetes.¹⁴⁹⁸
- (b) AWS said that it is incentivised to support, and does support, the development of open-source software. [REDACTED]. AWS also said that it is an active member of and contributor to the PostgreSQL community and sits on the board of the MariaDB foundation.¹⁴⁹⁹

¹⁴⁹² Notes of meetings with [REDACTED]; [REDACTED] response to the CMA's information request [REDACTED].

¹⁴⁹³ Responses to the CMA's information requests [REDACTED]; [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), paragraph 4.4.17.

¹⁴⁹⁴ Notes of meetings with [REDACTED]; [REDACTED] response to the CMA's information request [REDACTED]; [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), paragraph 4.4.17.

¹⁴⁹⁵ Note of meeting with [REDACTED].

¹⁴⁹⁶ Note of meeting with [REDACTED].

¹⁴⁹⁷ Note of meeting with [REDACTED].

¹⁴⁹⁸ [AWS' response to the Technical barriers working paper](#) dated 06 June 2024, paragraph 19(a); and Microsoft's response to the CMA's information request [REDACTED].

¹⁴⁹⁹ AWS' submissions to the CMA [REDACTED].

- (c) Microsoft also said that the open-source community was a significant disciplining force against proprietary moats and customers are likely to leave a cloud provider if it does not offer popular open-source software.¹⁵⁰⁰
- (d) Google said its cloud supports a variety of open standards, protocols and compatibility with open-source frameworks.¹⁵⁰¹
- (e) IBM and Oracle said they support open-source software.¹⁵⁰²

6.269 AWS said that while trade-offs in deciding whether to use IaaS and open-source services instead of proprietary PaaS alternatives and challenges of adopting open-source technologies compared to managed open-source software might exist, customers have the choice to use IaaS and open-source software if they want to and so there cannot be any customer harm.¹⁵⁰³ As an example, it said that it tends to see that even customers who at the start of their cloud journey may only use AWS, tend to opt for open-source solutions such as EKS, which keep those customers' options open and facilitate switching and multi-cloud if they choose to in the future.¹⁵⁰⁴

6.270 Cloud providers also explained how they support open-source software and standard protocols within their services.

- (a) AWS said that its services support various standard protocols to make it easier for customers to enable communications and interactions between services – whether all services are on AWS or not – in a common way. For example, AWS said that its foundational compute and database services run on a range of open-source software and third-party software which can be used on other cloud environments.¹⁵⁰⁵ AWS also said it makes many of its APIs and SDKs publicly available under open-source licences,¹⁵⁰⁶ and that these allow customers and third parties to build compatible software and solutions, encouraging interoperability by allowing third-party IT providers to create services and solutions that work with AWS and making it easier for customers to move workloads away from AWS.¹⁵⁰⁷
- (b) AWS also said that in 2024 it added Apache Iceberg table support for S3 (where AWS is adopting a de facto open-source standard) to enable interoperability with Snowflake and Databricks.¹⁵⁰⁸

¹⁵⁰⁰ [Microsoft's response to the Technical barriers working paper](#) dated 06 June 2024, paragraph 35.

¹⁵⁰¹ Google's response to the CMA's information request [X].

¹⁵⁰² IBM's response to the CMA's information request [X]; Oracle's response to the CMA's information request [X].

¹⁵⁰³ [AWS' response to the Technical barriers working paper](#) dated 06 June 2024, paragraph 64.

¹⁵⁰⁴ AWS' submission to CMA [X].

¹⁵⁰⁵ AWS' response to the CMA's information request under [X].

¹⁵⁰⁶ AWS' submission to CMA [X].

¹⁵⁰⁷ [AWS' response to the Technical barriers working paper](#) dated 06 June 2024, paragraph 66(a).

¹⁵⁰⁸ AWS' submission to CMA [X].

- (c) Microsoft said it ensures its services are compatible and interoperable with other public cloud infrastructure services, by following industry standards and best practices, such as using open-source software.¹⁵⁰⁹ Microsoft said it helps with multi-cloud and switching by using common protocols and formats and providing APIs and SDKs for developers.¹⁵¹⁰
- (d) Google said its APIs are openly accessible to all and its SDK is available to developers;¹⁵¹¹
- (e) IBM and Oracle said they also provide API and SDK access to third parties.¹⁵¹²
- (f) A cloud provider's internal documents say that its customers perceive that tools built on open-standard APIs reduce the risk of changes which could limit use or raise costs for customers but the fact that these APIs are typically available as open-source software puts the burden on customers to select, integrate and manage a complex set of APIs to achieve their goals.¹⁵¹³

6.271 Microsoft, AWS and Google said that they make investments in open-source projects.¹⁵¹⁴

- (a) AWS said that it helped create and launch the permissively-licensed open source software like OpenSearch and Valkey, after the maintainers of the original projects 'went proprietary'.¹⁵¹⁵ It added that Valkey is 'progressing faster' than its original counterpart, Redis.¹⁵¹⁶
- (b) AWS also said that it is a major contributor to the open-source OpenTelemetry.¹⁵¹⁷

¹⁵⁰⁹ Microsoft's response to the CMA's information request [REDACTED].

¹⁵¹⁰ Microsoft's response to the CMA's information request [REDACTED].

¹⁵¹¹ Google's response to the CMA's information request [REDACTED].

¹⁵¹² IBM's response to the CMA's information request [REDACTED]; Oracle's response to the CMA's information request [REDACTED].

¹⁵¹³ [REDACTED] response to the CMA's information request [REDACTED].

¹⁵¹⁴ [Microsoft's response to the Technical barriers working paper dated 06 June 2024](#), paragraph 32 and [Google's response to the Technical barriers working paper dated 06 June 2024](#), paragraph 6, and [REDACTED].

¹⁵¹⁵ AWS' submission to CMA [REDACTED].

¹⁵¹⁶ AWS' submission to CMA [REDACTED].

¹⁵¹⁷ AWS' submission to CMA [REDACTED].

- (c) Microsoft highlighted investments in tools like Radius,¹⁵¹⁸ Paraglider,¹⁵¹⁹ and Dapr¹⁵²⁰ that it said support application portability and multi-cloud network management.¹⁵²¹
- (d) Google highlighted its development of Kubernetes and Tensorflow.¹⁵²²
- (e) Google also said that conversely, Microsoft and AWS do not support and promote open standards and protocols to the same extent as Google Cloud. As an example, it said that Google offers a free 'Assured Open Source Software' service that helps developers obtain open source software packages from trusted and known suppliers. Google also said that while Microsoft describes itself as a 'major contributor' to open source efforts, it does not identify any examples of open-source technology that Microsoft has itself developed and offered to the community.¹⁵²³

Customer and other stakeholder views

- 6.272 Two ISVs, Blue Prism and Splunk, and a customer, Vodafone, said the use of open-source software improves the scope for portability across clouds.¹⁵²⁴ One ISV mentioned limiting the number of proprietary services it used to build its product to increase the scope of portability of the product.¹⁵²⁵ Similarly, two customers said that using generic IaaS increased their ability to switch between different clouds.¹⁵²⁶
- 6.273 However, some customers highlighted the trade-offs customers face when deciding whether to use IaaS and open-source services instead of proprietary PaaS alternatives, such as limited access to innovations, difficulty scaling open-source solutions and potential financial loss due to loss of access to critical features.¹⁵²⁷ Additionally, a supplier of professional services said that while working with open-source software, customers should understand that they are

¹⁵¹⁸ Microsoft said that Radius abstracts away some of the effort required to run a self-managed platform for deploying applications. It provides consistent tooling and interfaces across the infrastructure it can be run on (currently Azure, AWS or on-premises). See ['The Microsoft Azure Incubations Team launches Radius, a new open application platform for the cloud – Microsoft Azure blog'](#) and ['Radius project - GitHub'](#), both accessed on 9 June 2025.

¹⁵¹⁹ Microsoft said that Paraglider is a control plane for cloud networking designed to simplify multi-cloud networking. See [Paraglider project released as open source to simplify networking within and across clouds - Microsoft Community Hub](#) and [Paraglider documentation](#).

¹⁵²⁰ Microsoft said that Dapr is a portable, event-driven runtime designed for microservice applications. It may increase the portability of code between VMs, Kubernetes, and serverless containers on different providers, but its main aim looks to be formalising and simplifying microservice development. See ['Dapr - Microsoft Community Hub'](#) and ['Dapr - Distributed Application Runtime'](#), both accessed on 9 June 2025.

¹⁵²¹ [Microsoft response to the Technical barriers working paper](#), paragraph 32.

¹⁵²² [Google's response to the Technical barriers working paper dated 06 June 2024](#), paragraph 6. We consider Kubernetes to fall under our definition of customer led when it is used by the customer as self-managed software and provider led when it is provided as managed service by a cloud provider.

¹⁵²³ Google's response to the CMA's information request [§].

¹⁵²⁴ Notes of meetings with Blue Prism [§]; Splunk [§]; Vodafone [§].

¹⁵²⁵ [§] submission to the CMA [§].

¹⁵²⁶ Responses to the CMA's information requests [§].

¹⁵²⁷ Responses to the CMA's information requests [§] Notes of meetings with [§].

trading service portability for continuous management and support and that they must rely on a community for understanding and using the tool.¹⁵²⁸

- 6.274 We heard from a non-profit organisation that cloud providers make their managed open-source services interoperable between platforms.¹⁵²⁹ It also said that the large engineering costs associated with forking OSS is a major deterrent to such practices.¹⁵³⁰ Separately, it noted that while providers do make their SDKs open, some are seen to be overly complicated making it difficult for the community to contribute to provider projects.¹⁵³¹
- 6.275 KPMG said customers typically use PaaS for their high-value workloads since it is more economically viable than for low value workloads.¹⁵³² Separately, another supplier of professional services (Accenture) said there is no satisfactory answer to workload portability while implementing/using proprietary services as they provide technical and potential pricing advantages.¹⁵³³

Containers and Kubernetes

Cloud provider views

- 6.276 AWS said that due to inherent technical barriers customers will always need to change certain configurations for each IT services provider they use, but they can run containers everywhere they run their software. It said that containers and Kubernetes do not lock in customers, and containers do not reduce the breadth of functionality or make scaling more difficult. The use of containerised applications can be used to effectively address concerns about the amount of recoding required to recreate or integrate applications across different IT environments when switching or integrating multiple clouds.¹⁵³⁴
- 6.277 An internal document from another cloud provider said containerised workloads have a greater 'attach power' and managed workloads are more sticky. This document suggested that this stickiness is associated with managed container services (like ECS, EKS, and GKE Enterprise). This document said that some customers are opting to use such services instead of building their own.¹⁵³⁵
- 6.278 On the 'stickiness' of container workloads, AWS said that this is limited to those customers that used managed container services like ECS and EKS, for which the

¹⁵²⁸ Note of meeting with [REDACTED].

¹⁵²⁹ [REDACTED] submission to the CMA [REDACTED].

¹⁵³⁰ [REDACTED] submission to the CMA [REDACTED].

¹⁵³¹ [REDACTED] submission to the CMA [REDACTED].

¹⁵³² Note of meeting with [REDACTED].

¹⁵³³ [REDACTED] response to the CMA's information request [REDACTED].

¹⁵³⁴ AWS' submission to the CMA [REDACTED].

¹⁵³⁵ [REDACTED] response to the CMA's information request [REDACTED].

customers were compensated with management and orchestration capabilities.¹⁵³⁶
AWS said that [REDACTED].¹⁵³⁷

6.279 AWS also said [REDACTED].¹⁵³⁸

6.280 Internal documents from two cloud providers state that customers are using Kubernetes to standardise and modernise their workloads and are leveraging it for AI/ML workloads to avoid vendor lock-in to proprietary AI models and access openly available AI foundation models.¹⁵³⁹

6.281 We understand that Microsoft, AWS and Google have made investments to improve the quality of their container and Kubernetes offerings.¹⁵⁴⁰

Customer and other stakeholder views

6.282 Some customers said that containers improve the portability of their workloads¹⁵⁴¹ with one customer saying it helped them manage applications across multiple clouds.¹⁵⁴² Additionally, Accenture said that at the application layer, the only 'true' way to implement and ensure cloud switching or multi-cloud portability is through containerising the app and running on a container orchestration platform.¹⁵⁴³

6.283 The Jigsaw report found that there is an increasing use of containerisation software, and some customers use such software to improve portability (it noted that not all customers use containers specifically for portability reasons but the portability benefits were widely recognised).¹⁵⁴⁴

6.284 Some customers told us that Kubernetes helped streamline application development for portability and provided open APIs that support workload portability.¹⁵⁴⁵

6.285 However, some customers, an ISV and a supplier of professional services said that containers cannot effectively overcome the technical barriers to switching and multi-cloud.¹⁵⁴⁶ Consistent with this, a report from a market research firm found that containers cannot significantly improve portability.¹⁵⁴⁷

¹⁵³⁶ [AWS' response to the Technical barriers working paper](#) dated 06 June 2024, paragraphs 66a and 66b.

¹⁵³⁷ AWS' submission to the CMA [REDACTED].

¹⁵³⁸ AWS' submission to the CMA [REDACTED].

¹⁵³⁹ Responses to the CMA's information requests [REDACTED].

¹⁵⁴⁰ AWS' response to the CMA's information request [REDACTED]; Transcript of hearing with Microsoft [REDACTED] Google's response to the CMA's information request [REDACTED].

¹⁵⁴¹ Responses to the CMA's information requests [REDACTED]; Note of meeting with [REDACTED].

¹⁵⁴² [REDACTED] response to the CMA's information request [REDACTED].

¹⁵⁴³ Accenture's response to the CMA's information request [REDACTED]. Container orchestration platforms provide services and tools for automating container orchestration.

¹⁵⁴⁴ [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), paragraph 4.4.6.

¹⁵⁴⁵ Responses to the CMA's information requests [REDACTED].

¹⁵⁴⁶ Notes of meetings with [REDACTED].

¹⁵⁴⁷ [REDACTED] response to the CMA's information request [REDACTED].

- 6.286 KPMG said that, even when customers move containerised workloads that run on IaaS, they may still need to do a small amount of refactoring, particularly in relation to integrations with those containers. However, this compares favourably to a much larger overhaul when customers are ‘moving a workload out of a container and into a PaaS service’.¹⁵⁴⁸
- 6.287 Similarly, a report from a market research firm found that containers only address technical portability barriers in a limited way – masking some differences from developers, but not from operators – and Kubernetes represents an alternative point of lock-in. It found that containers have limited impact on the other elements of portability and do not address the lock-in created by adoption of PaaS, management capabilities or other higher-level services.¹⁵⁴⁹
- 6.288 Separately, a customer, Asda, said that using Kubernetes increases the management costs and requires significant upfront investment.¹⁵⁵⁰

Cloud-agnostic services and tools

Cloud provider views

- 6.289 Microsoft submitted that use of cross-cloud applications to centralise data across various clouds and on-premises environments, such as Databricks Data Lakehouse, Snowflake, and Apache Airflow, allows for more effective multi-cloud architectures.¹⁵⁵¹

Customer and other stakeholder views

- 6.290 The Jigsaw report found that a well-known technical enabler to reduce lock-in is to use cloud-agnostic services and tools. In particular, it found that some customers considered using such services and tools instead of proprietary PaaS services to be an important way to reduce lock-in.¹⁵⁵²
- 6.291 We heard from a customer that there is a cloud-agnostic tool for every type of tool it needs,¹⁵⁵³ while an ISV said it uses a cloud-agnostic billing tool, Flexera, to manage its billing processes across the three major cloud providers.¹⁵⁵⁴

¹⁵⁴⁸ Note of meeting with KPMG [REDACTED].

¹⁵⁴⁹ [REDACTED] response to the CMA's information request [REDACTED].

¹⁵⁵⁰ ASDA's response to the CMA's information request [REDACTED].

¹⁵⁵¹ [REDACTED].

¹⁵⁵² [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), paragraph 4.4.4.

¹⁵⁵³ Note of meeting with [REDACTED].

¹⁵⁵⁴ Note of meeting with [REDACTED].

- 6.292 BT Group said that its Global Fabric ‘Network as a Service’ service makes it easier for customers to move between cloud providers by being pre-connected and offering customers unified control over connections to multiple cloud providers.¹⁵⁵⁵
- 6.293 A supplier of professional services, Capgemini, said that most customers have adopted a position to use a combination of cloud-agnostic services and proprietary tools, where they may use a native Kubernetes and a provider’s managed services.¹⁵⁵⁶
- 6.294 However, a customer and a supplier of professional services said that deploying cloud-agnostic tools increases the costs associated with such tools, forcing customers to choose between service portability and functionality.¹⁵⁵⁷ For example, a customer said that for a hypothetical application workload that costs £1 million to develop, it might cost £200,000 to ensure that the workload could be easily switched between public clouds.¹⁵⁵⁸
- 6.295 Additionally, an ISV said that cloud-agnostic tools are not as integrated into a provider’s cloud as proprietary tools are,¹⁵⁵⁹ while market research found that cloud agnostic tools will rarely have the same breadth and depth of functionality across all clouds.¹⁵⁶⁰ Similarly, a supplier of professional services and a non-profit organisation said that using cloud-agnostic ISVs can still lead to vendor lock-in.¹⁵⁶¹
- 6.296 A party said that the scope for unliteral effectiveness of tools used to mitigated technical barriers is limited. [X]. [X] said that, [X], these tools need to interoperate with various functionalities of the relevant cloud service providers, including APIs. [X].¹⁵⁶²
- 6.297 Similarly, the Jigsaw report also found that:
- (a) some customers are worried that by using cloud-agnostic alternatives, they risk losing access to innovation;¹⁵⁶³
 - (b) there are potentially long-term financial implications associated with avoiding proprietary services, and some customers fear that by avoiding proprietary services it is more difficult to achieve committed spend discounts;¹⁵⁶⁴
 - (c) not all organisations possess the requisite engineering expertise to effectively use cloud-agnostic alternatives, and the increased effort that some

¹⁵⁵⁵ [X] response to the CMA’s information request [X].

¹⁵⁵⁶ Note of meeting with Capgemini [X].

¹⁵⁵⁷ Responses to the CMA’s information requests [X].

¹⁵⁵⁸ Note of meeting with [X].

¹⁵⁵⁹ Note of meeting with [X].

¹⁵⁶⁰ [X] response to the CMA’s information request [X].

¹⁵⁶¹ Notes of meetings with [X].

¹⁵⁶² [X] submission to the CMA [X].

¹⁵⁶³ [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), paragraph 4.4.13.

¹⁵⁶⁴ [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), paragraph 4.4.15.

customers associated with using such alternatives means that proprietary services can be more appealing (especially to start-ups).¹⁵⁶⁵

Adaptors

Cloud provider views

- 6.298 AWS said that customers can procure adaptors from cloud providers, often for free or as part of the cloud providers' other services. It said it offers over 200 adaptors to customers so they can connect to data and applications that reside in other clouds, SaaS applications and on-premises solutions. AWS also said that, in any case, the fact that customers bear the additional cost of developing or purchasing an adaptor does not mean that adaptors do not effectively support efficient multi-cloud and switching.¹⁵⁶⁶
- 6.299 AWS also said that since APIs cannot be patented due to the presence of substantial prior art in the field, it was unlikely that method patents were deterring developers developing adaptors for APIs. Additionally, it said that all its APIs were available under the Apache licence, making it open to all.¹⁵⁶⁷

Customer and other stakeholder views

- 6.300 A non-profit organisation said that due to proprietary APIs that may be covered by method patents, developers are wary of creating adaptors that work with these APIs.¹⁵⁶⁸
- 6.301 A customer said it had deployed an internal network tool to act as a bridge across its cloud providers but said that 'even that comes at a cost'.¹⁵⁶⁹

Services that facilitate switching and multi-cloud

Cloud provider views

- 6.302 Cloud providers provided examples of the services that they have developed which either specifically aim to assist customers in multi-clouding or switching, or may have auxiliary benefits to customers' ability to multi-cloud or switch.
- 6.303 AWS said that it has invested in solutions that allow customers to switch, if they wish to do so, such as AWS Application Migration Services, AWS Database

¹⁵⁶⁵ [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), paragraph 4.4.16.

¹⁵⁶⁶ [AWS' response to the Technical barriers working paper](#) dated 06 June 2024, paragraph 57.

¹⁵⁶⁷ [AWS' response to the Technical barriers working paper](#) dated 06 June 2024, paragraph 59.

¹⁵⁶⁸ Note of meeting with [X].

¹⁵⁶⁹ [X] response to the CMA's information request [X].

Migration services, AWS Direct Connect, AWS Site-to-Site VPN, Kubernetes, among many others.¹⁵⁷⁰

- 6.304 AWS also said that it has made significant investments in other services and tools which are designed for specific use cases – and are not intended to ‘solve’ switching or multi-clouding – but are designed in such a way that anticipates customers’ multi-cloud usage, and are very successful in doing so. These services and tools included Amazon Bedrock,¹⁵⁷¹ Amazon Q Business,¹⁵⁷² AWS Glue,¹⁵⁷³ AWS Outposts,¹⁵⁷⁴ and Amazon AppFlow,¹⁵⁷⁵ AWS DataSync,¹⁵⁷⁶ and certain connectors included in Amazon Athena.¹⁵⁷⁷
- 6.305 AWS also submitted that IT service providers have invested in automation and tools to enable easier migration, allowing transformation between virtualisation layers and between operating systems, amongst others, or supporting portable infrastructure-as code deployment.¹⁵⁷⁸ It also said that switching will likely become even easier with AI.¹⁵⁷⁹
- 6.306 In particular, AWS said that AI is rapidly affecting the ease by which customers can move workloads across different API boundaries. It said that it increasingly sees customers using Amazon Q Developer, Claude, or ChatGPT to redevelop software from one of AWS’ APIs to another provider’s API. It gave real-life examples of where Amazon Q Developer has significantly accelerated software migration projects.¹⁵⁸⁰

¹⁵⁷⁰ [AWS’ response to the Updated issues statement and working papers](#) dated 23 May 2024 and 6 June 2024, paragraph 9.

¹⁵⁷¹ AWS said that Amazon Bedrock makes available a collection of FMs, including Amazon’s own FMs and FMs from third-party model developers, through a common API. It said that customers opt to use Bedrock because of the additional security and privacy benefits they get from fully managed service, but if they want the ability to port their model, they can use Amazon SageMaker, which would allow a developer to start training a model on AWS, take the model to another IT services provider, and continue training on that provider. AWS’ submission to the CMA [X].

¹⁵⁷² AWS said that Amazon Q Business supports over 40 native integrations, 75% of which are to third-party services such as SharePoint, Google Drive, Salesforce, Confluence, and more. It said that in addition, a customer can recreate Q Business using AWS tools, such as Amazon Kendra, in conjunction with an external large language model AWS’ submission to the CMA [X].

¹⁵⁷³ AWS said that AWS Glue enables customers to perform data analytics, machine learning, and application development from AWS, other cloud sources, on-premises, and third-party SaaS sources. AWS’ submission to the CMA [X].

¹⁵⁷⁴ AWS said that AWS Outposts helps customers manage and govern multi-cloud deployments, but only for servers that are compatible. While Outposts itself is an extension of AWS cloud, workloads running on Outposts can interface with other cloud providers just as they would in the context of AWS Regions. AWS’ submission to the CMA [X].

¹⁵⁷⁵ AWS said that Amazon AppFlow regularly adds new data connectors to its repertoire, and can ingest data from third-party applications, enrich it through filters and validations, and transfer it to destinations both inside and outside of AWS. AWS’ submission to the CMA [X].

¹⁵⁷⁶ AWS said that AWS DataSync supports transfers between AWS and many other IT providers, which facilitates multi-cloud usage and customer migration in both directions. AWS’ submission to the CMA [X].

¹⁵⁷⁷ AWS said that Amazon Athena can run SQL queries on data stored in almost any format, wherever it resides. AWS said that it introduced a data source connector for Amazon Athena, allowing customers to query and surface data insights from over 25 external data sources, including relational, non-relational, object, and other cloud data stores, without copying or transforming data. AWS’ submission to the CMA [X].

¹⁵⁷⁸ [AWS response to the CMA’s Technical barriers working paper dated 06 June 2024](#), paragraph 5.

¹⁵⁷⁹ [Summary of hearing with AWS, 12 March 2025](#), paragraph 5.

¹⁵⁸⁰ AWS’ submission to the CMA [X].

- 6.307 AWS also said that agentic AI (an emerging technology) has the potential to (i) autonomously analyse a customer's existing workloads and generate a migration path to another provider, identifying service equivalents and configuration changes; (ii) autonomously reduce a customer's IT costs by reallocating workloads to different providers when they identify better offers; and (iii) interface with multiple cloud APIs.¹⁵⁸¹
- 6.308 Similarly, Microsoft said that AI-based solutions are making it easier for customers to switch and multi-cloud.¹⁵⁸²
- 6.309 Microsoft said some of its services are designed to facilitate interoperability between clouds and provided the example of Azure Arc.¹⁵⁸³ Azure Arc allows customers to manage resources hosted outside Azure across on-premises, hybrid and multi-cloud environments from the Azure Portal as though they were hosted on Azure. Microsoft also said its cloud service can integrate data from disparate data sources across cloud environments to be used with Azure services, such as AI.¹⁵⁸⁴ Microsoft said it has developed Azure Fabric¹⁵⁸⁵ to support multi-cloud querying and analytics.¹⁵⁸⁶
- 6.310 Google said it is developing services specifically to support easy integration. Google Cloud also offers solutions such as Dataproc, which allows customers to run open-source data analytics at scale and TensorFlow, a free and open-source software library for machine learning and AI tools.¹⁵⁸⁷ Additionally, Google has also worked with third parties like Apache Cassandra, Datadog and VMware to facilitate integrations with its own services.¹⁵⁸⁸
- 6.311 Oracle also said that it designs products to be interoperable across clouds and that one of the ways it supports multi-cloud is by offering Oracle Database@Azure where OCI provisions the Exadata Database Service within Azure. This means that a customer can access Oracle's Exadata Database Service within Azure just as well as within OCI, which provides the lowest latency for the customer.¹⁵⁸⁹
- 6.312 IBM said that it designs products to be interoperable across clouds and that its cloud offering is based on a hybrid multi-cloud approach, meaning that IBM

¹⁵⁸¹ AWS' submission to the CMA [§]. [§].

¹⁵⁸² Microsoft's submission to the CMA [§].

¹⁵⁸³ 'Azure Arc overview - Azure Arc - Microsoft Learn', accessed on 7 July 2025; Microsoft's response to the CMA's information request [§].

¹⁵⁸⁴ Microsoft's response to the CMA's information request [§].

¹⁵⁸⁵ 'What is Microsoft Fabric - Microsoft Fabric - Microsoft Learn', accessed on 7 July 2025.

¹⁵⁸⁶ AWS' response to the Updated issues statement and working papers dated 23 May 2024 and 6 June 2024, paragraphs 4 and 5; Microsoft's response to the CMA's information request [§].

¹⁵⁸⁷ Google's response to the Technical barriers working paper dated 06 June 2024, paragraph 6.

¹⁵⁸⁸ Google's response to the CMA's information request [§].

¹⁵⁸⁹ Oracle's response to the CMA's information request [§].

provides cloud-related services largely irrespective of the customers' choice of cloud service provider.¹⁵⁹⁰

Customer and other stakeholder views

- 6.313 Two customers identified specific services that may help facilitate switching and multi-cloud:
- (a) A customer said that it uses BigQuery Omni, which allows it to query AWS billing data from GCP, and had explored using Amazon Athena's Google Cloud Storage connector. It said that overall, the ability to query data held in another provider's systems is incredibly valuable, and makes evaluation of the services, and any potential migration, far easier.¹⁵⁹¹
 - (b) Another customer told us that Amazon Q Developer has been embraced as a core developer productivity tool, and accelerates its integration of AWS services.¹⁵⁹²

Conclusion on technical mitigations

- 6.314 The mitigations that customers most frequently mentioned, such as IaC, self-managed open-source software, and containers/Kubernetes, are typically developed by ISVs or the open-source community to support switching and multi-cloud. We view these mitigations as 'customer-led mitigations' (and refer to them as such below), in that customers shoulder most of the implementation effort and cost, which can affect customers' decisions to use them.
- 6.315 The evidence relating to the customer-led mitigations is mixed and in each case, there was some evidence to suggest that the mitigation does to an extent improve the ability of some customers to switch and multi-cloud. However, there was also evidence showing that these mitigations are not fully effective in doing so, for example because the costs outweighed the benefits, customers were locked into the ISVs that supply these mitigations, the breadth of functionality was reduced and scaling was more difficult.
- 6.316 The cloud providers submitted evidence relating to the customer-led mitigations, but also noted a range of other potential mitigations that were generally not mentioned by customers. We view these other mitigations as 'provider-led mitigations' (and refer to them as such below), meaning services and tools designed by cloud providers to facilitate switching and multi-cloud. Such services

¹⁵⁹⁰ IBM's response to the CMA's information request [§].

¹⁵⁹¹ Note of meeting with [§].

¹⁵⁹² [§] response to the CMA's information request [§].

and tools may emerge where cloud providers' incentives align with customers' needs.

6.317 We have examined the range of provider-led mitigations further. We found that some provider-led mitigations sufficiently addressed a particular and narrow problem, some seemed insufficient and for some their effectiveness is uncertain. For example, we found that:

- (a) AWS Glue, AWS AppFlow, and Microsoft Fabric allow customers to query and analyse data stored within other providers' IaaS, PaaS, and SaaS services but these solutions only address multi-cloud data analytics.¹⁵⁹³ Similarly, Azure Arc and AWS Outposts help customers manage and govern multi-cloud deployments for servers that are 'Arc-enabled'¹⁵⁹⁴ or compatible with AWS Outposts (which allows customers to run their workloads using AWS infrastructure on-premises).¹⁵⁹⁵ Overall, these services do not seem to address broader customer needs when it comes to multi-cloud.
- (b) while Amazon Bedrock allows customers to access multiple foundation models, it is still a proprietary cloud service and customers must modify their code if they switch to or from a competing service, such as OpenAI. We note, however, that Amazon provides guidance on how to use other AWS services to set up an API adapter.¹⁵⁹⁶
- (c) while Amazon Q Business supports many integrations, these seem to be mostly with third-party SaaS services rather than third-party cloud infrastructure services.¹⁵⁹⁷
- (d) tools like Dapr, Radius and Paraglider that support application portability and multi-cloud network design can potentially support switching and multi-cloud. However, these services are nascent and there is insufficient evidence to conclude on their practical effectiveness and usability.

6.318 We recognise the contributions of cloud providers to open source software projects and the fact that in some cases cloud providers have made source code publicly available, which facilitates interoperability.

6.319 Cloud providers, ISVs and open-source communities have developed some mitigations to the technical barriers to switching and multi-cloud. Many of these, including some that have emerged relatively recently (such as Microsoft Fabric),

¹⁵⁹³ 'What is AWS Glue? – AWS documentation', accessed on 19 May 2025; 'What is Amazon AppFlow? – AWS documentation', accessed on 19 May 2025; and 'Data Analytics Platform | Microsoft Fabric', accessed on 19 May 2025.

¹⁵⁹⁴ AWS explained that customers can register non-Azure virtual machines (servers) with Azure Arc by installing the Azure Connected Machine agent on each one, which simplifies enablement and data collection for supported Azure services (see [AWS' response to the Technical barriers working paper](#) dated 06 June 2024, paragraph 92).

¹⁵⁹⁵ [AWS' response to the Technical barriers working paper](#) dated 06 June 2024, paragraph 46.

¹⁵⁹⁶ 'bedrock-access-gateway, aws-samples – GitHub', accessed on 3 April 2025.

¹⁵⁹⁷ 'Supported connectors – Amazon Q User Guide', accessed on 3 April 2025.

facilitate interoperability in specific use cases. However, current mitigations, where one addresses a particular barrier (or instances of a particular barrier) at least to a limited extent, are not sufficient to mitigate the overall effects of technical barriers on switching and multi-cloud. Switching and multi-cloud remain technically challenging processes for customers even in light of these mitigations.

Conclusions on technical barriers

- 6.320 There are substantial technical barriers facing customers who wish to switch or use multiple clouds and the incentives for cloud providers to reduce these barriers are mixed.
- 6.321 Some customers can and do successfully multi-cloud but technical barriers to multi-cloud negatively affect many customers' ability to use and integrate multiple public clouds. This limits customers' ability and/or incentive to exercise choice of cloud provider.
- 6.322 Whilst customers are able to switch if they are willing to expend all the required resources to do so, the costs associated with overcoming these technical barriers currently outweigh the benefits for most and some would not switch even if they were forgoing benefits in failing to do so.
- 6.323 A greater ability for customers to switch and multi-cloud would improve competition, whether or not they choose to do so, because this ability would increase competitive pressure on providers.¹⁵⁹⁸ Our view of technical barriers does not rest on an assumption that all customers want to switch or multi-cloud or that they should be able to do so without any friction. Rather, our view is that if barriers were reduced, then the ability to switch and multi-cloud would improve and this would drive competition amongst providers.¹⁵⁹⁹
- 6.324 Technical differentiation of both features and interfaces in core services and ancillary services and tools means that products cannot be easily substituted and this harms customers' ability to switch and/or multi-cloud. Further disincentives to switching and multi-cloud include latency, a lack of transferable skills across clouds and insufficient transparency around potential technical barriers and how to mitigate or overcome them.
- 6.325 There are some mitigations to some of these technical barriers, including those from cloud providers and third parties and customers' own mitigations. Many of these, including some that have emerged over the course of our investigation, facilitate interoperability in specific use cases. However, the patchwork of

¹⁵⁹⁸ See the above section titled Providers' incentives to facilitate switching and multi-cloud.

¹⁵⁹⁹ See AEC chapter.

mitigations does not amount to a comprehensive solution to the barriers found and switching and multi-cloud remain technically challenging in many cases.

Egress fees


Introduction

- 6.326 We have considered whether, and to what extent, egress fees are a barrier to switching and multi-cloud and a feature that, in combination with other features (such as technical barriers), give rise to an AEC in cloud services in the UK. This section presents our analysis of the impact of egress fees on switching and multi-cloud and on competition among cloud providers.
- 6.327 We have set out in turn:
- (a) background information on egress fees in general;
 - (b) our conceptual framework for assessing the impact of egress fees on competition; and
 - (c) our analysis and views of:
 - (i) the relevance of egress fees in customers' switching and multi-cloud decisions; and
 - (ii) potential benefits in relation to egress fees.
- 6.328 We have focused on egress fees incurred during the switching and/or multi-cloud process. Where we use the term 'egress fees', we refer to fees charged to transfer data out of one cloud provider's cloud into another cloud provider's cloud, unless stated otherwise (ie 'egress fees in general').

Background

- 6.329 Cloud providers offer networking services to allow cloud customers to transfer data:
- (a) into a cloud provider's infrastructure ('ingress');
 - (b) within a cloud provider's infrastructure ('internal transfers'); and
 - (c) out of a cloud provider's infrastructure ('egress').
- 6.330 As set out in Figure 6.2, cloud providers charge for the use of most types of their data transfer services but cloud providers generally do not charge for data ingress.

Figure 6.2: Types of data transfers



*A content distribution network (CDN) is a geographically distributed network that delivers content to end users or applications.

Source: CMA analysis of Google's submission to Ofcom [8]. ¹⁶⁰⁰

6.331 Most cloud providers charge customers when transferring data within ('internal transfer fees') and out of a cloud provider's infrastructure. Egress fees in general (ie not just for switching and multi-cloud) are incurred when a customer transfers data out of a cloud provider's cloud.¹⁶⁰¹ This can occur when:

- customers move data between different cloud providers as part of a multi-cloud architecture or as part of switching;
- customers move their data from their cloud to their on-premises data centres; and
- customers deliver content to end users or applications, either directly or via content distribution networks ('CDN') (eg Cloudflare).

6.332 For the purposes of our assessment on whether egress fees constitute a barrier to switching and multi-cloud, we have focused on the impact of egress fees when customers move data between different cloud providers as part of a multi-cloud architecture or as part of switching (ie transfer type (a) above). In other words, where we use the term 'egress fees', we refer to fees charged to transfer data out of one cloud provider's cloud into another cloud provider's cloud. Where we use

¹⁶⁰⁰ Google explained that within the cloud industry, egress refers to any data transfer, whether that transfer occurs within a provider's cloud infrastructure, from one provider's cloud infrastructure to another's or from one provider's cloud infrastructure to an on-premises environment or end-customer. We adopt a different definition for the purposes of this report.

¹⁶⁰¹ The terminology and the definition of these fees are not consistently used by all cloud providers. For instance, AWS refer to egress fees as DTO (Data Transfer Out) fees. For consistency across providers, we will continue to refer to them as egress fees.

the term ‘egress fees in general’, we refer to all forms of egress fees (ie transfer types (a) to (c) above).

- 6.333 We refer to the wider set of fees charged by cloud providers for data transfer services (ie internal transfer fees and egress fees in general) as data transfer fees.
- 6.334 The cost of egressing data in general varies based on several factors. For example, egress fees’ pricing is generally volume-based and region-specific. Some cloud providers provide a monthly data egress ‘free tier’, within which a customer can transfer up to a certain amount of data without incurring a charge. Above this, there is a stepped per-GB pricing structure, with per-GB prices decreasing with higher volumes of data egress. This means that the amount paid for egress (and hence the significance of egress spend) for a customer depends on a number of different factors, such as a customer’s industry, whether it has regular or intermittent data traffic, the geographic regions data is egressed to, the customer’s IT architecture (including whether they also use other clouds and/or on-premises) and any applicable discounts the customer has negotiated.
- 6.335 For the purposes of our assessment of whether egress fees constitute a barrier to switching and multi-cloud, we have focused on the impact of egress fees when customers move data between different cloud providers as part of a multi-cloud architecture or as part of switching (ie transfer type (a) above). This is because customers are likely to consider these fees when switching cloud provider or using multiple clouds.
- 6.336 However, we have been told by more than one cloud provider that they cannot:¹⁶⁰²
- (a) identify the purpose of a customer’s data transfer out of their cloud infrastructure; or
 - (b) consistently identify the peer (ie the company and/or the relevant business unit/subsidiary within that company) to which the data is transferred and whether the peer is the end destination.

Conceptual framework

- 6.337 As set out at the start of this chapter, customers considering whether to switch or multi-cloud face a trade-off between the expected benefits and expected costs of doing so.
- 6.338 Egress fees represent a cost to customers both when switching and when using multiple clouds:

¹⁶⁰² Responses to the CMA’s information requests to [X].

- (a) when switching cloud provider, a customer will need to transfer data from one cloud to the other and will incur egress fees if the volume of data transferred is large enough (unless they are a customer of a cloud provider that offers a free switching programme and they meet the relevant cloud provider's eligibility criteria).
- (b) when using an integrated multi-cloud architecture, a customer will transfer data back and forth between public clouds and will incur ongoing egress fees over time if the volume of data transferred reaches a sufficient level.
- (c) in some cases, customers may incur both switching and multi-cloud costs, eg where a customer performs a partial switch (moves some of their workloads from one cloud to another) and then runs both clouds in parallel in a way that involves continued data transfers back and forth. In such a case, both costs would need to be considered.

6.339 Where the aggregate expected costs of switching and/or multi-cloud exceed the expected benefits, customers are likely to choose not to switch to or to multi-cloud, including when alternative cloud providers would otherwise have a better offer (in terms of price, quality, range of features/capabilities, or other considerations).

6.340 In considering whether egress fees act as a barrier to switching and multi-cloud we have assessed the following issues:

- (a) first, we have considered a range of evidence on the likely relevance to customers' decisions of the egress fees a customer would experience when switching or using multiple clouds. The costs of switching and/or multi-cloud are more likely to reduce customers' propensity to switch when they are large in magnitude. We have also considered the effect of free switching programmes on our analysis of switching costs.
- (b) second, we have considered whether there are potential benefits, such as efficiencies, which may remove or mitigate any competitive harm arising from egress fees both when switching or using multiple clouds. In particular, we have considered cloud providers' submissions that costs recovered from egress fees in general result in investment and innovation and lower prices for customers and/or avoid unnecessary network investment due to inefficient egress usage.

Relevance of egress fees to customers' switching and multi-cloud decisions

6.341 As set out at the start of this chapter, barriers to switching providers or barriers to using multiple providers may restrict customers from exercising effective choice. They directly affect customers' ability and incentives to choose the provider(s) that

best fits their needs. Switching from one provider to another, so as to respond to attractive offers, may be made difficult for customers by the costs of doing so.¹⁶⁰³

- 6.342 The higher the level of egress fees associated with transferring data between clouds, the higher the cost of switching and/or using a multi-cloud architecture becomes. This cost may further increase as a multi-cloud setup becomes more integrated and the levels of data transfer increases. Thus, these higher costs for customers as a result of higher egress fees can be expected to make customers less responsive to the efforts of rivals to compete by improving their offer.
- 6.343 The extent to which egress fees represent a barrier to both switching and/or multi-cloud depends on the amount of data a customer seeks to transfer and the frequency with which the data must be transferred. Data transfer follows different patterns between switching and multi-cloud:
- (a) in relation to switching, customers could move all the data they seek to transfer as a 'one-off' transfer.¹⁶⁰⁴ When transferring as a 'one-off', the impact of egress fees are incurred only during the switching process rather than on a recurring basis.
 - (b) in relation to the use of multiple clouds, the extent to which egress fees are a barrier to multi-cloud depends on the degree of integration of customers' multi-cloud architecture. Under a more siloed multi-cloud architecture, where workloads have low levels of interdependence, smaller and less frequent data transfer between clouds may imply lower egress fees. Under a more integrated multi-cloud architecture, with larger and more frequent data transfers across clouds, egress fees are likely to be larger.
- 6.344 In view of the above, we have considered the extent to which egress fees are relevant to customers' switching and/or multi-cloud decisions. In particular, we have considered evidence on:
- (a) cloud provider views on the relevance of egress fees, including their responses to our PDR. We have also reviewed cloud providers' internal documents relevant to egress fees;
 - (b) the cost to customers of egressing data through our analysis of switching and multi-cloud scenarios;
 - (c) customer views on the impact of egress fees on their decisions to switch and/or multi-cloud; and

¹⁶⁰³ CC3 (Revised), paragraphs 297 and 316.

¹⁶⁰⁴ This might involve a period of running two services in parallel with each other, which would contribute to the cost of switching, particularly if data must be egressed back and forth between those clouds.

- (d) the impact of the free switching programmes introduced by cloud providers in 2024 on our analysis.

6.345 We have focused much of our analysis on Microsoft, AWS and Google as they collectively represent the majority of revenue in the IaaS and PaaS markets. Whilst egress fees for other cloud providers vary,¹⁶⁰⁵ we have considered that the findings of our analysis would also apply more widely to any cloud provider that charges egress fees, as any level of egress fees can be expected to represent a barrier to switching and/or multi-cloud. However, the impact of any such barrier is likely to vary with the magnitude of the egress fees charged for switching and/or multi-cloud, the size of the provider and the scale of customer data transfers.

Cloud provider views

6.346 Cloud providers have submitted views on the relevance of egress fees. We have grouped these by theme below, along with our analysis of these submissions:

- (a) the observed prevalence of egress fees;
- (b) customers' levels of concern about egress fees;
- (c) trends in egress fee prices; and
- (d) justification for egress fees in the context of free ingress.

6.347 We also summarised our review of cloud providers' internal documents, and the views they contain in relation to egress fees.

6.348 The cloud providers' submissions on customers' levels of concern about egress fees were focused on egress fees associated with switching and multi-cloud (ie item (b) above). The cloud providers' internal documents and submissions relating to the observed prevalence of egress fees, to trends in egress fee prices, and to the relevance of free ingress (ie items (a), (c) and (d) above) tended to relate to egress fees in general (ie egress fees associated with any data transfers out of the cloud provider's cloud, rather than egress fees specifically associated with switching and/or multi-cloud use). Our analysis of items (a), (c) and (d) therefore referred to egress fees in general, except where switching and multi-cloud egress was specifically contemplated, ie item (b).

¹⁶⁰⁵ For example, IBM has egress prices which are similar to the largest cloud providers, whilst Oracle has a lower flat fee for data transfers, and UK cloud provider Civo has recently removed all egress fees. See appendix L; [Civo abolishes all egress fees to advance cloud simplicity strategy](#), accessed on 24 September 2024.

The observed prevalence of egress fees

- 6.349 Microsoft, AWS and Google told us the proportion of their customers who paid egress fees in general in 2023. These varied widely.
- (a) AWS submitted that [REDACTED] of UK customers with at least \$1,000 of annual spend paid egress fees in 2023. AWS submitted that this figure overestimates the share of UK customers paying egress fees because it excludes some customers who are more likely to fall into the AWS free tier and does not account for the fact that the AWS free tier is allocated to customers on a monthly basis.¹⁶⁰⁶
 - (b) A separate analysis by AWS which included customers that spent under \$1,000, estimated that approximately [REDACTED] of UK customers paid egress fees in 2023. AWS noted that this may still overestimate the share of UK customers paying DTO fees,¹⁶⁰⁷ because it excludes certain customers, namely customers [REDACTED] and customers where AWS has not been able to match [REDACTED].¹⁶⁰⁸
- 6.350 We note that the analysis in this submission includes customers who spent under \$1,000 on AWS services in 2023. These include very small customers which, overall, make up a small proportion of cloud spend on AWS and some of whom may be experimenting with AWS services. This submission also considers customers with a total data egress volume of under 1,200GB in 2023 as falling within AWS' free tier. This is based on the sum of all data covered by the free tier for each of the 12 months in 2023 (12/100GB). We think this approach underestimates the number of customers who pay egress fees in a year.
- 6.351 Microsoft submitted that [REDACTED] of its UK customers with at least \$1,000 of annual spend on first-party public cloud infrastructure services in 2023 paid non-zero egress fees in general, unweighted by yearly cloud spend. When weighted by cloud spend, this rose to [REDACTED].¹⁶⁰⁹
- 6.352 Google submitted that just over [REDACTED] of its UK customers with a cloud spend of over \$1,000 in 2023 paid egress fees in general, unweighted by yearly cloud spend. When weighted by cloud spend, this rose to [REDACTED]. Google noted that since data transfer fees are part of the ordinary course provision of cloud infrastructure services (with the majority of external data transfer volumes relating to ordinary

¹⁶⁰⁶ AWS' response to the CMA's information request [REDACTED]. AWS also submitted that over 90% of AWS' global customers which egress data fall within the free tier and therefore do not pay anything for egress fees. ([AWS' response to the CMA's PDR dated 28 January 2025](#), paragraph 51). We have focused on UK-specific analysis.

¹⁶⁰⁷ AWS refer to egress fees as DTO (Data Transfer Out) fees. For consistency across providers, we will continue to refer to them as egress fees.

¹⁶⁰⁸ AWS' response to the CMA's information request [REDACTED].

¹⁶⁰⁹ Microsoft's response to the CMA's information request [REDACTED].

course 'serving' egress), it is to be expected that most customers pay egress fees.¹⁶¹⁰

6.353 Two other stakeholders submitted that, in their experience, egress fees only impact some customers.¹⁶¹¹

6.354 The evidence above shows that egress fees are widespread. In considering the significance of egress fees as a barrier to switching and multi-cloud, we note that:

- (a) the level of spending on egress fees is likely influenced by the level of egress fees. For example, a customer may decide to egress less data than it otherwise would in order to avoid those egress fees. Thus, using the actual levels of spend on egress fees as evidence would risk underestimating the potential for egress fees to have an impact on customers' switching and multi-cloud decisions. In other words, if a customer is deterred from switching or multi-cloud because of egress fees, then they will not show up in the analyses submitted by the cloud providers above. As set out in chapter 3, we have found that full switching between public cloud providers is extremely rare in the market and whilst there is some multi-cloud, generally amongst larger customers, its overall prevalence in conjunction with other evidence (eg on customers' split of spend across cloud providers) indicates that customers' ability to integrate workloads on more than one cloud is subject to barriers. We therefore consider that the cloud providers' analysis will therefore tend to be biased in favour of customers whose behaviour has not been affected, or has been less affected, by egress fees.
- (b) the analyses presented by the cloud providers do not tell us about the observed prevalence of egress fees specifically for switching and/or multi-cloud. This is because cloud providers cannot reliably distinguish between the purpose of the data transfer.

6.355 This is why we have estimated what customers would have to pay in egress fees if they were to either move all their data to another cloud provider or adopt a multi-cloud setup. This allowed us to consider the significance of egress fees for switching and multi-cloud for all customers, rather than only for a subset of customers that have not been deterred by egress fees. These estimates are described in a later section.

6.356 Microsoft submitted that we should not dismiss data on the prevalence and spend by UK customers on egress usage. It said the PDR provided no evidence of

¹⁶¹⁰ Google's response to the CMA's information request [8].

¹⁶¹¹ Rayo, a cloud technology consulting company, submitted that egress fees affect only a small number of very large customers moving substantial amounts of data (typically over 100GB per month), with most businesses never reaching these thresholds. It also said that moving data in bulk can be done with offline devices. [Rayo's response to the CMA's PDR dated 28 January 2025](#), page 1. Startup Coalition also submitted that startups are rarely impacted by egress fees as currently structured in the market. [Startup Coalition's response to the CMA's PDR dated 28 January 2025](#), page 2.

customer avoidance behaviour and ignored that customers opt not to switch or multi-cloud for a number of more material reasons than egress fees.^{1612, 1613} AWS also submitted that there is a lack of evidence of such customer avoidance behaviour and that observational data is informative.¹⁶¹⁴

- 6.357 We agree that, in addition to egress fees, there are several factors affecting a customer's decision to switch or multi-cloud. This does not remove our concern that the level of egress fees affects the volume of data egressed by customers and therefore how much they spend on it. As set out below, we found that some customers, competitors and industry bodies identified egress fees as a challenge when considering, or a factor disincentivising, switching and/or multi-cloud. This suggests that egress fees contribute to some customers avoiding switching and/or multi-cloud, which will not appear in the data and in the analyses submitted by the cloud providers above.

Customers' levels of concern about egress fees

- 6.358 Google submitted that the majority of customer evidence in our Egress fees working paper and in the Jigsaw report showed that egress fees are not a meaningful barrier to switching.¹⁶¹⁵

- 6.359 AWS submitted:

- (a) that the evidence gathered on customer views indicated that egress fees are a minor consideration for customers thinking about switching and/or multi-cloud, rather than a material barrier. AWS added that 'only a few [customers] spontaneously identified egress fees [as a barrier to switching and/or multi-cloud]'.¹⁶¹⁶ AWS further submitted that survey evidence does not support the claim that egress fees distort competition,¹⁶¹⁷ nor support the conclusion that egress fees are significant for any "substantial minority of customers".¹⁶¹⁸ Also, AWS submitted that public responses to the PDR from key stakeholders do not support the PDR's conclusions on egress fees.¹⁶¹⁹

¹⁶¹² [Microsoft's response to the CMA's PDR dated 28 January 2025](#), paragraph 97.

¹⁶¹³ Microsoft submitted that when making decisions about switching or multi-clouding, customers are focused on the various services available and their inherent cost, plus the implicit costs of training, operations, time, and security. [Microsoft's response to the CMA's PDR dated 28 January 2025](#), paragraph 23.

¹⁶¹⁴ Transcript of hearing with AWS, 12 March 2025 [REDACTED]. AWS' submission to the CMA [REDACTED].

¹⁶¹⁵ [Google's response to Egress Fees Working Paper](#), paragraph 31. [Google's response to the CMA's PDR dated 28 January 2025](#), paragraph 25.

¹⁶¹⁶ [AWS' response to the CMA's updated issues statement and working papers](#), paragraph 35(d).

¹⁶¹⁷ AWS' submission to the CMA [REDACTED].

¹⁶¹⁸ AWS' submission to the CMA [REDACTED]. We note that AWS [REDACTED].

¹⁶¹⁹ AWS' submission to the CMA [REDACTED].

(b) that the PDR did not make clear the composition of the ‘substantial minority’ of customers affected by egress fees.¹⁶²⁰

(c) that the level of concern expressed by customers falls below the balance of probabilities standard to find an AEC,¹⁶²¹ and that the relevant question that the CMA should ask is whether switching and multi-cloud costs are substantial enough to prevent customers from adopting a multi-cloud approach or switching providers, should they see benefits to doing so.¹⁶²²

6.360 Similarly, Microsoft submitted that it was erroneous to consider it sufficient to find an AEC on the basis that egress fees constitute a meaningful barrier to switching or multi-cloud for a substantial minority of customers.¹⁶²³

6.361 IBM submitted that egress fees are a cost that people will consider when deciding whether to move, but they are not the main driver of the decision and not a cost that they would expect anyone else to pay. Specifically, IBM explained that the Egress fees working paper noted that most large customers had not identified egress fees as the main challenge to switching cloud provider or implementing a multi-cloud architecture and that almost none of the customers interviewed in the Jigsaw report viewed egress fees as the main or one of the main barriers to switching or multi-cloud.¹⁶²⁴

6.362 Microsoft submitted that:

(a) the Jigsaw report’s customer evidence did not demonstrate weak competition and, in fact, illustrated that egress fees were not a meaningful barrier to switching nor to multi-cloud.¹⁶²⁵

(b) the PDR’s consideration of broader factors affecting customer switching decisions – including the fact that free switching programmes are (in the PDR’s view) ‘voluntary’, customers’ awareness of the programmes and requirements to submit switching applications in advance – is inconsistent with the PDR’s one-dimensional assessment of egress fees’ influence on customer switching / multi-cloud decisions.¹⁶²⁶

6.363 We have considered these submissions. Whilst egress fees are one of many factors that customers take into consideration when making their switching and multi-cloud decisions, this assessment focuses on understanding the impact of

¹⁶²⁰ AWS’ submissions to the CMA [§§]. We note that AWS quotes the terminology of a ‘substantial minority of customers’ from our PDR, which we have revised in this report.

¹⁶²¹ [AWS’ response to the CMA’s PDR dated 28 January 2025](#), paragraph 53.

¹⁶²² AWS’ submission to the CMA [§§].

¹⁶²³ [Microsoft’s response to the CMA’s PDR dated 28 January 2025](#), paragraph 102. We note that Microsoft quotes the terminology of a ‘substantial minority of customers’ from our PDR, which we have revised in this report.

¹⁶²⁴ [IBM’s response to the CMA’s working paper on egress fees](#), page 1.

¹⁶²⁵ [Microsoft’s response to the Competitive landscape, Committed spend agreements and Egress fees working papers](#), paragraph 89.

¹⁶²⁶ [Microsoft’s response to the CMA’s PDR dated 28 January 2025](#), paragraph 99.

egress fees on these decisions and the extent to which they represent a barrier to switching and multi-cloud. As part of this assessment, we consider the free switching programmes below. We consider the combined impact of egress fees with other factors (such as technical barriers) and whether they lead to harm to competition, either on their own or in combination in chapter 9.

- 6.364 We assessed a range of evidence below in reaching a view on the relevance of egress fees to customer decisions. When assessing customer views, and in light of the submissions by cloud providers described above, we consider that although egress fees were not referred to as a key concern for a majority of customers, they can still act as a barrier to switching and/or multi-cloud if they affect some customers' decisions to do so. Egress fees can be one of a number of concerns that are relevant to some customers. In our assessment of the relevance of egress fees to customers' switching and multi-cloud decisions, we explain the types of customers who are more likely to be impacted by egress fees in their decision to switch and/or multi-cloud. In the conclusions of this chapter as well as in chapter 9, we consider the impact of egress fees in combination with technical barriers, as barriers to switching and/or multi-cloud and their impact on competition in the market.

Trends in egress fee prices

- 6.365 As set out below, Microsoft, AWS and Google submitted that egress fee prices (for egress fees in general) have decreased and that this is evidence of market competition.

AWS

- 6.366 AWS submitted that its egress fees for transferring data out to the internet had fallen by over 30% globally between 2018 and 2022,¹⁶²⁷ and by almost 40% between 2018 and 2023 and that this is evidence of a competitive market.¹⁶²⁸
- 6.367 AWS submitted data which it said showed that the effective price (ie net of discounts) charged for egress in general has decreased globally by 37% and by 25% for UK customers, between 2018 and 2023.¹⁶²⁹ AWS also submitted that this data showed that from 2019 to 2023 [REDACTED].¹⁶³⁰ AWS also noted that its egress margin [REDACTED], consistent with a pass-on rate exceeding [REDACTED].¹⁶³¹ AWS has since submitted an update of this data which it said showed egress prices have decreased by almost 40% between 2018 and 2024, and that this showed a pass-

¹⁶²⁷ AWS' response to the CMA's updated issues statement and working papers, paragraph 35(a).

¹⁶²⁸ AWS' response to the CMA's PDR dated 28 January 2025, paragraph 51.

¹⁶²⁹ AWS' response to the CMA's Egress fees working paper issued 23 May 2024, paragraph 5.

¹⁶³⁰ AWS' submission to the CMA [REDACTED].

¹⁶³¹ AWS' submission to the CMA [REDACTED].

on rate of over 100% between 2019 and 2024.¹⁶³² We consider AWS' submissions on pass-on cost savings later on in this chapter.¹⁶³³

- 6.368 AWS submitted that falling costs do not necessarily translate to falling prices for customers, as providers could keep cost reductions as profits. It said that AWS has lowered egress prices for customers in proportions equal to its efficiency gains which demonstrates competition.¹⁶³⁴
- 6.369 AWS submitted that if egress fees could lock in customers, then AWS would have no incentive to decrease egress prices, let alone decrease prices in line with costs, and would instead be incentivised to keep egress prices at a high level and increase egress margins, thereby maintaining its market power over existing customers.¹⁶³⁵
- 6.370 AWS submitted that there are diseconomies of scale in egress data transfers. AWS submitted that the cost of transferring a byte increases with the growth of the network due to maintenance and other costs, and that AWS has instead decreased costs through innovation.¹⁶³⁶

Microsoft

- 6.371 Microsoft submitted that competition has driven egress fees in general down and that quality-adjusted egress fees have fallen for customers in the UK, while the volume of data transferred had [REDACTED]. Microsoft also submitted that nominal effective prices for Microsoft's egress services in general had remained stable despite increases in product quality and infrastructure since 2018.¹⁶³⁷ We consider Microsoft's submissions on falling list prices for egress products further later in this chapter.
- 6.372 Microsoft submitted that the list prices for its egress products (Bandwidth and ExpressRoute¹⁶³⁸) have fallen over time. Microsoft submitted the following as evidence of this:¹⁶³⁹
- (a) Microsoft's introduction of internet egress routed via public ISP network as a cheaper egress option to routing via its Premium Network.

¹⁶³² AWS' submission to the CMA [REDACTED].

¹⁶³³ See section 'Role of egress fees in passing on cost savings'.

¹⁶³⁴ [AWS response to the CMA's PDR dated 28 January 2025](#), paragraph 53.

¹⁶³⁵ AWS' submission to the CMA [REDACTED].

¹⁶³⁶ AWS' submission to the CMA [REDACTED].

¹⁶³⁷ Microsoft's submission to the CMA [REDACTED].

¹⁶³⁸ 'Bandwidth' offers services to move data in (ingress), within (internal) and out of Azure data centres (egress). ExpressRoute is a private connection service that connects Microsoft data centres to customers' on-premises or colocation facility infrastructure.

¹⁶³⁹ Microsoft's submission to the CMA [REDACTED].

- (b) Microsoft said it decreased the price of one add-on for its ExpressRoute service by 50% in June 2021 and prices of all other egress services have remained stable, having not changed since 2018.
- (c) Microsoft's expansion of the Bandwidth free tier from 5GB a month to 100GB a month in Q2 2022. Microsoft submitted that this decreased the effective price paid by customers for Bandwidth egress.

6.373 Microsoft also submitted that on average, real prices for both metered and unmetered egress plans have fallen since 2018.¹⁶⁴⁰ It submitted that the CMA's PDR placed limited weight on evidence of falling egress prices in general, but the CMA's Market Investigation Guidelines state that (inter alia) price is an important and measurable outcome for customers, which helps to determine whether there is an adverse effect on competition. Microsoft submitted that the UK cloud market dynamics have led to free ingress due to intense competition for workloads and falling quality-adjusted prices for egress.¹⁶⁴¹

Google

- 6.374 Google submitted that the effective price¹⁶⁴² for standard transfers via the internet had been trending down over time. Google further submitted that this is 'clear evidence' that there was effective price competition in the market and pointed to the introduction of its free tier in October 2023 as an example of this competition.¹⁶⁴³ We also consider Google's submissions on prices later in this chapter.¹⁶⁴⁴
- 6.375 Google submitted that most of its networking SKUs have either experienced no price change or an effective price decrease and that this is clear evidence that there is effective price competition in the market. Google cited its introduction of a free tier across all Standard Tier SKUs in October 2023, keeping the prices of all other tiers the same and its free tier for standard internet transfers is more extensive than other free tiers.¹⁶⁴⁵

Our analysis of submissions on falling egress fee prices

- 6.376 Our guidelines say that there may be several factors affecting prices and that we take this into account when considering inferences from this type of analysis.¹⁶⁴⁶

¹⁶⁴⁰ Microsoft submission to the CMA [X].

¹⁶⁴¹ [Microsoft's response to the CMA's PDR dated 28 January 2025](#), paragraph 98.

¹⁶⁴² Here, Google uses 'effective price' to mean unit price implied by the actual usage and pricing tier.

¹⁶⁴³ [Google's response to Egress fees working paper](#), paragraph 15.

¹⁶⁴⁴ See 'Role of egress fees in passing on cost savings'.

¹⁶⁴⁵ [Google's response to the CMA's Egress fees working paper](#), paragraph 15.

¹⁶⁴⁶ [CC3 \(Revised\)](#), paragraph 108.

- 6.377 In addition, observing falling prices for egress, on its own, is not enough to conclude that the presence of egress fees cannot constitute a barrier to switching and/or multi-cloud. In the absence of prices falling to zero for all customers, egress fees have the potential to be a cost to customers which may constitute a barrier to switching and/or multi-cloud. Also, as noted in Market outcomes, price trends in isolation do not constitute direct evidence of the competitive nature of the market, especially in a growing industry with strong economies of scale where we would expect costs to fall.
- 6.378 The submissions made by cloud providers focus on 'effective prices' for egress fees in general, ie total revenues from egress fees (net of discounts) divided by quantity of egressed data. This does not tell us what the effective price has been or would be for switching and/or multi-cloud egress data transfers specifically, given it represents an overall average across all customers and all egress use cases. It is also difficult to interpret any customer views or impacts on customer decision-making from trends in prices, and hence we find it more useful to consider the level of egress fees in relation to switching and/or multi-cloud rather than price trends.
- 6.379 We note that the trends in effective prices submitted by cloud providers are also largely driven by the expansion of the egress fee-free tier and individually negotiated discounts on either egress fees or across customers' cloud spend. As the data is not specific to switching and/or multi-cloud egress, this means a fall in effective prices could reflect larger discounts on other types of egress data transfers (eg CDN transfers to end users) and/or changes in customer mix, for example a greater proportion of customers falling within the free tier. Neither of these scenarios would necessarily mean the price for switching and/or multi-cloud egress decreased for customers.
- 6.380 Microsoft's data on real prices for egress in general falling over time simply reflects that inflation has been positive (and significant in 2022) whilst nominal prices have remained constant. This does not provide evidence on the prices for switching and/or multi-cloud egress specifically, or show that the level of egress fees paid by customers has reduced.
- 6.381 We also do not accept AWS' submission that falling effective egress prices and/or margins in general would show that egress fees do not lock in customers because if they did, cloud providers would keep egress prices at a high level and increase egress margins to maintain market power over existing customers. Whilst increasing prices and/or margins is something a firm could do if it had some form of lock-in power, the absence of these indicators would not prove an absence of

any lock-in.¹⁶⁴⁷ Cloud providers' pricing decisions may be based on a number of factors and we find it more useful to consider the level of egress fees in relation to switching and/or multi-cloud, rather than the price or margin trends.

- 6.382 Our review of internal documents (set out further below) indicated that Microsoft, AWS and Google set egress fee prices not only based on costs and margins, but considering other factors, including competitors' pricing and the pricing of their other networking services. Microsoft, AWS and Google have also cited the increase or introduction of free tiers and/or introduction of tiered egress routing options as evidence of price reductions,¹⁶⁴⁸ and we note that these providers have introduced free switching programmes (discussed later in this chapter). This provides some evidence of price competition for egress fees.
- 6.383 However, in the last five years Microsoft and AWS have each only changed the free tier available to customers once and list prices for AWS non-CDN egress and Microsoft Standard and Premium tier egress have remained unchanged.¹⁶⁴⁹ Google introduced a free tier for Standard tier egress in 2023 and some other new free tiers in the last five years and increased its Premium Tier egress pricing in February 2024.¹⁶⁵⁰
- 6.384 As noted above, we find it more useful to consider the level of egress fees and have found that the list prices for Microsoft's, AWS' and Google's egress fees in general have largely not changed in the past five years. We also note that Microsoft, AWS and Google continue to have higher marginal list prices for egress outside of the free tier compared to other cloud providers.¹⁶⁵¹
- 6.385 On AWS' submission that there are diseconomies of scale in egress data transfers, we acknowledge that the growth of a cloud providers' network in order to meet additional demand may involve stepped-up investment in infrastructure and/or capacity with an associated increase in costs such as maintenance. However, once the investment has occurred, we consider there to be minimal incremental costs per additional egress data transfer, thereby leading to economies of scale.¹⁶⁵² It is unclear why maintenance or other costs would scale

¹⁶⁴⁷ We note that our analysis of AWS' egress fee revenues and costs also found that AWS' margins for its egress data transfer services overall [§] from 2019 to 2022, and [§] in 2023 and 2024. [§]. Further analysis of cloud provider egress margins is set out in appendix Q.

¹⁶⁴⁸ AWS' submission to the CMA [§]; Microsoft's submission to the CMA [§], [Google's response to the CMA's Egress fees working paper](#), paragraph 15.

¹⁶⁴⁹ Note, Microsoft cites one instance of decreasing the price of its ExpressRoute service, in 2021 and [§]. Microsoft submission to the CMA [§]; Microsoft's response to the CMA's information request [§]; AWS' response to the CMA's information request [§].

¹⁶⁵⁰ Google's response to the CMA's information request [§]; [Announcement of pricing change for egress traffic](#).

¹⁶⁵¹ See appendix L.

¹⁶⁵² See appendix Q for our analysis of cloud providers' costs of providing egress data transfer services. We have found that that there is shared network infrastructure for providing data transfer services so the cost of additional network investment will also be incurred for a range of services including egress, ingress and internal data transfer services. We also found that whilst there may be some peering or IP transit costs that are closer to being incremental in nature, these costs comprise a relatively small proportion of wider networking and/or total costs for cloud providers and for cloud-to-cloud egress data transfers may be avoidable.

increasingly per egress data transfer once the additional investment has been made.¹⁶⁵³ In any case, the presence of economies and/or diseconomies of scale makes evidence about price trends and their significance for customers' decision-making more difficult to interpret. As noted above, we find it more useful to consider the level of egress fees in relation to switching and/or multi-cloud rather than price trends.

- 6.386 Overall, for the reasons explained above, we do not consider the submitted data on effective prices and price trends for egress fees in general allows us to draw conclusions about competition in the market.

Justification for egress fees in the context of free ingress

- 6.387 Cloud providers generally charge for egressing data out of the cloud, but not for ingressing data into the cloud, despite both processes using the same fixed assets and therefore represent common costs incurred by a cloud provider.^{1654,1655} This raises the question of the justification for charging egress fees (in particular, when customers are transferring data for switching or multi-cloud use) and the level of the fees, given that data ingress is free.
- 6.388 AWS submitted that the same network supports inbound and outbound transfers at the same time, but that their capacity is built to support the requirements of the dominant direction of traffic (in this case outbound) rather than the sum of both directions. The cloud provider accounts for data transfer costs based on this dominant direction of traffic.¹⁶⁵⁶
- 6.389 Google submitted that to avoid double-charging customers, cloud providers had to choose whether to charge for data transfer at the point of entry (ingress) or the point of exit (egress). The industry standard is to charge at the point of exit as it is easier to monitor data on customers, projects and products for data egress than it is for data ingress.¹⁶⁵⁷
- 6.390 Microsoft submitted that, historically, cloud providers charged for both egress and ingress. It said that competitive practices have driven out charging for ingress fees on private premium networks, and instead accommodated the egress fees that pay for the networks.¹⁶⁵⁸ Microsoft also submitted that there are rational economic reasons to expect that ingress fees become free and egress fees remain positively priced in a competitive cloud market – specifically as this facilitates migration to

¹⁶⁵³ Note, we analyse the cost of egress data transfers in general, including providers' submissions on incremental costs, in appendix Q.

¹⁶⁵⁴ See appendix Q for information common vs egress-specific data transfer costs.

¹⁶⁵⁵ Responses to the CMA's information requests [38].

¹⁶⁵⁶ AWS' response to the CMA's information request [38].

¹⁶⁵⁷ [Google's response to Egress fees working paper](#), paragraph 5, footnote 3.

¹⁶⁵⁸ [Microsoft's response to the CMA's PDR dated 28 January 2025](#), paragraph 26.

the cloud, while mitigating overall costs for integrating multi-cloud solutions and for customer switching.¹⁶⁵⁹

- 6.391 Oracle submitted that if egress fees for all cloud providers were based on a cost-recovery model, the costs of egress fees in general should be variable based on cloud providers' actual underlying costs. Oracle submitted that instead, excessive egress fees are used by some cloud providers as an easy way to extract unreasonable profit margin while also serving as a gating factor, helping to lock in their customers.¹⁶⁶⁰
- 6.392 We consider that these submissions offer useful context for our assessment of egress fees in relation to switching and multi-cloud. However, they do not directly address whether egress fees influence customers' decisions to switch and/or multi-cloud, and therefore whether egress fees act as a barrier to competition for cloud services.
- 6.393 Nevertheless, we note that Microsoft, AWS and Google previously charged for both ingress and egress fees.
- (a) [REDACTED] and [REDACTED] described ingress fees as having been waived in response to competitive pressure.¹⁶⁶¹
 - (b) a [REDACTED] internal document showed [REDACTED].¹⁶⁶²
 - (c) further, Microsoft's submission that data ingress being free facilitates migration to the cloud suggested that data egress being paid hinders migration out of the cloud.
- 6.394 Overall, these points do not contribute to our assessment of whether egress fees are likely to harm competition. Whilst AWS' submission refers to the pricing for egress and ingress as being based on costs, there is some evidence that pricing is also influenced by competition considerations.¹⁶⁶³

Internal documents

- 6.395 We reviewed Microsoft's, AWS' and Google's internal documents related to egress fees. We note that the internal documents reviewed generally related to egress fees in general, rather than switching and/or multi-cloud egress fees specifically. Our analysis in this section therefore refers to egress fees in general, except where switching and multi-cloud egress is specifically contemplated.

¹⁶⁵⁹ Microsoft's response to the Competitive landscape, Committed spend agreements and Egress fees working papers, paragraph 85.

¹⁶⁶⁰ Oracle's response to the Updated Issues Statement and working papers, page 2.

¹⁶⁶¹ [REDACTED] note this was also [REDACTED]. [REDACTED] response to the CMA's information request [REDACTED]; [REDACTED] submission to Ofcom [REDACTED].

¹⁶⁶² [REDACTED] response to the CMA's information request [REDACTED].

¹⁶⁶³ Note, this also aligns with our review of internal documents, set out further below.

6.396 Some of the cloud providers' internal documents indicated that egress fees are likely to be a concern and a barrier for some customers and can be a factor that can influence customer behaviour and/or decisions.¹⁶⁶⁴ For example:

(a) [REDACTED];¹⁶⁶⁵

(b) [REDACTED].¹⁶⁶⁶ It also indicates that egress fees may be or could become relevant and/or material for some customers, such as those that multi-cloud or operate in high egress usage industries.¹⁶⁶⁷

6.397 Our review of internal documents also indicated that Microsoft, AWS and Google set egress fee prices not only based on costs and margins, but they also consider other factors, including competitors' pricing and the pricing of their other networking services. This indicated that there are strategic considerations in the setting of egress fee prices and that cloud providers may manage their revenue/cost recovery from networking services overall, rather than primarily focusing on this for individual services (eg egress) in isolation.

Analysis of customers' costs for switching and multi-cloud scenarios

6.398 As explained above, when assessing the impact of egress fees on switching and/or multi-cloud, we did not rely on data on actual spend on egress fees, as the level of spending was likely influenced by the level of egress fees. As such, we considered evidence on how influential the costs would be if a customer were to switch or multi-cloud (ie regardless of whether they actually switch or multi-cloud). For this purpose, we have analysed the cost of egressing data in a range of scenarios of switching and of multi-cloud architectures, which we have set out below.¹⁶⁶⁸

6.399 The submissions and analysis set out in this section relate to egress fees for switching and/or multi-cloud use specifically.

Switching

6.400 We have assessed the egress fees that customers would face to switch cloud provider by considering the following evidence:

¹⁶⁶⁴ Our review is set out in more detail in appendix P.

¹⁶⁶⁵ [REDACTED] response to the CMA's information request [REDACTED].

¹⁶⁶⁶ [REDACTED] submission to the CMA [REDACTED].

¹⁶⁶⁷ We have set out the evidence from cloud providers' internal documents on their commercial strategies in relation to egress fees in appendix P.

¹⁶⁶⁸ We previously described this as an analysis of 'hypothetical' switching and multi-cloud costs. We have removed the word 'hypothetical' to make clear that we have used actual customer data as opposed to hypothetical cost data. Details of the methodology for this analysis can be found in appendix N.

- (a) cloud providers' estimates of the cost of switching arising from egress fees; and
- (b) data on the egress fees that would be incurred based on a 'one-off' switch of cloud provider.

Cloud provider views on 'one-off' switching costs

6.401 Microsoft, AWS and Google submitted that egress fees incurred by switching customers were a small percentage of their total annual spend. For example:

- (a) AWS, using its own pricing calculator, estimated that the egress fees incurred by switching customers represented less than [0-5%] of their annual spend on AWS. AWS noted that its approach to calculating this figure is similar to that of a 2017 IDC Report for the European Commission but did not explicitly explain how they calculated the estimate.¹⁶⁶⁹ AWS submitted that this is a level that is unlikely to deter customers from switching;¹⁶⁷⁰
- (b) AWS further submitted an analysis of data transfer out charges incurred by customers in the process of switching based on its own data, which estimated that these charges represent less than [0-5%] to [0-5]% of these customers' annual spend;¹⁶⁷¹
- (c) Microsoft referenced the same 2017 IDC Report as AWS, which estimated that the one-off egress fee cost of a customer performing a full switch was approximately 0.3% – 0.35% of the expected annual operating cost using either Azure or AWS for cloud services.¹⁶⁷² Microsoft also submitted that it provides customers with an option to use direct connections to Azure which allows customers transferring significant amounts of data to reduce their overall egress fee charges and to transfer at faster latency (eg Azure ExpressRoute);¹⁶⁷³ and
- (d) Google estimated that egress fees as a direct result of switching exceed no more than [0-5%] of a customer's contract spend over a typical three-year contract, based on customers it had identified as switching away from GCP. Google further submitted that egress fees covering all types of egress traffic are still a small proportion of an average customer's spend.¹⁶⁷⁴

¹⁶⁶⁹ See [AWS Pricing Calculator](#). AWS noted that the approach used to reach this estimate is similar to the one followed in the Final Report of the study "SMART 2016/0032 – Switching of Cloud Services Providers" entrusted by the European Commission to IDC and Arthur's Legal (available at [Switching of Cloud Services Providers - Publications Office of the EU](#)) (accessed on 10 September 2024).

¹⁶⁷⁰ [AWS' response to Ofcom's Interim Report](#), paragraph 17.

¹⁶⁷¹ AWS' submission to Ofcom [38].

¹⁶⁷² [Microsoft's response to Ofcom's Interim Report](#), paragraph 237.

¹⁶⁷³ [Microsoft response to the Issues Statement](#), paragraph 41.

¹⁶⁷⁴ Google's submission to Ofcom [38].

- 6.402 A customer submitted that egress fees not only hamper competition but also stifle innovation by discouraging customers from using their data and/or exploring alternative solutions, resulting in vendor lock-in.¹⁶⁷⁵
- 6.403 We note that the 2017 IDC analysis referenced by both Microsoft and AWS used an individual ‘mid-sized’ customer with a specific workload as the base for its analysis.¹⁶⁷⁶ No reference was made to this customer being representative of any wider customer group and there is no reason to believe this is the case. We also consider that the 2017 IDC analysis is likely outdated given how much the markets have evolved in the eight years since its publication.¹⁶⁷⁷ We therefore consider this analysis is unlikely to be indicative of the egress fees the average customer may expect to be charged when switching.
- 6.404 We further note that both AWS’ and Google’s analyses are based on customers that have been identified as having switched. Basing such an analysis on customers that have switched will only capture those customers that considered egress fees low enough that they do not disincentivise switching and will tend to exclude any customers that found egress fees too high and were deterred from switching. We therefore consider that these analyses may underestimate the cost of switching, and therefore we place less weight on these results relative to our own analysis of switching costs. However, we note that these figures indicate that customers may face meaningful egress fees for switching their data between cloud providers.
- 6.405 AWS submitted a further analysis which uses customer level data, and found that the average one-off switching costs across customers is around [REDACTED] of total spend on AWS services. We discuss this analysis in appendix M.

Our analysis of switching costs

- 6.406 We calculated the ‘one-off’ switching cost for an average customer based on a scenario where they transfer all of their data out of their current cloud and express this cost as a percentage of their annual spend. This was based on the specific characteristics of real-world UK customers of Microsoft, AWS and Google. In particular:
- (a) we collected UK customer-level data on the mean volumes of data they stored for each year from 2018 to 2022;¹⁶⁷⁸

¹⁶⁷⁵ Wasabi’s response to the CMA’s working papers, pages 1 to 2.

¹⁶⁷⁶ Switching of Cloud Services Providers - Publications Office of the EU, page 42.

¹⁶⁷⁷ See chapter on competitive landscape.

¹⁶⁷⁸ CMA analysis of AWS, Microsoft and Google data. Responses to the CMA’s information requests [REDACTED].

- (b) we used the list prices on cloud providers' websites for data transfer from the applicable UK region,¹⁶⁷⁹ via the internet;¹⁶⁸⁰
- (c) we multiplied each customer's mean data storage volume by the price per GB for egress that would apply were they to transfer all of this data at once; and
- (d) we then calculated the percentage of annual spend this would make up if spread over the duration of a typical contract (which we have found to be three years).¹⁶⁸¹ This is because switching costs can be spread over multiple years. A customer may incur switching costs once but incur the benefits of switching over multiple years.

6.407 Our analysis shows that:

- (a) some customers face meaningful egress fees to switch their cloud provider. For example, [0-5]% to [5-10]% of customers would have had to pay egress fees of more than [5-10]% of their annual cloud spend per year over the course of a typical three-year contract, and [0-5]% of customers would have had to pay fees over [10-20]% per year.
- (b) across all providers, the smallest customers face higher egress fees than other customers as a proportion of their annual cloud spend in order to switch cloud providers. The average egress fees faced by customers with annual spends of £1,000-£10,000 is [0-5]% which is, between [0-5]% percentage points higher per year than the averages of higher spend bands.¹⁶⁸²

6.408 This means that around [0-5]% to [5-10]% of customers would not have the incentive to switch to a rival cloud provider even if, over the course of a typical contract, that provider offered a price that was [5-10]% better or equivalently better quality or functionality. The disincentive to switch would be even higher for the smallest customers.

6.409 We consider that some customers may have to take into account the up-front cost in the initial year, rather than exclusively considering the cost of switching spread out over the course of a contract. For example, budget constraints may mean that customers are unable to spread the cost over multiple years. We have therefore also calculated the egress fees which would be faced by customers as a

¹⁶⁷⁹ We have used the 'London' region for Google and AWS as this encompasses all of the UK.

¹⁶⁸⁰ Google: [All networking pricing - Virtual Private Cloud - Google Cloud](#), Microsoft: [Pricing – Bandwidth | Microsoft Azure](#), AWS: [Amazon S3 Simple Storage Service Pricing - Amazon Web Services](#) and [EC2 On-Demand Instance Pricing – Amazon Web Services](#), all accessed on 29 Jan 2024.

¹⁶⁸¹ CMA analysis of AWS and Microsoft data. Responses to the CMA's information requests [§<].

¹⁶⁸² This compares the average cost of switching for the £1,000-10,000 spend band, to the lowest and highest average costs for the other spend bands, for each of the three cloud providers in our analysis. See details of our analysis on average costs of switching by spend band in appendix M.

percentage of annual cloud spend, if they do not spread the cost over multiple years.

- 6.410 We have found that [10-20]% to [20-30]% of customers would face an up-front egress fee cost of more than [5-10]% of their annual cloud spend in a year and around [0-5]% to [5-10]% of customers would have had to pay egress fees of more than [10-3]%.¹⁶⁸³
- 6.411 In interpreting this analysis, we have considered the following factors:
- (a) our analysis is likely to understate the effect of egress fees on switching. This is because we have calculated these fees as a proportion of total cloud spend, but egress fees are relevant only for data-intensive workloads. By including unrelated spend in the denominator, we are likely to underestimate (potentially by a significant degree)¹⁶⁸³ the effect of egress fees on switching for these data-intensive workloads.
 - (b) we used list prices for the most standard or basic egress data transfer service for each provider that we identified. Our analysis therefore does not take into account various cost-mitigating or premium alternatives which may be offered by cloud providers (eg direct connect services). However, our understanding is that these are often for specialised use cases (eg customers with very large data transfer requirements). Therefore, our analysis does not account for all the products or use cases for customers paying egress fees.
 - (c) our analysis does not account for discounts. We have not found a robust methodology to consistently apply discounts to egress fees (we assess [X] proposed methodology in appendix M). Overall, our methodology does not account for discounts, including cross-service discounts like CSAs, which are common mostly among the largest customers.¹⁶⁸⁴ Accounting for discounts in our analysis would lower the estimated egress fees customers who receive discounts would have to pay when switching.
 - (d) we have considered egress fees in the context of switching being a complex and lengthy decision, potentially requiring a period of using multiple clouds at the same time, and where egress fees are one barrier among others. We interpret the results of our analysis acknowledging that egress fees are not the only consideration made by customers when considering switching cloud providers.

¹⁶⁸³ Analysis presented in the licensing chapter indicates spending on storage (the most obvious data related service) accounts for only a fraction of the overall cloud spend.

¹⁶⁸⁴ For our analysis of the prevalence of committed spend discounts (CSDs), see chapter 8.

6.412 Overall, we consider these factors have opposing effects on our estimates but the combined effect of these different factors is ambiguous. We have not seen evidence that suggests that these factors, taken together, would significantly impact our estimates.¹⁶⁸⁵ Therefore, we consider our estimates to be of value in the overall assessment.

6.413 We consider the detailed submissions we have received on our analytical approach and the interpretation of our analysis in appendix M.

Multi-cloud

6.414 Below, we have analysed how costly it would be to egress data for customers running a multi-cloud architecture as a proportion of their total monthly cloud spend. Given multi-cloud covers a wide range of use cases, we modelled these costs under a variety of assumptions and break down our results by customers' spend.

6.415 Below, we consider:

- (a) cloud provider views on multi-cloud costs, including a Google submission on the costs of multi-cloud; and
- (b) our analysis of multi-cloud costs.

Cloud provider views on multi-cloud costs

6.416 Google submitted that egress fees constitute a meaningful barrier to multi-cloud adoption only when they increase costs relative to a situation where a customer already moves data between multiple locations within a single cloud provider.¹⁶⁸⁶

6.417 AWS submitted that, when considering hypothetical scenarios related to multi-cloud, we do not consider whether integrated multi-cloud is the most efficient outcome for customers.¹⁶⁸⁷ AWS added that, to determine whether the costs from egress fees have an impact on ability to multi-cloud, they should be compared to the benefits customers receive from multi-cloud. AWS submitted that this was not incorporated into our analysis.¹⁶⁸⁸ We agree comparing any multi-cloud costs to the benefits of multi-cloud is the right approach. We do this in the round in chapter 9.

¹⁶⁸⁵ For example, AWS' analysis (DEREF) suggests the impact of discounts (point (c) above) on our analysis is unlikely to be large. Separately, even considering the analysis discussed in the licensing chapter, the magnitude of the effect of not including unrelated spend in the denominator (point (a) above) is unclear. These two factors have opposing effects on our estimates.

¹⁶⁸⁶ Google's submission to the CMA [38].

¹⁶⁸⁷ AWS submission to the CMA [38].

¹⁶⁸⁸ AWS submission to the CMA [38].

- 6.418 AWS submitted that customers who multi-cloud usually employ architectures, cost-mitigating services and other specific tools to reduce data transfer costs or reduce the need for data transfer. It submitted that these cost-mitigating services have a significant impact on the actual price paid by customers who multi-cloud.¹⁶⁸⁹
- 6.419 Microsoft submitted that it has designed egress fee strategies to minimise the effects of egress fees on multi-cloud customers,¹⁶⁹⁰ but did not provide information on what these strategies entail.
- 6.420 IBM submitted that it may be difficult to accurately model the proportion of egress costs out of annual spend. It considers egress costs more suited to a qualitative analysis than a quantitative analysis.¹⁶⁹¹
- 6.421 Google also submitted an analysis which sought to quantify the difference between a customer's internal data transfer cost between multiple regions in a single-cloud setup and a customer's cross-cloud cost in a multi-cloud setup.¹⁶⁹² Google submitted that its results show that the cost difference between the two data transfer types is 'de minimis' in the context of overall cloud spend.¹⁶⁹³ We discuss this submission in further detail in appendix M.
- 6.422 We consider that Google's submission sets an artificially high benchmark for finding a concern due to its choice of hypothetical customer and the modelling choices it made in a single-cloud scenario:
- (a) first, we do not consider the customer chosen to be representative of the wider market. The customer chosen, a large bank, is likely to be subject to various requirements (eg resiliency/availability thresholds, cross-region data replication, data sovereignty) that other customers will not need to consider. This inflates single-cloud egress costs, as this customer transfers more data between regions than we may otherwise expect.
 - (b) second, we consider that the assumptions made when modelling the customer's single-cloud architecture artificially inflate egress fees in the single-cloud case, therefore reducing the difference in egress fees between the single-cloud and multi-cloud cases.
- 6.423 We further consider that internal transfer fees are not a relevant benchmark against which to assess the impact of egress fees on multi-cloud. That's because not all customers will incur internal transfer fees, eg if they do not operate across

¹⁶⁸⁹ AWS submission to the CMA [X].

¹⁶⁹⁰ [Microsoft's response to the Competitive landscape, Committed spend agreements and Egress fees working papers](#), paragraph 87.

¹⁶⁹¹ IBM's response to the CMA's information request [X].

¹⁶⁹² Google's submission to the CMA [X].

¹⁶⁹³ Google's submission to the CMA [X].

multiple regions. This has the effect of distorting the results. As such, we consider this analysis does not show that egress fees are not a barrier to multi-cloud.

Our analysis of multi-cloud costs

- 6.424 In this analysis, we have focussed on customers' ability to multi-cloud should they wish to, ie should the benefit they would obtain from it outweigh the costs of doing so. In that context, egress fees add some costs for the customer to adopt their preferred architecture.
- 6.425 With respect to Microsoft, AWS and Google individually:
- (a) we collected UK customer-level data on the mean volumes of data customers stored for each year from 2018 to 2022;¹⁶⁹⁴
 - (b) we used the list prices on cloud providers' websites for data transfer from the applicable London region via the internet (ie standard tier egress);¹⁶⁹⁵ and
 - (c) we then performed a series of calculations to assess the egress fees which customers would incur under various multi-cloud setups. The details of these calculations can be found in appendix M, which also includes the feedback on our methodology from cloud providers, where appropriate.
- 6.426 As there is no single use-case for multi-cloud and a multi-cloud architecture can, in theory, involve the transfer of any proportion of a customer's stored data, we calculated the egress fees a customer may expect to pay (as a proportion of monthly cloud spend) based on:
- (a) the customer's size (proxied by their annual spend on cloud services); and
 - (b) the proportion of their stored data that they may choose to transfer.
- 6.427 In our analysis, within each spend band, we calculated the average customer's annual cloud spend and the average volume of stored data. We use the mean and, in response to AWS submissions (see appendix M) the median as well. We consider that both are metrics are informative for our assessment.
- 6.428 Overall, our analysis shows that customers spending more than £10,000 face meaningful egress fee costs as a percentage of their cloud spend across different data transfer amounts. We considered the costs faced by the median customer of each spend band (in terms of volume of data stored and total cloud spend) and found that for a customer spending more than £10,000, the cost of egressing 20%

¹⁶⁹⁴ CMA analysis of AWS, Microsoft and Google data. Responses to the CMA's information requests [X].

¹⁶⁹⁵ Google: [All networking pricing - Virtual Private Cloud - Google Cloud](#), Microsoft: [Pricing – Bandwidth - Microsoft Azure](#), AWS: [Amazon S3 Simple Storage Service Pricing - Amazon Web Services](#) and [EC2 On-Demand Instance Pricing - Amazon Web Services](#), all accessed on 29 Jan 2024.

of their data per month ranges from [0-5]% to [10-20]% of their monthly cloud spend in egress fees alone.¹⁶⁹⁶ This increases to at least [5-10]% when egressing 50% of data stored.

6.429 Our analysis shows that the smallest customers (those who spend between £1k - £10,000) also face meaningful egress fees when transferring larger volumes of data. Based on our analysis of the costs faced by the median customer of each spend band (in terms of volume of data stored and total cloud spend), we found that:

- (a) these smallest customers tend to be covered by their cloud provider's free tier when transferring up to [10-20]% to [30-40]% of their data per month.
- (b) however, these customers face meaningful fees for egressing amounts greater than this. For example, we found that to egress 50% of their data per month, these customers would face an average cost of between [0-5]% and [5-10]% of their monthly spend in egress fees alone.

6.430 The results of our analysis are broadly consistent whether we use the mean or median to estimate the average volume of data and cloud spend. However, we observe that:

- (a) the results using the mean tend to be higher than the median. This difference is most significant amongst the smallest customers (those who spend between £1,000-£10,000) when egressing at least [20-30]% of their data per month; and
- (b) compared to larger customers, the difference between the median and the mean egress fees faced by the smallest customers increases at a faster rate as the volume of data being transferred increases.

6.431 This suggests that some of the smallest customers face meaningful egress fees when transferring at least [20-30]% of their data per month. These customers may have a particularly high disincentive to increase the level of integration between clouds or multi-cloud. This is because the costs they would face to transfer data between clouds at higher data transfer levels, in terms of percentage of their annual spend, are higher than larger customers.

6.432 In interpreting these results we account for the same points made in the analysis for switching, plus:

¹⁶⁹⁶ We note that, as set out in appendix M, some spend bands have very few observations because not many customers fell within that spend bracket in 2022. Therefore, as a sensitivity we considered the egress fees faced by customers in spend bands with at least [30] customers and found that this range was [0-5]% to [5-10]%.

- (a) Microsoft, AWS and Google all submitted that our data transfer volumes are not reflective of customers' current architectures.¹⁶⁹⁷ We consider that the current level of integration (and therefore the level of data transferred across clouds) in multi-cloud architectures is likely to be influenced by the current level of egress fees. This is because the higher egress fees are, the more expensive it is to multi-cloud. Therefore, the current level of data transferred across cloud does not provide a reliable indication of how much data would be transferred in the absence of such fees.

6.433 Overall, we consider these factors have opposing effects on our estimates but the combined effect of these different factors is ambiguous. As explained in our analysis of customers' costs for switching, we consider our estimates to be of value in the overall assessment.

6.434 We consider the detailed submissions we have received on our analytical approach and the interpretation of our analysis in appendix M.

Customer views on the impact of egress fees

6.435 We assessed whether egress fees are having an impact on customer behaviour by asking customers a range of questions about factors influencing their use of cloud services. In this section, we set out the evidence gathered from large customers on the relevance of egress fees in the context of switching and operating a multi-cloud architecture.

6.436 In line with the qualitative nature of the evidence we gathered, we have provided a narrative summary of the key points that we consider emerge from the evidence. Further detail on this is presented in the appendix O.

6.437 The Jigsaw report captured a wider range and a different set of customers from those we contacted directly and the key points from this are also presented below.

6.438 The evidence and analysis set out in this section relate to egress fees for switching and/or multi-cloud use specifically.

Switching

Customer evidence

6.439 We asked questions on switching to a variety of customers.

¹⁶⁹⁷ AWS submission to the CMA [§]; Google's submission to the CMA [§]; Microsoft's submission to the CMA [§]. See appendix M for further details of these submissions.

- (a) among customers that have switched or have considered switching, only some spontaneously identified egress fees as a challenge.¹⁶⁹⁸ Others said that they had not found egress fees to be an issue or they had not explicitly considered them and that there were other costs to switching that were far more significant.¹⁶⁹⁹
- (b) among those customers who had never considered switching, some had not done so because they were satisfied with or had only recently moved to their current cloud provider.¹⁷⁰⁰ However, some other customers that had never considered switching explicitly identified egress fees as a factor that disincentivises switching.¹⁷⁰¹

6.440 Other customers indicated that egress fees would impact their decision-making when considering a change of cloud provider, as these fees could be substantial if switching application from one cloud provider to another.¹⁷⁰²

Evidence from the Jigsaw report

- 6.441 The Jigsaw report highlights that in almost no cases were egress fees considered the main or even one of the main barriers to switching. Technical challenges and the lack of a clear business case were cited more often as a reason not to switch.¹⁷⁰³
- 6.442 However, some customers did say that egress fees played a part in disincentivising them from potentially considering a migration to another cloud provider. One-off 'exit costs' were mentioned by some customers as a potential source of concern when switching cloud provider and transferring large amounts of data out of the incumbent cloud and into a new one.¹⁷⁰⁴ These concerns typically involved a company with relatively high data volumes, or mature companies who have accumulated a large amount of data over the years.¹⁷⁰⁵
- 6.443 For some customers who have switched, or use multiple cloud providers, the Jigsaw report shows that they consider the egress fees they incurred a price worth paying to deliver the cloud strategy that makes most sense for their business.¹⁷⁰⁶

¹⁶⁹⁸ Responses to the CMA's information requests [§].

¹⁶⁹⁹ Notes of meetings with [§].

¹⁷⁰⁰ Responses to the CMA's information requests [§].

¹⁷⁰¹ Responses to the CMA's information requests [§].

¹⁷⁰² [§] response to the CMA's information request [§]; Notes of meetings with [§].

¹⁷⁰³ [Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw \(2024\)](#), paragraphs 5.2.2 and 5.2.3.

¹⁷⁰⁴ [Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw \(2024\)](#), paragraph 5.2.5.

¹⁷⁰⁵ [Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw \(2024\)](#), paragraph 5.2.2-4.

¹⁷⁰⁶ [Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw \(2024\)](#), paragraph 5.3.2.

6.444 When assessing customer views, although egress fees are not referred to as a key concern for a majority of customers, we consider they can still act as a barrier to switching and/or multi-cloud as they affect some customers' decisions to do so.

Other stakeholder views

6.445 In response to our PDR, we received submissions from other stakeholders.¹⁷⁰⁷ Some of them agreed with our provisional finding that egress fees influence customers' decisions and represent a barrier to switching and/or multi-cloud decisions.

- (a) BlackBox Hosting, a UK-based private cloud hosting provider, submitted that egress fees are a significant disincentive to switch to a private cloud provider, and that they are part of a commercial strategy to lock customers in.¹⁷⁰⁸
- (b) The Open Cloud Coalition,¹⁷⁰⁹ an industry body, submitted that high egress fees make it financially punitive for customers to migrate workloads away from dominant providers.¹⁷¹⁰
- (c) OVHcloud, a smaller European cloud provider, submitted that cloud users' ability to switch or multi-cloud is significantly hindered by technical and commercial practices from market leaders, with egress fees as an example of such a commercial practice.¹⁷¹¹

Multi-cloud

Customer evidence

6.446 Several customers we contacted identified themselves as using an integrated form of multi-cloud architecture, with different levels of integration.¹⁷¹²

6.447 Some customers with integrated multi-cloud architectures said that they did not consider anything, including egress fees, to be a challenge in setting up a multi-cloud architecture.¹⁷¹³

6.448 Some of the customers with integrated multi-cloud architectures said that egress fees have been a challenge to multi-cloud and/or they took egress fees into consideration when making their decision to multi-cloud.¹⁷¹⁴ Reasons for this

¹⁷⁰⁷ For example, Dr George R Barker's response to the CMA's PDR dated 28 January 2025, pages 94-96.

¹⁷⁰⁸ [Blackbox Hosting's response to the CMA's PDR dated 28 January 2025](#), paragraph 17.

¹⁷⁰⁹ Open Cloud Coalition members include Google Cloud.

¹⁷¹⁰ [Open Cloud Coalition's response to the CMA's PDR dated 28 January 2025](#), page 4.

¹⁷¹¹ [OVHcloud's response to the CMA's PDR dated 28 January 2025](#), paragraph 8.

¹⁷¹² Responses to the CMA's information requests [REDACTED].

¹⁷¹³ Responses to the CMA's information requests [REDACTED].

¹⁷¹⁴ Responses to the CMA's information requests [REDACTED].

included that egress fees make it difficult to use the lowest cost service across cloud provider and egress fees incentivise the use of one cloud provider.

- 6.449 Of those customers who use a siloed multi-cloud architecture, a few suggested that egress fees had a significant impact on their multi-cloud architecture or that egress fee costs were significant enough to be a consideration in the design phase of a multi-cloud architecture.¹⁷¹⁵
- 6.450 A customer said at the time of submission that egress fees were not a material challenge because they were not a material disincentive to multi-cloud for its use cases.¹⁷¹⁶ However, it has since developed a use case that involves the synchronisation of large amounts of data between two repositories. It submitted that, in this case, egress fees do provide a disincentive to multi-cloud.¹⁷¹⁷
- 6.451 For some large customers, egress fees related to multi-cloud were also mentioned when discussing other perceived issues. These customers did not self-categorise by cloud architecture.
- (a) one customer said that it had managed to reduce egress fees through the architecture it had employed.¹⁷¹⁸
 - (b) some customers indicated that egress fees would impact their decision-making when considering a switch to multi-cloud or changing cloud providers, as egress fees are expensive given the volume of data that they would need to transfer, or were a barrier more generally.¹⁷¹⁹
 - (c) some customers we contacted said they used a single cloud provider.¹⁷²⁰ We asked these customers why they used a single public cloud and if there were any potential challenges that they would encounter if moving to a multi-cloud architecture. None of the customers commented on whether egress fees were a reason to adopt a single cloud architecture or whether egress fees had been a challenge to start using multi-cloud.

Evidence from the Jigsaw Report

- 6.452 In relation to egress fees as a barrier to multi-cloud, the Jigsaw report is consistent with the rest of our customer evidence: some customers said that egress fees played a part in disincentivising them from potentially considering a multi-cloud approach. The Jigsaw report highlights that in almost no cases were egress fees

¹⁷¹⁵ Responses to the CMA's information requests [X].

¹⁷¹⁶ [X] response to the CMA's information request [X].

¹⁷¹⁷ [X] submission to the CMA [X].

¹⁷¹⁸ [X] response to the CMA's information [X].

¹⁷¹⁹ Responses to the CMA's information requests [X]; Notes of meetings with [X].

¹⁷²⁰ Responses to the CMA's information requests [X].

considered as the main or even one of the main barriers to a multi-cloud approach.¹⁷²¹

- 6.453 However, in some cases, customers expressed that egress fees presented a challenge to maintaining a multi-cloud architecture. While this was usually mentioned in the context of a hypothetical multi-cloud scenario, this consideration does suggest that egress fees could act as a barrier or deterrence to even considering a multi-cloud strategy. When egress fees did contribute to this reluctance, this typically involved a company with relatively high data volumes, or mature companies who have accumulated a large amount of data over years.¹⁷²²
- 6.454 The Jigsaw report found cases of customers voicing concern about egress fees' impact on a potential multi-cloud strategy. These customers were particularly concerned about the costs of keeping large databases synchronised across different cloud providers.¹⁷²³ Customers also noted that the increasing importance of AI, alongside the expected larger data volumes, would bring egress fees more into focus in the future. We briefly explore this topic in the section 'Impact of the emergence of AI', below.
- 6.455 Customers also voiced concern about the lack of transparency and overall complexity of egress fees, which make it difficult for some customers to predict costs for their general cloud usage. This is relevant in the switching and multi-cloud context, as any lack of transparency and control makes it difficult for customers to include egress fees in a potential cost-benefit calculation for switching or using multiple cloud providers.¹⁷²⁴ This perception was particularly strong among customers with high data volumes.¹⁷²⁵
- 6.456 Some customers and a cloud provider said that the rise of generative AI means that, in the future, egress fees may become more relevant.¹⁷²⁶ We understand that this would be because customers wanting to develop or use AI services may need to move their data, either as part of a switch between clouds or as part of a multi-cloud architecture. We have assessed the impact of AI on switching and multi-cloud in general in the Impact of AI chapter. In that chapter we reviewed evidence related the importance of data transfers to uses cases related to AI. However, we consider the evidence we reviewed not to be indicative as to whether egress fees will become more or less important in the future.

¹⁷²¹ [Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw \(2024\)](#), paragraph 5.2.2.

¹⁷²² [Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw \(2024\)](#), paragraph 5.2.4

¹⁷²³ [Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw \(2024\)](#), paragraph 5.2.5.

¹⁷²⁴ [Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw \(2024\)](#), paragraph 5.3.6.

¹⁷²⁵ [Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw \(2024\)](#), paragraph 5.1.7.

¹⁷²⁶ [Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw \(2024\)](#), paragraph 5.2.10; Oracle's response to the Updated Issues Statement and working papers, page 5.

Free switching programmes

- 6.457 The European Union's (EU) Data Act (the 'EU Data Act')¹⁷²⁷ entered into force in January 2024. The EU Data Act governs the use and exchange of data within the EU and imposes obligations on cloud providers via provisions that refer to 'data processing services'.¹⁷²⁸ The EU Data Act contains provisions relevant to egress fees for EU customers:¹⁷²⁹
- (a) Article 29 requires that any 'switching charges' (which includes data egress charges relating to switching) 'cannot exceed costs incurred by the provider of data processing services that are directly linked to the switching process concerned' and applies from 11 January 2024 until 12 January 2027.¹⁷³⁰ From 12 January 2027, a full removal of switching charges is required.¹⁷³¹
 - (b) Article 34(2) requires that egress data charges for customer use of data processing services in parallel with another data processing service (which we consider to cover multi-cloud use) cannot exceed costs incurred by the cloud provider. Our understanding is that this will take effect along with other main provision of the EU Data Act on 12 September 2025.¹⁷³²
- 6.458 Since the start of the application of Article 29(2) of the EU Data Act, programmes providing free egress for switching ('free switching programmes') have been introduced globally (and therefore include the UK) by Microsoft, AWS and Google. Following the announcement of these free switching programmes, egress fees have been removed by Civo (a UK-based smaller cloud provider).¹⁷³³
- 6.459 IBM said it plans to charge egress at cost for EU-based customers that are switching cloud provider between 11 January 2024 and 11 January 2027, after which this egress will reduce to nil in line with the EU Data Act.¹⁷³⁴

¹⁷²⁷ [Regulation \(EU\) 2023/2854 of the European Parliament and of the Council of 13 December 2023 on harmonised rules on fair access to and use of data and amending Regulation \(EU\) 2017/2394 and Directive \(EU\) 2020/1828 \(Data Act\)](#) ('EU Data Act').

¹⁷²⁸ A data processing service is defined in the Data Act as 'a digital service that is provided to a customer and that enables ubiquitous and on-demand network access to a shared pool of configurable, scalable and elastic computing resources of a centralised, distributed or highly distributed nature that can be rapidly provisioned and released with minimal management effort or service provider interaction' (Article 2(8)). See, also, recital 81 of the Data Act: 'The generic concept 'data processing services' covers a substantial number of services with a very broad range of different purposes, functionalities and technical set-ups. As commonly understood by cloud providers and users and in line with broadly used standards, data processing services fall into one or more of the following three data processing service delivery models, namely Infrastructure as a Service (IaaS), Platform as a service (PaaS) and Software as a Service (SaaS). Those service delivery models represent a specific, pre-packaged combination of ICT resources offered by a provider of data processing service. Those three fundamental data processing delivery models are further complemented by emerging variations, each comprised of a distinct combination of ICT resources, such as Storage as a Service and Database as a Service.'

¹⁷²⁹ [EU Data Act](#), Articles 29(1), 29(2) and 29(3).

¹⁷³⁰ [EU Data Act](#), Articles 29(2) and 29(3).

¹⁷³¹ [EU Data Act](#), Article 29(1).

¹⁷³² [EU Data Act](#), Article 50.

¹⁷³³ [Civo abolishes all egress fees to advance cloud simplicity strategy](#), accessed on 24 September 2024.

¹⁷³⁴ IBM's response to the CMA's information request [🔗].

- 6.460 Oracle said that [REDACTED].¹⁷³⁵ We understand this to mean that Oracle [REDACTED] already compliant with the current EU Data Act provisions that apply until January 2027 in relation to egress fees for switching.
- 6.461 We considered whether the introduction of these free switching programmes means that we may be overestimating the switching costs incurred by a UK customer that wants to switch to another cloud provider. If the scope and eligibility criteria of these free switching programmes are sufficiently wide, UK customers who are switching will generally not incur egress fees.
- 6.462 To assess this, we gathered evidence on the nature and applicability of Microsoft's, AWS' and Google's free switching programmes from these cloud providers and from information published on their websites. The coverage of each cloud provider's free switching programme is set out in detail in appendix N along with each cloud provider's views and information on uptake, rationale and cost of the programmes.
- 6.463 We note that the submissions and analysis set out in this section relate to egress fees for switching and/or multi-cloud use specifically.

Stakeholder views on switching programmes

- 6.464 Google submitted that the launch of free switching programmes by Microsoft, AWS and Google has effectively taken the issue of egress fees as a potential barrier to switching off the table.¹⁷³⁶ Google further submitted that as a result of these programmes, most customers can now freely switch their entire cloud spend to another cloud provider or an on-premises environment (with Google and AWS also supporting free partial switching) and that therefore egress fees do not pose any real or perceived barrier to customer switching.¹⁷³⁷
- 6.465 Once Google introduced its free switching programme, AWS also removed its egress fees for switching. AWS submitted that one of the reasons for this was it [REDACTED].¹⁷³⁸
- 6.466 AWS submitted that, to the extent that there was any perceived concern that egress fees act as a barrier to switching or multi-cloud, AWS' elimination of egress fees globally (which includes the UK) for switching customers removes one of the two potential concerns around egress fees identified in our Egress fees working paper.¹⁷³⁹

¹⁷³⁵ Oracle's response to the CMA's information request [REDACTED].

¹⁷³⁶ [Google's response to Egress fees working paper](#), annex response (b).

¹⁷³⁷ [Google's response to Egress fees working paper](#), paragraph 42.

¹⁷³⁸ Note of meeting with AWS [REDACTED].

¹⁷³⁹ [AWS' response to the Egress fees working paper](#), paragraph 32; [AWS' response to the CMA's PDR dated 28 January 2025](#), paragraph 54.

- 6.467 Microsoft cited the free switching programmes by Microsoft, AWS and Google as indication that there is no evidence of harm to consumers as a result of egress fees and submitted that customers remain able to switch or to multi-cloud when it makes commercial sense to do so.¹⁷⁴⁰
- 6.468 Google and AWS submitted that [REDACTED].¹⁷⁴¹ AWS also submitted that the PDR's assessment of the uptake of the programmes does not consider how AWS has publicised its programme through numerous AWS channels and that its programme has received significant media coverage.¹⁷⁴²
- 6.469 Microsoft also submitted that there has been a low uptake on the major cloud providers' elimination of egress fees for switching over the past year.¹⁷⁴³
- 6.470 Microsoft and AWS said that if egress fees had been a barrier to switching, then eliminating them would have resulted in material switching.¹⁷⁴⁴
- 6.471 Oracle said that the 'voluntary commitments' of some cloud providers should not be accepted at face value as they are crafted to cause minimal impact on their most lucrative anticompetitive practices. Notably, Oracle highlighted that AWS' free switching programme would not apply to multi-cloud architecture,¹⁷⁴⁵ and customers large enough to exceed the AWS programme's 100GB of data per month eligibility threshold are likely to require a multi-cloud strategy.¹⁷⁴⁶
- 6.472 In its announcement about removing egress fees, Civo said that its removal of egress fees comes with no requirements or limits, which it said was in contrast to other cloud providers that only ended egress fees with significant caveats, including requiring a customer to exit the platform.¹⁷⁴⁷
- 6.473 A customer told us that the cloud providers' 'changed policies' for switching do not address data egress charging in non-switching scenarios such as where an

¹⁷⁴⁰ [Microsoft's response to the Competitive landscape, Committed spend agreements and Egress fees working papers](#), paragraphs 90-91.

¹⁷⁴¹ AWS' response to the CMA's information request [REDACTED]; Google's response to the CMA's information request [REDACTED].

¹⁷⁴² [AWS' response to the CMA's PDR dated 28 January 2025](#), paragraph 54. Note, the PDR's provisional assessment was that customer uptake data is inconclusive as, amongst other factors, there may be low customer awareness of the programmes.

¹⁷⁴³ [Microsoft's response to the CMA's PDR dated 28 January 2025](#), paragraph 99.

¹⁷⁴⁴ [Microsoft's response to the CMA's PDR dated 28 January 2025](#), paragraph 99. [AWS' response to the CMA's PDR dated 28 January 2025](#), paragraph 54.

¹⁷⁴⁵ In response to the finding that AWS' free switching programme would not apply to multi-cloud architecture, AWS submitted that insofar as this is intended to address the fact that the free switching programme does not cover data transfers for the purposes of multi-clouding, that is correct. However, it submitted that customers are eligible to use the credits they receive under the programme for egress across all services and data types, including for the purpose of effecting a multi-cloud strategy during the process for a partial or total switch from AWS. [AWS' response to the CMA's PDR dated 28 January 2025](#), paragraph 56.

¹⁷⁴⁶ [Oracle, consolidated commentary to the CMA's Updated Issues Statement, its working papers, and the Market investigation qualitative customer research final report](#), page 4.

¹⁷⁴⁷ [Civo abolishes all egress fees to advance cloud simplicity strategy](#), accessed on 24 September 2024.

enterprise has demand to use a variety of cloud providers. This customer said it supports a wider application of data egress policy.¹⁷⁴⁸

Our analysis of relevance of free switching programmes to UK customers

6.474 We assessed the scope and conditions of eligibility for Microsoft's, AWS' and Google's free switching programmes. In particular, we looked at:

- (a) the types of customer switching covered by the programmes (full or partial switching, temporary multi-cloud use to enable switching);
- (b) the requirements and restrictions that impact which customers are eligible (eg the eligible geographies and data products, exceptions based on free tiers) and how a switch is completed to be eligible under the terms of the programmes (eg time limits, services that can be used to transfer data, requirements to submit requests in advance and close accounts);
- (c) the aspects of the programmes that are uncertain or at cloud providers' discretion; and
- (d) the requirement for customers to be aware of the existence of the programmes.

6.475 Our analysis of free switching programmes is set out in more detail in appendix N. In particular, we considered the scope of these programmes, the extent to which the programmes have restrictions or limitations and whether these are relevant to our assessment of egress fees as a barrier to switching and/or multi-cloud use (as submitted by cloud providers).

6.476 In relation to scope, we have found that the free switching programmes are for customers that are switching cloud providers, not ongoing multi-cloud use. We consider the exclusion of ongoing multi-cloud egress fees to be a substantial limitation in the scope of the programmes.

6.477 We have also found that AWS' and Google's programme terms would not provide sufficient credits to cover both the transfer of all data once from a customer's account or single service (for a partial switch) and multiple transfers of data due to temporary multi-cloud during switching. Some customer evidence has suggested switching involves periods of multi-cloud running,¹⁷⁴⁹ so we consider the

¹⁷⁴⁸ [Vodafone's response to the CMA's working papers](#), page 1.

¹⁷⁴⁹ The Jigsaw report found that switching typically involved running two sets of services concurrently, then switching off the legacy system when ready. Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw Research (2024). [X] said that a switch would involve 'leaving the old environment running for a while in case there are problems with the new one'. Note of meeting with [X]; [X] response to Ofcom's information request [X].

insufficiency of credits for this to also be a limitation to the scope of AWS' and Google's programmes.

- 6.478 Microsoft's programme does not allow for partial switching, only complete account switches. We consider the ability for a customer to switch part of its service to another provider to provide a similar competitive constraint to a customer's ability to enact a complete switch. We therefore also consider this to be a limitation in the scope of Microsoft's programme.
- 6.479 In addition to the above limitations, we have found that Microsoft's, AWS' and Google's free switching programmes are all limited to a 60-day time period. The 60-day time limit means the free switching programmes are likely to be ineffective in practice for many customers. Customer evidence suggests switching usually takes longer than 60 days,¹⁷⁵⁰ and any extension of the free switching programmes beyond 60 days is entirely at the respective cloud provider's discretion. As at 3 April 2025, AWS' uptake data also shows that [redacted] of the customers approved as eligible for its free switching programme have applied for extensions.¹⁷⁵¹
- 6.480 We have found that the cloud providers have a substantial degree of discretion over the application of the programmes (for example, over extensions to the 60 day time limit; in Google's programme, over the eligibility of partial switches; in Microsoft's programme, over the determination of whether required processes have been followed for a customer to receive credits).¹⁷⁵² We also consider that the programmes have restrictions on service or data eligibility,¹⁷⁵³ and all three providers' programmes require customers to submit their switching applications beforehand.
- 6.481 AWS submitted that the restrictions in its free switching programme exist as protective measures to prevent abuse and to provide clarity and certainty to AWS and to customers.¹⁷⁵⁴ Whilst cloud providers may consider it necessary to have measures in the programmes to limit potential customer misuse, we consider these measures may impact the effectiveness of the programmes in enabling a customer to switch, and it is not clear to what extent customer misuse is a risk or would impact other customers' ability to switch using the programmes. AWS was unable to provide evidence of the customer behaviour(s) that the protective

¹⁷⁵⁰ Responses to the CMA's information requests [redacted] Notes of meetings with [redacted]; [Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw \(2024\)](#), paragraphs 3.6.5, 3.6.7, 3.6.8; and page 43.

¹⁷⁵¹ CMA analysis of [redacted] submission to the CMA [redacted]; [redacted] response to the CMA's information request [redacted]. Note, Microsoft and Google did not provide data on the number of extension requests received for their free switching programmes.

¹⁷⁵² AWS' website FAQs also note that 'AWS may make changes with respect to free data transfers out to the internet at any time' and Microsoft's website notes 'Azure might make changes regarding the egress credit policy in the future'. [Amazon EC2 FAQs – AWS; Cancel and delete your Azure subscription - Microsoft Cost Management, accessed on 15 May 2025.](#)

¹⁷⁵³ Types of service and data eligible and/or amount of data stored to be eligible.

¹⁷⁵⁴ [AWS' response to the CMA's PDR dated 28 January 2025](#), paragraph 55.

measures in the AWS programme are seeking to prevent (we consider AWS' submission in more detail in appendix N).

- 6.482 Microsoft, AWS and Google submitted that there has been low uptake of their free switching programmes and that this shows that egress fees are not a meaningful barrier to switching.¹⁷⁵⁵ It is difficult to interpret what programme uptake data means for barriers to customer switching, given that 'low' uptake could suggest that customers do not want and/or need the free switching programmes, but could also mean that customers are unable or deterred from using the programmes (for example, because of eligibility restrictions), or face other barriers to switching that would also need to be addressed to enable switching. We also do not expect lowering barriers to switching to necessarily translate into actual switching; it can also work to enhance the threat of switching and thereby increase the competitive pressure on providers.
- 6.483 Following the PDR, we collected updated data from Microsoft, AWS and Google on their programme uptake since introduction.¹⁷⁵⁶ Our analysis of this data is set out in more detail in appendix N.
- 6.484 We found that AWS has approved [X] switching requests; Microsoft has had [X] complete a switch under its programme; and Google has approved [X] switching requests.¹⁷⁵⁷ AWS has received [X] requests for its programme and determined [X] of these requests to be [X];¹⁷⁵⁸ and Google received [X] requests for its programme, of which [X] of these requests have not been approved (either cancelled by customer or under ongoing review) and [X] requests were made in error by customers who wanted to close their billing accounts but had no data to transfer out as part of closing down their accounts.¹⁷⁵⁹
- 6.485 We therefore found that uptake of the free switching programmes has been low relative to the size of these providers' customer bases, although the number of requests received (of which a number have been determined to be ineligible or in error) by AWS and Google may indicate some broader customer interest in switching. We also identified a number of potential explanations for this low uptake, such as other factors (eg technical barriers) remaining to deter customer

¹⁷⁵⁵ AWS' response to the CMA's information request [X]; [AWS' response to the CMA's PDR dated 28 January 2025](#), paragraph 54; [Microsoft's response to the CMA's PDR dated 28 January 2025](#), paragraph 99; Google's response to the CMA's information request [X].

¹⁷⁵⁶ Microsoft provided the number of customers that had completed a switch up to the date of its information request response of 7 April 2025. Google provided the number of customer applications and approvals between January 2024 (when its programme was introduced) and 28 February 2025. AWS provided the number of customer applications and approvals as well as information the number and reasons for ineligible applications and the number of applications pending further information, for the period from its programme launch to 3 April 2025.

¹⁷⁵⁷ Responses to the CMA's information requests [X].

¹⁷⁵⁸ The most common reason for [X] in the data provided was because the [X] requested a [X] of the [X]. CMA analysis of AWS' response to the CMA's information request [X].

¹⁷⁵⁹ Google's response to the CMA's information request [X]. Google said that [X] of the [X] requests it received were made in error by customers who wished to close their billing accounts but had no data to transfer out which could indicate switching interest or switches that have already taken place.

switching or the programme requirements and restrictions, in line with the challenges for interpretation noted above.

6.486 The information on programme uptake is inconclusive and does not support the view that egress fees are not a barrier to switching.

6.487 We also considered the nature of Microsoft's, AWS' and Google's free switching programmes. The free switching programmes have been voluntarily introduced in the UK by Microsoft, AWS and Google, in response to the EU Data Act and have only been in place since January/March 2024. Given their voluntary nature, they could change at any time, as the requirements set out in the EU Data Act only apply to egress fees with respect to customers in the EU.

6.488 The three main parties made submissions on the voluntary nature of the programmes:

(a) Google submitted that it would not be viable to retrospectively narrow the applicability of its free switching programme now that free switching is widespread, because the reputational and commercial damage to its business would be severe.¹⁷⁶⁰

(b) AWS submitted that a loss of public face or potential revenue would make it very difficult for them to remove or narrow the programmes outside of the EU.¹⁷⁶¹

(c) Microsoft submitted that it is [REDACTED].¹⁷⁶²

6.489 AWS and Google submitted that a loss of public face or potential revenue would make it very difficult for them to remove or narrow the programmes outside of the EU.¹⁷⁶³ Our view is that:

(a) First, as the programmes have been voluntarily introduced and have only been in place since January/March 2024, we consider it would be relatively easy for cloud providers to change or remove their applicability in the UK.¹⁷⁶⁴

(b) Second, whilst risk of reputational damage may provide an incentive for cloud providers to retain the programmes, it will not necessarily prevent the narrowing of the programmes or imposition of further requirements and restrictions. We note that Microsoft's and AWS' website terms specifically

¹⁷⁶⁰ [Google's response to Egress fees working paper](#), Annex response (b).

¹⁷⁶¹ Transcript of hearing with AWS 2 July 2024 [REDACTED].

¹⁷⁶² Microsoft's response to the CMA's information request issued 17 March 2025, paragraph 22.

¹⁷⁶³ [Google's response to Egress fees working paper](#), Annex response (b); Transcript of hearing with [REDACTED].

¹⁷⁶⁴ We note also that there have been very few applications to the programmes from UK customers: AWS received [REDACTED] UK applications and Google received [REDACTED] UK applications. CMA analysis of AWS and Google's responses to the CMA's information requests [REDACTED].

note that changes could be made to the programmes,¹⁷⁶⁵ and we consider that only customers that are applying or undergoing a switch under the programmes are likely to have a detailed knowledge of current programme terms (and uptake data indicates this to be a relatively low number of customers). We therefore do not consider reputational damage to constitute a sufficiently strong deterrent against the narrowing of the programmes. We also note that, given that any narrowing of programme terms or geographic application may be followed by other main cloud providers, the impact of any reputational damage arising from this on revenues may be limited.

- 6.490 For the above reasons, we consider that the introduction of the providers' free switching programmes does not remove egress fees as a barrier to customer switching and multi-cloud use. We consider that the scope limitations and restrictions in the free switching programmes limit their application and impact. The low uptake of the programmes may be on account of other factors and therefore cannot be interpreted as evidence that egress fees are not a barrier or concern for customers. Finally, we note the voluntary nature of the programmes would enable the providers to change or narrow them at their discretion.

Our assessment of the relevance of egress fees to customers' switching and multi-cloud decisions

- 6.491 Egress fees have some influence on customers' decisions and represent a barrier to switching and/or multi-cloud. Switching from one provider to another that has a better offering may be harder for customers as a result of, among other things, higher costs of making a switch.¹⁷⁶⁶ Egress fees are therefore likely to disincentivise and/or hinder customers who might be considering to, or attempting to, switch and/or multi-cloud. We summarise our reasoning below.
- 6.492 Although our customer research suggests that egress fees are typically not seen as the most important factor in decisions relating to switching/multi-cloud, some customers, competitors and industry bodies explicitly identified egress fees as a challenge when considering, or a factor disincentivising, switching.
- 6.493 In addition, some customers said that egress fees had been a challenge to setting up multi-cloud architectures and/or took them into consideration when taking their decision to multi-cloud. These views were mirrored in the Jigsaw report, where some customers expressed concerns about egress fees for switching and/or multi-cloud.

¹⁷⁶⁵ AWS' website FAQs also note that 'AWS may make changes with respect to free data transfers out to the internet at any time' and Microsoft's website notes 'Azure might make changes regarding the egress credit policy in the future'. [Amazon EC2 FAQs – AWS; Cancel and delete your Azure subscription - Microsoft Cost Management](#), accessed on 15 May 2025.

¹⁷⁶⁶ CC3 (Revised), paragraphs 297 and 316.

- 6.494 When assessing customer views, although egress fees are not referred to as a key concern for a majority of customers, we consider they can still act as barrier to switching and/or multi-cloud as they affect some customers' decisions to do so.
- 6.495 Our analysis of these scenarios indicates that some customers may incur meaningful costs from switching and multi-cloud. Specifically:
- (a) our estimates of 'one-off' switching costs indicates that, in order to move all their data from one cloud provider to another, [0-5]% to [5-10]% of customers would have had to pay egress fees of more than [5-10]% of their annual cloud spend per year if they spread the cost over the course of a typical contract. If they do not spread the cost over multiple years, around [10-20]% to [20-30]% of customers would have had to pay egress fees of more than [5-10]% of their annual cloud spend in a year.
 - (b) the smallest customers face higher egress fees than other customers as a proportion of their annual cloud spend in order to switch cloud providers. For example, over the course of a typical contract, customers with annual spends of £1,000-£10,000 pay egress fees which are on average between [0-5]% more per year than those in other spend bands.
 - (c) this cost could influence the decision on whether or not to switch cloud provider. Around [0-5]% to [5-10]% of customers would not have the incentive to switch to a rival cloud provider even if, over the course of a typical contract, that provider offered a price that was [5-10]% better or equivalently better quality or functionality, thereby reducing the level of competition in the markets.
 - (d) our estimates of multi-cloud costs show that customers spending more than £10,000 face non-negligible egress fee costs as a percentage of total monthly cloud spend across different data transfer amounts. For example, for the median customer in each spend band over £10,000, egressing 20% of their data would cost [0-5]% to [10-20]% of their total monthly cloud spend.
 - (e) we have found that whilst the smallest customers (those who spend between £1,000 -£10,000) tend to be covered by their cloud provider's free tier when transferring up to [10-20]% to [30-40]% of their data, they face non-trivial fees for egressing amounts greater than this. For example, to egress 50% of their data, the median customer in this spend band would face an average cost of between [0-5]% and [5-10]% of its monthly spend in egress fees alone. We have also found that some of the smallest customers face meaningful egress fees.
 - (f) these costs could make it more difficult for a customer to multi-cloud or increase the level of integration of an existing multi-cloud architecture. For

reasons discussed in detail in appendix M, we do not consider that Google's submission on multi-cloud costs¹⁷⁶⁷ shows that egress fees are not a barrier to multi-cloud.

- 6.496 Our evidence from customers and from our analysis of customers' costs in switching and multi-cloud scenarios together indicate that there are some customers for whom egress fees are a more significant barrier to switching and multi-cloud. In particular, our analysis of switching and multi-cloud scenarios indicates that:
- (a) smaller customers are likely to face the highest egress fees as a proportion of their total cloud spend; and
 - (b) customers egressing larger proportions of their stored data are likely to pay higher egress fees as a proportion of their total cloud spend.
- 6.497 Internal documents of the main providers indicate that egress fees are likely to be a concern and a barrier for some customers and can be a factor which influences their behaviour and/or decisions and that these egress fees are most significant for those that multi-cloud or operate in high egress usage industries. Internal documents also indicate that there are strategic considerations in the setting of egress fee prices, including competitors' strategies and customers' behaviour.
- 6.498 While cloud providers have submitted that falling effective prices for egress are indicative of competition in the markets, we do not consider this is enough to conclude that the presence of egress fees cannot constitute a barrier to switching and multi-cloud. We do not find the data submitted on effective prices and price trends in general allows us to draw conclusions about competition in the markets.
- 6.499 We consider that analyses of the observed prevalence of egress fees are not wholly representative of customers' liability to pay them. This is because some customers, when faced with egress fees they consider too high, might take actions to avoid paying them. This would reduce the overall spend on egress fees but not the impact they have on customers' decision-making.
- 6.500 We consider that cloud providers' justifications for data ingress being free while data egress is paid, despite using the same fixed assets, are not relevant considerations as to whether egress fees are a barrier to switching and/or multi-cloud. Regardless of the validity of the explanations for egress being charged while ingress is not, we consider that these points do not contribute to our overall conclusion.

¹⁷⁶⁷ Google's submission to the CMA [X].

- 6.501 We have found that Microsoft's, AWS' and Google's free switching programmes have scope limitations and restrictions, including that the programmes generally do not cover multi-cloud use and are limited to a 60-day switching period, as well as cloud providers having a substantial degree of discretion over the application of the programmes. We also do not consider that the relatively low uptake of the programmes can be interpreted as evidence that egress fees are not a barrier or concern for customers. We therefore consider that these programmes do not impact the conclusions of our analysis of either switching costs or multi-cloud costs.
- 6.502 We focus much of our analysis on Microsoft, AWS and Google as they collectively represent the majority of revenue in the IaaS and PaaS markets. Whilst we note that egress fees for other cloud providers vary, we consider that the findings of our analysis would still apply more widely to any cloud providers that charge egress fees, as any level of egress fees can be expected to represent a barrier to switching and/or multi-cloud. However, the impact of any such barrier is likely to vary with the magnitude of the egress fees charged for switching and/or multi-cloud, the size of the provider and with the scale of customer data transfers.

Potential benefits of egress fees

- 6.503 Cloud providers submitted that there are some benefits from egress fees, notably that:
- (a) egress fees fund investment and innovation;
 - (b) cloud providers pass on cost savings to customers; and
 - (c) egress fees deter inefficient egress usage and thus avoid inefficient network investment.
- 6.504 We considered each of these arguments to determine if there is evidence of benefits arising from egress fees, and if so what the nature of the benefits are and whether this indicates that there are rivalry-enhancing efficiencies which may remove or mitigate the impact of egress fees on customers' decisions to switch and/or multi-cloud.¹⁷⁶⁸
- 6.505 The submissions received on potential benefits generally related to egress fees in general, rather than switching and/or multi-cloud egress fees specifically. Our analysis in this section therefore refers to egress fees in general, except where switching and multi-cloud egress is specifically contemplated.

¹⁷⁶⁸ In some circumstances, the positive effects of efficiencies on competition associated with a particular market feature may outweigh the harmful effects of that feature. Efficiencies can enhance rivalry ('rivalry-enhancing efficiencies') when they induce one or more firms to follow a course of action of benefit to customers (eg lowering prices or increasing innovation) in response to actual or expected actions by rivals. [CC3 \(Revised\)](#), paragraph 174.

Role of egress fees in funding investment

6.506 In this section we set out cloud providers' submissions that egress fees fund investment and innovation. We then set out our analysis of the role of egress fees in funding investment, where we examined the evidence on investment and innovation.

Cloud provider views on investment and innovation

6.507 Microsoft, AWS and Google, IBM and Oracle submitted that egress fees recover costs.¹⁷⁶⁹ Microsoft, AWS and Google also submitted that data transfer fees, including egress fees, enable and/or incentivise their investment in networking services and innovation:

- (a) AWS submitted that data transfer fees have enabled it to invest in high-quality proprietary infrastructure to provide a premium service, and AWS' continuous investment in physical infrastructure has yielded benefits to customers. AWS submitted that as a result of its investments, customers experienced a faster and more reliable network. AWS said that regulation which prevents its ability to make these investments would lead to reduced innovation and poorer outcomes for UK businesses.¹⁷⁷⁰
- (b) AWS submitted that it needs to retain the possibility of earning a return on its investment to justify the investments and innovations in its network. AWS submitted that removing its ability to earn profit would remove any incentive to invest and innovate in the network in the way that has yielded substantial benefits for customers and the economy more generally.¹⁷⁷¹
- (c) Microsoft submitted that in order to provide competitive data transfer services, Microsoft's annual investments have been comprehensive and extended beyond data transfer connectivity to improve, maintain and innovate in the data transfer services that customers experience.¹⁷⁷² Microsoft submitted that cloud providers must be allowed to recoup their costs and return on investment in order to retain incentives for cloud providers to invest and innovate in these data transfer services.¹⁷⁷³ Microsoft highlighted that without the ability to charge egress fees, it would lost the

¹⁷⁶⁹ [AWS response to the Issues Statement](#) dated 17 October 2023, paragraphs 24 and 29; [Microsoft's response to the Issues Statement](#) dated 17 October 2023, paragraph 38; [Google's response to the Issues Statement](#) dated 17 October 2023, paragraph 26; [Google's response to the Egress fees working paper](#), paragraph 22; [IBM response to the Issues Statement](#) dated 17 October 2023, page 3; [Oracle, consolidated commentary to the CMA's Updated Issues Statement, its working papers, and the Market investigation qualitative customer research final report](#), page 1.

¹⁷⁷⁰ [AWS' response to the CMA's updated issues statement and working papers](#), paragraph 32; [AWS' response to the CMA's PDR dated 28 January 2025](#), paragraph 51; AWS' submission to the CMA [§].

¹⁷⁷¹ [AWS' response to the CMA's Egress fees working paper](#), 31 July 2024, paragraph 11.

¹⁷⁷² [Microsoft's response to the Competitive landscape, Committed spend agreements and Egress fees working papers](#), footnote 80.

¹⁷⁷³ [Microsoft's response to the Competitive landscape, Committed spend agreements and Egress fees working papers](#), paragraphs 97-99.

incentive to develop and offer premium bandwidth services that cater to the needs of fewer but more demanding customers.¹⁷⁷⁴

- (d) Google submitted that data transfer fees (not egress fees specifically) allow it to continue to invest in developing a range of networking products as well as maintaining, expanding and upgrading its network infrastructure. Google submitted that this, in turn, generates significant benefits for customers by fostering a competitive market that provides greater choice and gives customers credible alternatives to the two large incumbent cloud providers.¹⁷⁷⁵ Google submitted that a market-wide intervention in egress fees puts the customer benefits of investment in networking products and infrastructure at risk and competitive market at risk.¹⁷⁷⁶
- (e) Google submitted that its networking prices reflect the quality and innovation of its networking products, as well as its cost and ongoing investment.¹⁷⁷⁷ Google also noted that removal of egress fees risks decreased innovation and product quality.¹⁷⁷⁸

6.508 Microsoft, AWS and Google each submitted examples of business developments to support their statements regarding investment and innovation:

- (a) AWS submitted that its most significant investment of recent years consisted of designing and rolling out its own hardware and software which forms part of the AWS global network and that it has made significant investments in backbone and metro fibre.¹⁷⁷⁹ AWS also cited the development of [REDACTED].¹⁷⁸⁰
- (b) Microsoft submitted that it has made significant investments in its Premium Global Network, which will allow it to build low-latency infrastructure at higher costs and technical performance. Microsoft cited its cumulative number of feature updates for its top 20 UK cloud services as evidence of innovation, but also submitted that much of the innovation in technology happens outside of Microsoft, AWS and Google and often in the 'commons' (notably, open source).¹⁷⁸¹
- (c) Microsoft also submitted that many cloud providers, including Microsoft, AWS and Google, have invested billions of pounds in building out extensive and

¹⁷⁷⁴ Microsoft's response to the CMA's PDR dated 28 January 2025, paragraph 107.

¹⁷⁷⁵ Google response to the Egress fees working paper, Annex response (h).

¹⁷⁷⁶ Google's response to Egress fees working paper, Annex response (h).

¹⁷⁷⁷ Google's response to Egress fees working paper, paragraph 3(c).

¹⁷⁷⁸ Google response to the Egress fees working paper, paragraph 57.

¹⁷⁷⁹ AWS' submission to the CMA in response [REDACTED].

¹⁷⁸⁰ AWS' submission to the CMA in response [REDACTED].

¹⁷⁸¹ Microsoft's response to the Competitive landscape, Committed spend agreements and Egress fees working papers, paragraphs 12 and 45 and footnote 115.

sophisticated private networking infrastructure spanning the globe for the benefit of their customers and to supplement the public network.¹⁷⁸²

- (d) Google submitted that it has made and continues to make significant investments in its high-quality global fibre-optic software-defined infrastructure network, citing its presence in over 200 countries and territories, including 187 network edge locations, 40 regions, 120+ zones, 113 interconnect locations and 14 sub-sea cables. Google also submitted that it offers innovative networking products such as Cross-Cloud Interconnect and a range of quality and service level options, including networking product add-ons such as additional security and Standard Tier and Premium Tier networking products.¹⁷⁸³

Analysis of the role of egress fees in funding investment

- 6.509 Microsoft, AWS and Google submitted that data transfer fees, including egress fees, enable and/or incentivise their investment in networking services and innovation.¹⁷⁸⁴ We considered if there is evidence of customer benefits arising from egress fees in the form of investment and innovation and if so, whether there is evidence that the investment for this is dependent on egress fees, such that egress fees should be considered the key driver of the benefits.
- 6.510 We note that even if benefits are identified, such benefits would need to be found to be rivalry-enhancing efficiencies in order for them to potentially remove or mitigate the impact of egress fees on customers' decisions to switch and/or multi-cloud.¹⁷⁸⁵

Analysis of potential benefits

- 6.511 Based on our analysis of Microsoft's, AWS' and Google's submissions,¹⁷⁸⁶ we have found that these cloud providers have made substantial investments in their global network infrastructure. These investments will have benefit to customers of cloud services broadly in the quality and global reach of the network used to provide cloud services.
- 6.512 In terms of innovation, the examples cited by Microsoft, AWS and Google that appear potentially to be innovation are: AWS' development of an automated traffic

¹⁷⁸² Microsoft's response to the Issues statement dated 17 October 2023, paragraph 36.

¹⁷⁸³ Google's response to the Egress fees working paper, paragraphs 19 and 20.

¹⁷⁸⁴ AWS' response to the CMA's updated issues statement and working papers, paragraph 32; AWS' response to the CMA's Egress fees working paper dated 23 May 2023, paragraphs 5 and 11; Microsoft's response to the Competitive landscape, Committed spend agreements and Egress fees working papers, paragraph 9; Google response to the Egress fees working paper, Annex response (h)

¹⁷⁸⁵ CC3 (Revised), paragraph 174.

¹⁷⁸⁶ AWS' response to the CMA's updated issues statement and working papers, paragraph 33; Microsoft's response to the Competitive landscape, Committed spend agreements and Egress fees working papers, footnote 115; Google's response to the Egress fees working paper, paragraphs 19 and 20; [X] submission to Ofcom [X].

optimisation and management service for network traffic between AWS and the internet;¹⁷⁸⁷ Microsoft's investments allowing it to build low-latency infrastructure at higher costs and technical performance and feature updates for its top 20 UK cloud services;¹⁷⁸⁸ and Google's networking products such as Cross-Cloud Interconnect.¹⁷⁸⁹

6.513 In response to our requests for information on investments that cloud providers have made in improving the quality of customer data transfer services and/or innovation that cloud providers have made in relation to customer data transfer services, we received the following submissions:

- (a) in its response, AWS referred to the largest capital infrastructure investments in AWS' Public Cloud Infrastructure Services business that it submitted to the CMA.¹⁷⁹⁰ Our review of these investments identified that some are for network assets, but none appear to be investments specific to providing, maintaining or improving egress data transfer services.¹⁷⁹¹
- (b) Microsoft submitted that to provide competitive data transfer services, including those in the UK, its annual investments need to be comprehensive and need to extend beyond data transfer connectivity. Microsoft said these investments aim to improve, maintain and innovate the data transfer services that customers experience, but noted that while these services are visible to UK customers, the investments are global in nature and not limited to data transfer.¹⁷⁹²
- (c) Google submitted that [REDACTED]. Google submitted that more broadly, it invests in a premium network infrastructure globally to ensure that it provides the best standard of networking experience for customers. Google also cited its introduction of new and enhanced data transfer products such as Cross-Cloud Interconnect (a direct connection service), Private Service Connect (a virtual private cloud service) and Media CDN (for streaming video to end users).¹⁷⁹³ We note that these are different products to the egress data transfer via the internet, which is the focus of our analysis.

6.514 It is difficult to assess innovation and we had limited detail and evidence from these submissions to determine the extent to which there is network innovation to the benefit of customers. However, we consider that Microsoft's, AWS' and

¹⁷⁸⁷ AWS' submission to the CMA [REDACTED].

¹⁷⁸⁸ [Microsoft's response to the Competitive landscape, Committed spend agreements and Egress fees working papers](#), footnote 115; [Microsoft's response to the Competitive landscape, Committed spend agreements and Egress fees working papers](#), paragraphs 12 and 45 and footnote 115; [Google's response to the Egress fees working paper](#), paragraph 19.

¹⁷⁸⁹ [Google's response to the Egress fees working paper](#), paragraph 20.

¹⁷⁹⁰ AWS' response to the CMA's information request [REDACTED].

¹⁷⁹¹ AWS response to the CMA's information request [REDACTED].

¹⁷⁹² Microsoft's response to the CMA's information request [REDACTED].

¹⁷⁹³ Google's response to the CMA's information request [REDACTED].

Google's submissions indicate that they have made substantial investments and potentially some innovation in their network overall, which may benefit customers in the terms of the quality of network used to deliver cloud services.

Extent to which investment and innovation is dependent on egress fees

- 6.515 Of the submissions that Microsoft, AWS and Google have cited as evidence of investment in egress services or data transfer services more broadly, we noted that:
- (a) For AWS' submissions,¹⁷⁹⁴ our understanding is that hardware and software for AWS' global network and backbone and metro fibre are common infrastructure used across data transfer services and potentially more widely across cloud services and investments in this infrastructure will therefore be funded not just by egress fees.
 - (b) Microsoft's cited investment in its global network and Microsoft's feature updates for its top 20 UK cloud services relate to cloud services generally.¹⁷⁹⁵
 - (c) Google's cited investment in its global network and product innovations mostly relate to other networking services rather than egress and we consider its Standard and Premium tier services to be a price differentiation offer, rather than evidence of innovation.¹⁷⁹⁶
- 6.516 The submissions that we set out above as potential evidence of innovation were also largely for network investment and data transfer services overall. Only AWS' [REDACTED] seems likely to directly benefit egress data transfer services as it [REDACTED] (although this would likely also apply to ingress data transfers and potentially other services).
- 6.517 Given that the investments and innovation that Microsoft, AWS and Google submitted largely all related to the provision of their broader networking services (and/or cloud services overall), we consider this investment to be funded by these services overall.¹⁷⁹⁷ Whilst investment in providing a 'high quality' network will benefit egress data transfer services, our analysis has not indicated that egress data transfer services are the main user or beneficiary of this. For example, Microsoft's, AWS' and Google's internal data transfer volumes are [REDACTED] their

¹⁷⁹⁴ AWS submitted that egress fees have enabled it to invest in high-quality proprietary infrastructure, but also that it used the profits that it generates from data transfer fees (not egress fees specifically) to invest in infrastructure. AWS' submission to the CMA [REDACTED].

¹⁷⁹⁵ [Microsoft's response to the Competitive landscape, Committed spend agreements and Egress fees working papers](#), paragraphs 12 and 45 and footnote 115; [Google's response to the Egress fees working paper](#), paragraph 19.

¹⁷⁹⁶ [Google's response to the Egress fees working paper](#), paragraph 20.

¹⁷⁹⁷ Our analysis of cloud provider revenues indicates that egress fees contribute [0-5]%-[0-5]% of AWS', Microsoft's and Google's UK cloud revenues. Source: CMA analysis of AWS, Microsoft and Google data. Responses to the CMA's information requests [REDACTED].

egress volumes (excluding CDN egress, as this is a type of egress with different pricing and use case).¹⁷⁹⁸ [X].¹⁷⁹⁹ AWS' CDN egress volumes are [X] its non-CDN egress volumes.¹⁸⁰⁰

- 6.518 This is supported by our analysis of cloud providers' costs, as set out in appendix Q, which has found that there is shared network infrastructure for providing data transfer services (as well as wider cloud services in the case of some infrastructure, such as data centres) and network investment will be used by both ingress and egress services and to a large extent internal data transfers. Microsoft, AWS and Google each allocate a share of costs associated with network and wider cloud infrastructure to data transfer services.¹⁸⁰¹ The costs of connectivity to the public internet (ie internet transit and peering costs) are the only costs commonly identified as being specific to egress and reflect operating costs rather than investment in innovation. Therefore, any investment is being funded by data transfer and other cloud services more broadly and is not seen to be investment in assets that are solely or specifically for egress.
- 6.519 In our review of Microsoft's, AWS' and Google's internal documents (see appendix P for more detail), we did not see evidence that cloud providers considered future investment or innovation in making their data transfer pricing decisions.
- 6.520 We consider that the investment and innovation cloud providers have cited is not dependent on egress fees, such that egress fees should be considered the key driver of any associated benefits. For the same reasons, we do not consider there to be rivalry-enhancing efficiencies from egress fees in relation to investment and innovation.
- 6.521 We also do not agree with the implicit argument made by some cloud providers that profits generated from egress fees are necessary to incentivise investment in innovation. For a provider to invest in improving the quality of different components of a wider service or product offering, it is not necessary for a fee to be charged for each of those components in order for investment to arise. The observation that there are investments does necessarily mean that this came out of funding raised by egress fees, or that egress fees were necessary for the investment to arise. This is especially true when the infrastructure is used for multiple purposes, not just egress. We also note that egress data transfer is one service of many provided by Microsoft, AWS and Google and so not the only source of profits.¹⁸⁰²

¹⁷⁹⁸ CMA analysis of responses to the CMA's information requests [X].

¹⁷⁹⁹ Microsoft's response to the CMA's information request [X].

¹⁸⁰⁰ AWS' response to the CMA's information request [X]. AWS' response to Ofcom's information request [X].

¹⁸⁰¹ Google only allocates costs to its 'Networking' segment overall, which is [X]. Google's response to the CMA's information request [X].

¹⁸⁰² We have assessed overall cloud business profitability in chapter 3, market outcomes section.

Role of egress fees in passing on cost savings

6.522 In this section we set out cloud providers' submissions that egress fee prices have decreased over time and cost savings have been passed on to customers. We then set out our analysis of the extent to which customers are benefitting from pass-on cost savings as a result of egress fees.

Cloud provider views on pass-on cost savings

6.523 As set out below, AWS has submitted that cost savings resulting from investment are passed on to customers through lower egress fee prices. Google has submitted that [REDACTED] and Microsoft and Google both submitted that egress fee prices have decreased.

6.524 Oracle submitted that data transfers between cloud providers via peering incur lower costs but these cost savings are not passed on to the customer in cloud providers' egress fees.¹⁸⁰³

AWS

6.525 AWS submitted that it has used the profits that it generates from data transfer fees to invest in infrastructure which ensures that AWS' network is of a premium quality and ultimately results in lower prices for customers. As set out earlier in this chapter, AWS submitted as evidence data which it said showed that the effective price (ie net of discounts) charged for egress has decreased globally by 37% and by 25% for UK customers, between 2018 and 2023.¹⁸⁰⁴ AWS also submitted that this data showed that from 2019 to 2023 [REDACTED].¹⁸⁰⁵ AWS also noted that its egress margin [REDACTED], consistent with a pass-on rate exceeding [REDACTED].¹⁸⁰⁶ AWS further submitted that its egress prices have decreased by almost 40% with a pass-on rate of over 100% between 2019 and 2024.¹⁸⁰⁷

6.526 AWS cited its investments in designing and rolling out its own hardware and software across the AWS business and investments in proprietary networking solutions, backbone and metro fibre and terrestrial and undersea cable as investments which have achieved cost efficiencies.¹⁸⁰⁸

6.527 AWS submitted that its investment, and the resulting cost savings, have allowed AWS to expand its network to serve more customers in more locations and handle a significant increase in volume of network traffic, without a significant increase in

¹⁸⁰³ Oracle's response to the Updated Issues Statement and working papers, pages 1-2.

¹⁸⁰⁴ AWS' response to the CMA's Egress fees working paper issued 23 May 2024, paragraph 5.

¹⁸⁰⁵ AWS' submission to the CMA [REDACTED].

¹⁸⁰⁶ AWS' submission to the CMA [REDACTED]; AWS' response to the CMA's PDR dated 28 January 2025, paragraph 51.

¹⁸⁰⁷ AWS' submission to the CMA [REDACTED].

¹⁸⁰⁸ AWS' submission to the CMA [REDACTED]; AWS' response to the CMA's updated issues statement and working papers, paragraph 33.

cost. AWS said that these cost savings have been passed on to customers, for example through the expansion of AWS' free tier for egress, [REDACTED] for customers [REDACTED] and the introduction of free data transfer for customers who switch.¹⁸⁰⁹

- 6.528 AWS submitted that it is difficult to prove that the cause of its egress price and cost decreases is specifically investments funded by egress revenues, but also that the Egress fees working paper does not prove that egress price and cost decreases are not due to investments funded by egress revenues. AWS submitted that (i) the current structure provides AWS with the ability and incentive to invest in its network to improve its quality and efficiency; (ii) investments are costly and need to be funded; (iii) positive margins on egress allow AWS to fund such investments; (iv) AWS' investments have resulted in higher efficiencies and lower egress costs; and (v) lower egress costs have resulted in lower egress prices. AWS submitted that this indicates a well-functioning structure, irrespective of whether it can be established that these positive outcomes result from investments funded by egress revenues specifically.¹⁸¹⁰
- 6.529 AWS also submitted that AWS' investments into its network enables it to receive and deliver traffic close to the source and destination and thereby allows AWS to [REDACTED].¹⁸¹¹

Microsoft

- 6.530 As set out earlier in this chapter, Microsoft submitted that the list prices for its egress products (Bandwidth and ExpressRoute¹⁸¹²) have fallen over time, and submitted the introduction of internet egress routed via public ISP network, a price decrease for ExpressRoute in 2021 and other prices remaining stable, and the expansion of the Bandwidth free tier.
- 6.531 Microsoft also submitted that on average, real prices for both metered and unmetered egress plans have fallen since 2018.¹⁸¹³

Google

- 6.532 As set out earlier in this chapter, Google submitted that the effective price paid for standard internet transfers has been trending down over time. Google submitted that most of its networking SKUs have either experienced no price change or an effective price decrease. Google cited its introduction of a free tier across all Standard Tier SKUs in October 2023, keeping the prices of all other tiers the

¹⁸⁰⁹ AWS' submission to the CMA [REDACTED].

¹⁸¹⁰ AWS' submission to the CMA [REDACTED].

¹⁸¹¹ AWS response to the CMA's information request [REDACTED].

¹⁸¹² 'Bandwidth' offers services to move data in (ingress), within (internal) and out of Azure data centres (egress). ExpressRoute is a private connection service that connects Microsoft data centres to customers' on-premises or colocation facility infrastructure.

¹⁸¹³ Microsoft submission to the CMA [REDACTED].

same, and its free tier for standard internet transfers which it said is more extensive than other free tiers.¹⁸¹⁴

6.533 Google submitted that [REDACTED].¹⁸¹⁵

Analysis of pass-on cost savings

6.534 In this section, we have examined the extent to which there may be a benefit from cost savings being passed on to customers.

6.535 AWS submitted that cost savings resulting from investment are passed on to customers through lower egress fee prices and submitted data analysis to support this. Microsoft and Google submitted that egress fee prices have decreased and Microsoft provided some pricing data to support this.

6.536 For there to be a benefit arising from egress fees for this, there would need to be evidence that the charging of egress fees funds investment, which results in cost savings, which are then passed on to the customers. We examined these submissions to consider whether customers are benefitting from pass-on cost savings as a result of egress fees.

Extent to which there are cost savings arising from egress fee-funded investment

6.537 As set out above, we consider that the investment and innovation cloud providers have cited is not dependent on egress fees and therefore not the key driver of any investment or innovation benefits.

6.538 If broader network investment, which egress fees have contributed to, has resulted in cost savings which are passed on in egress fee prices, this could potentially be relevant to consider. We note that AWS has submitted that its network investments have achieved cost efficiencies but also acknowledged that it is [REDACTED].¹⁸¹⁶

6.539 AWS submitted cost data as evidence that decreases in average costs of providing egress have been passed on. We do not consider this to show that there have been cost savings arising from egress-fee funded investment. Average costs could decrease in the ordinary course as a result of economies of scale. We also note that, having analysed the underlying data from the AWS submission for 2019 to 2023, a [REDACTED] of the costs that AWS include in their analysis of the movement of

¹⁸¹⁴ [Google's response to the CMA's Egress fees working paper](#), paragraph 15.

¹⁸¹⁵ Google's submission to Ofcom [REDACTED].

¹⁸¹⁶ AWS' submission to the CMA [REDACTED]. [AWS' response to the CMA's updated issues statement and working papers](#), paragraph 33; AWS' submission to the CMA [REDACTED].

average egress costs per unit compared to average prices [REDACTED].¹⁸¹⁷ These costs therefore do not reflect cost efficiencies or savings being passed on to customers but simply reflect [REDACTED].

- 6.540 Having reviewed AWS' data further, we have not found that AWS has lowered egress prices for customers in proportions equal to its efficiency gains. Excluding the costs allocated based on revenue contribution, our analysis has found that AWS' average egress revenues per GB (which it refers to as the 'effective price') do not move in a clear pattern in line with average egress costs. Our analysis of AWS' egress margins, both excluding the costs allocated based on revenue contribution and including all costs as submitted by AWS, has found that AWS' egress margins (as a percentage of revenues) were [REDACTED] in years 2019 to 2022, with margins [REDACTED] in 2023 and 2024.¹⁸¹⁸
- 6.541 AWS also submitted that its investments in its network [REDACTED].¹⁸¹⁹ This could represent cost savings arising from investment; however, as set out above, we do not consider such network investment to be dependent on or primarily driven by the charging of egress fees. [REDACTED] and [REDACTED].¹⁸²⁰
- 6.542 Google's example of [REDACTED].¹⁸²¹ We do not find this to be cost savings arising from investment.
- 6.543 Similarly, as noted earlier in the chapter, Microsoft's data on real prices for egress falling over time simply reflects that inflation has been positive (and significant in 2022) whilst nominal prices have remained constant. This does not provide any evidence of cost savings arising from investment.
- 6.544 We consider that there is no clear link between cloud providers charging egress fees to customers and investments which deliver egress cost savings.

Extent to which cost savings arising from investment are passed on

- 6.545 Microsoft, AWS and Google submitted that egress fee prices have decreased. As set out earlier in this chapter, we identified some issues with the data Microsoft and AWS submitted on egress fee prices,¹⁸²² but irrespective of this we do not consider falling prices to be evidence of a benefit arising from egress fees unless there is evidence that the lower prices resulted from egress fees.

¹⁸¹⁷ CMA analysis of AWS' submission to the CMA dated 31 July 2024, [REDACTED] 'Analysis of DTO costs prices and margins'; AWS response to the CMA's information request [REDACTED].

¹⁸¹⁸ CMA analysis of AWS' submission to the CMA dated 31 July 2024, [REDACTED] 'Analysis of DTO costs prices and margins'; AWS response to the CMA's information request [REDACTED].

¹⁸¹⁹ AWS response to the CMA's information request [REDACTED].

¹⁸²⁰ See appendix Q for further details. Source: CMA analysis of responses to the CMA's information requests [REDACTED]

¹⁸²¹ Google submission to Ofcom [REDACTED].

¹⁸²² For example, AWS' and Microsoft's 'price' data uses average revenue net of discounts which includes broader discounts such as CSAs. This means decreasing average revenues could reflect the growing application of broader discounts, rather than any particular changes in egress fee prices.

- 6.546 We would expect prices to fall in the presence of economies of scale, so we do not consider evidence of falling prices to necessarily indicate that there are efficiencies in the form of cost savings being passed on to customers arising specifically from the charging of egress fees.¹⁸²³ As set out earlier in the chapter on AWS' submission that there are diseconomies of scale in egress data transfers, we consider that once investment has occurred, there would be minimal incremental costs per additional egress data transfer, and it is unclear why maintenance or other costs would scale increasingly per egress data transfer once the additional investment has been made.¹⁸²⁴
- 6.547 As set out above, we have not seen evidence of cost savings arising from egress fee-funded investment and therefore do not consider there to be evidence that any decrease in Microsoft's, AWS' and/or Google's egress fee prices has been the passing on of cost savings resulting from investment.
- 6.548 We also note that Microsoft, AWS and Google have each cited the increase or introduction of free tiers as evidence of price reductions. Increased free tiers (and otherwise stable prices) will reduce the average price paid by customers for egress. However, we do not see a causal link between the charging of egress fees and these increases in the free tiers available. We also note that, as set out earlier in the chapter, in the last 5 years Microsoft and AWS have each only changed the free tier available to customers once and the list prices for Microsoft's, AWS' and Google's egress fees have largely not changed in the past 5 years.
- 6.549 Additionally, we note that AWS cites [redacted] for customers [redacted] and the introduction of free data transfer for customers who switch as evidence of passing on costs, whilst Microsoft cites the introduction of its cheaper ISP network routing option as evidence of reduced prices. As discussed above, we have found the free switching programmes introduced by Microsoft, AWS and Google to be in response to the EU Data Act and have not seen any evidence from our review of internal documents relating to the decisions to introduce these programmes that indicates that the programmes arose from cost savings. Additionally, [redacted] and Microsoft's cheaper routing option appear to be business responses to customer demand for lower egress pricing, rather than reflecting cost savings passed on. We also note that [redacted] are not necessarily permanent price changes for customers.
- 6.550 We consider that there is no clear evidence of a customer benefit arising from the charging of egress fees funding investment, which results in cost savings, which are then passed on the customers. For the same reasons we do not consider

¹⁸²³ We also note that while falling prices may be indicative of competition in a market, prices trending downwards does not mean that egress fees cannot still constitute a barrier to multi-cloud and switching.

¹⁸²⁴ Note, we analyse the cost of egress data transfers in general, including providers' submissions on incremental costs, in appendix Q.

there to be rivalry-enhancing efficiencies in the form of lower prices arising from the charging of egress fees.

Role of egress fees in deterring inefficient egress usage

6.551 In this section we set out cloud providers' submissions that egress fees deter inefficient egress usage. We then set out our analysis of the role of egress fees in deterring inefficient egress usage, in particular the extent to which there are possible benefits arising from this.

Cloud provider views

6.552 Microsoft, AWS and Google and IBM submitted that removing egress fees could result in inefficient and excessive egress usage by customers.¹⁸²⁵

6.553 AWS submitted that without egress fees, customers transferring data out of the AWS network would have no reason to care about efficient network usage or architecture, as there would not be a financial component to the transfer.¹⁸²⁶ [X]. AWS also submitted that customers' increased use may also be inadvertent, as without an egress fee customers may ignore or remove any alerts for volume of network use, making it easy for customers to accumulate higher network use without realising it.¹⁸²⁷

6.554 Microsoft submitted that inefficient usage is an incontrovertible consequence when the price for a service that provides intrinsic value to customers is reduced to zero, and that customers would have no incentive to be efficient in their data transfer practices, highlighting that customers would be incentivised to use the highest performing service, such as premium routing, that is available for free.¹⁸²⁸

6.555 Another cloud provider submitted that the existence of egress fees is a key tool to incentivise customers to design efficient architectures and that banning egress fees risks increasing the number of inefficient IT architectures and volume of unnecessary egress.¹⁸²⁹

6.556 IBM submitted that designing an efficient IT infrastructure requires some effort from a customer and it is very easy to accumulate a high volume of egress data very quickly. [X]. IBM noted that customers' need for an efficient IT infrastructure

¹⁸²⁵ AWS response to the Issues Statement dated 17 October 2023, page 14; Google response to the Egress fees working paper, Annex response (h); Microsoft's response to the Competitive landscape, Committed spend agreements and Egress fees working papers, paragraph 9; [X] submission to the CMA [X].

¹⁸²⁶ AWS' response to the CMA's Egress fees working paper dated 23 May 2024, paragraph 18.

¹⁸²⁷ AWS' submission to the CMA [X].

¹⁸²⁸ Microsoft's response to the CMA's PDR dated 28 January 2025, paragraph 109.

¹⁸²⁹ [X] submission to the CMA [X].

is only relevant to multi-cloud use, as this is less relevant when switching cloud provider entirely.¹⁸³⁰

6.557 Cloud providers identified the following as potential impacts of inefficient or excessive egress usage:

- (a) Google submitted inefficient usage of finite networking capacity and resources would negatively affect the availability of network capacity for customers across the industry and would also unnecessarily give rise to a greater carbon footprint and related harm to the environment.¹⁸³¹
- (b) AWS submitted that there may be incentives to encourage efficient network use (such as latency, security, data governance) for some sophisticated customers, but that its view was that free egress would encourage far greater use of cloud providers' networks, increasing strain on these.¹⁸³²
- (c) AWS also submitted that eliminating egress fees makes it cheaper for malicious actors to conduct Distributed Denial of Service ('DDoS') attacks, creating a very significant risk to the Internet more broadly. AWS said that if usage were free, anyone could generate huge amounts of traffic in order to overwhelm websites and online services for users.¹⁸³³
- (d) Microsoft submitted that banning or restricting egress fees could lead to data resilience security risks arising from the already significantly high and increasing volume of data traffic via cloud infrastructure.¹⁸³⁴

6.558 In response to the PDR, AWS submitted that the comparison to the egress free tier made in the PDR fails to acknowledge the fundamental difference between a service that is completely free and a service that is free up to a given limit, for which it said customers' incentives are likely to be entirely different. AWS said that the PDR's comparison to ingress usage also fails to understand that ingress and egress are used very differently by customers to serve different needs, with ingress typically a one-off event whereas customers may egress for a variety of reasons and may export data multiple times. AWS also submitted that the main risk is from customers that transfer huge amounts of data daily.¹⁸³⁵

6.559 In relation to the PDR's reference to some cloud providers offering free egress in some cases, AWS submitted that this pricing discrepancy stems from differences in business and investment strategy and reflects the costs providers face as a result of the size and quality of their networks. AWS submitted that it has made

¹⁸³⁰ [IBM's response to Egress fees working paper](#), page 3.

¹⁸³¹ [Google response to the Egress fees working paper](#), Annex response (h).

¹⁸³² AWS' submission to the CMA [redacted].

¹⁸³³ [AWS' response to the CMA's Egress fees working paper dated 23 May 2024](#), paragraph 19.

¹⁸³⁴ [Microsoft's response to the competitive landscape, committed spend agreements and egress fees working papers](#), paragraph 95.

¹⁸³⁵ AWS' submission to the CMA on [redacted].

substantial investments in its proprietary network, whereas other providers rely heavily on transit providers to enable data transfers.¹⁸³⁶

Our analysis of inefficient egress risk

- 6.560 In this section we have considered first the extent to which inefficient egress usage by customers is a potential risk for cloud providers to seek to deter. We then considered whether there are benefits arising from egress fees deterring inefficient egress usage.

Extent to which inefficient egress usage by customers is a potential risk

- 6.561 We recognise that each cloud provider's network capacity will generally be a finite resource and that inefficient or excessive egress usage could take up more of this finite network capacity.
- 6.562 Whilst cloud providers' submissions related to egress fees in general, our analysis has focused on egress fees incurred during switching and/or multi-cloud.
- 6.563 The risk of inefficient or excessive egress usage for customers that are switching is likely to be low given that customers' switching egress will not be indefinite; some customer data may only need to be egressed once (although we have found that some customers may require a transition period of using multiple clouds); and the broader costs involved in switching (eg time and resource) mean customers are unlikely to switch frequently and should be incentivised to complete their switch as quickly and efficiently as possible. We therefore consider this to potentially be a risk relevant only to multi-cloud egress usage and unlikely to be relevant to switching.
- 6.564 Cloud providers' submissions have been largely based on inefficient egress usage as an anticipated potential risk, rather than evidenced issues and costs arising under a counterfactual where egress fees are nil.¹⁸³⁷
- 6.565 AWS submitted the example [redacted].¹⁸³⁸ This provides a piece of evidence of costs deterring inefficient network usage. However, we also note that this is an example of the behaviour of an AWS internal service rather than a customer. This may be indicative of the behaviour of an egress user, but we also note that there could be other factors influencing an internal service's behaviour which are different to a

¹⁸³⁶ AWS' submission to the CMA on [redacted].

¹⁸³⁷ Note: In the [Egress fees working paper](#), paragraph 4.52, we invited evidence from cloud providers on inefficient egress usage by customers to address (a) how likely it may be that customers would continue to have an incentive to design their IT architectures efficiently and avoid significant volumes of unnecessary egress, for example for latency, availability, security and/or data governance reasons; (b) how easy or difficult it may be for a customer to inadvertently accumulate a high volume of egress via the internet; and (c) whether there is any existing evidence of free egress leading to inefficient egress usage, given that several cloud providers currently provide a free tier of egress per month.

¹⁸³⁸ AWS' submission to the CMA [redacted].

customer. For example, [REDACTED] may have been incentivised to have data traffic growth as a measure of business performance.

- 6.566 We also note that this is an example where [REDACTED].¹⁸³⁹ While relevant, for the reverse (eg egress fees being removed) this would mean customers would be actively changing from efficient to inefficient egress, which is not necessarily the same. AWS also does not mention any other internal services as having [REDACTED] when charged at a fixed rate. Additionally, whilst AWS submitted that the ‘example shows that, when AWS [REDACTED],¹⁸⁴⁰ AWS does not actually state that there were any [REDACTED] as a result of CloudFront’s approach, just that it was [REDACTED].
- 6.567 Microsoft and Google did not submit any evidence of inefficient or excessive egress usage. We do not accept Microsoft’s submission that inefficient usage is an incontrovertible consequence when the price for a service that provides intrinsic value to customers is reduced to zero, as no evidence was provided to support this statement and, as noted below, we consider that customers would have some non-cost incentives to have efficient data practices.
- 6.568 We found some evidence in internal documents that Microsoft and AWS consider data transfer overall and switching egress respectively to have a potential risk of abuse by some customers. In our review of internal documents we found that Microsoft considered ‘potential abuser traffic and security risks’ to be a factor when [REDACTED].¹⁸⁴¹ [REDACTED].¹⁸⁴²
- 6.569 We note that AWS submitted that the main risk is from customers that transfer large amounts of data.¹⁸⁴³ Whilst larger customers may represent a larger cost risk to providers for any increased egress usage, it is not clear why they would be more likely to inefficiently egress, or how this affects our assessment. We also consider this to be somewhat inconsistent with AWS’ previous submission that there may be incentives to encourage efficient network use (such as latency, security, data governance) for some sophisticated customers, which we generally expect to be larger customers.
- 6.570 If egress fees deter inefficient or excessive egress usage, by the same logic we would also expect to see instances of inefficient or excessive networking usage where there is no cost constraint. We have not seen evidence that there is currently inefficient or excessive customer egress usage up to the monthly free tier that Microsoft, AWS and Google provide. We accept that customer behaviours could to some extent be different for a service that is completely free compared to a service that is free up to a given limit. However, we would expect the incentives

¹⁸³⁹ AWS’ submission to the CMA [REDACTED].

¹⁸⁴⁰ AWS’ submission to the CMA [REDACTED].

¹⁸⁴¹ Microsoft’s submission to the CMA [REDACTED].

¹⁸⁴² AWS’ response to the CMA’s information request [REDACTED].

¹⁸⁴³ AWS’ submission to the CMA [REDACTED].

to use the free service to work in similar ways, and for there to be some evidence of distortion or skewing of usage around the free tier boundary.

- 6.571 The risks to network quality raised also seem applicable to all data traffic, ie egress as well as ingress and internal data transfers, given all require and use network capacity. We are not aware of any evidence that customers have had inefficient or excessive ingress since cloud providers made ingress free. We note that AWS submitted that ingress and egress are used very differently by customers, with ingress typically a one-off event whereas customers may egress data multiple times.¹⁸⁴⁴ While this may be true, the nature of egress involving multiple transfers does not make these transfers inefficient. Whilst this could mean there is more opportunity for potential inefficient usage for egress, it is not clear why this would mean there would not be any inefficient ingress, and hence why the lack of evidence of inefficient or excessive ingress is irrelevant.
- 6.572 We also note that this risk would appear to contradict Microsoft's and Google's submissions that egress fees are not a key driver of customer behaviour.¹⁸⁴⁵
- 6.573 Additionally, some cloud providers offer free egress in some cases, which indicates that there is unlikely to be an excessive or prohibitive risk associated with this. Members of the Bandwidth Alliance¹⁸⁴⁶ such as Oracle, Alibaba, Tencent and Scaleway provide free egress for shared customers that route transfers via Cloudflare's network (Cloudflare itself does not charge egress fees). This suggests that these cloud providers do not consider cross-cloud egress to have a significant risk of inefficient egress usage. Similarly, Civo has removed all egress fees and OVHcloud includes egress in the price of instances.¹⁸⁴⁷ We would expect smaller cloud providers such as Civo and OVHcloud to have less ability to absorb the impact of inefficient egress usage, and therefore would seek to deter this if it was a significant risk.
- 6.574 We accept that, as submitted by AWS, pricing differences between providers may reflect differences in business and investment strategy, with some providers relying more on transit providers and potentially having lower costs for providing an overall network as a result.¹⁸⁴⁸ However, we do not consider this to mean that inefficient egress would not be a concern for these providers. We would still expect inefficient egress usage by customers to have potential risks and effects on such

¹⁸⁴⁴ AWS' submission to the CMA [38].

¹⁸⁴⁵ [Microsoft's response to the competitive landscape, committed spend agreements and egress fees working papers](#), paragraph 83; [Google's response to the Egress fees working paper dated 23 May 2024](#), paragraph 34

¹⁸⁴⁶ The Bandwidth Alliance is a group of cloud providers and networking companies that have agreed to discount or waive data transfer fees for shared customers, for data transfers routed via Cloudflare's network. [The Cloudflare Blog](#), accessed on 27 February 2024.

¹⁸⁴⁷ [Civo abolishes all egress fees to advance cloud simplicity strategy](#) (accessed on 24 September 2024).

OVHcloud charges egress fees for storage services but includes egress in the price of instances on all locations, except the Asia-Pacific region. [Price list: a comparison of our Public Cloud offers - OVHcloud UK](#), accessed on 22 October 2024.

¹⁸⁴⁸ AWS' submission to the CMA [38]. Note, we considered the extent to which investment in network infrastructure relates specifically to the provision of egress or data transfer services more widely in appendix Q.

providers' networks, for example impacting the amount of data traffic capacity needing to be procured from transit providers. As such, we considered the pricing of other providers to be relevant.

- 6.575 We consider there to be some non-cost incentives for customers to design their IT architectures efficiently and avoid significant volumes of unnecessary egress, for example for latency, availability, security and/or data governance reasons. As noted above, AWS acknowledges that this may be the case for 'some sophisticated customers' but is of the view that customers' increased egress usage may be inadvertent.¹⁸⁴⁹ We considered AWS' suggestion that customers would actively ignore or remove any alerts for volume of network use to be unsubstantiated, as we are not aware of evidence that customers regularly ignore or remove usage alerts.
- 6.576 We consider that there may be some potential risk of inefficient or excessive use of network capacity by customers, but that this risk applies to all networking service usage eg egress as well as ingress and internal data transfers. As customers have non-cost incentives to efficiently egress and we have not seen evidence of customers inefficiently or excessively using ingress or the egress free tier, and some cloud providers are currently willing to provide free egress, we consider the risk of inefficient or excessive egress usage for switching to be low and for multi-cloud use to be unclear but have mitigating factors.

Analysis of potential benefits

- 6.577 As noted above, cloud providers' network capacity is a finite resource. Inefficient egress usage by customers could potentially require cloud providers to invest more in capacity and/or network resiliency to handle unpredictability resulting from inefficient or excessive egress usage. Therefore, we considered whether there is a customer benefit, by incentivising efficient egress usage that allows providers to avoid unnecessary network investment and/or reducing the risk of reduced network quality (ie network availability and resiliency).
- 6.578 It is unclear if there are direct benefits to customers from any deterrent effect from egress fees. The potential impacts that cloud providers have identified are all based on egress fees decreasing potential risks to network quality, through deterring inefficient and excessive egress usage.¹⁸⁵⁰ The main benefit of this would seem to be potentially allowing cloud providers to avoid additional network investment and/or further developing tools to manage data traffic unpredictability and network security. Avoiding unnecessary investment and/or negative impacts on network quality may benefit customers but it is unclear at what level of

¹⁸⁴⁹ AWS' submission to the CMA [38].

¹⁸⁵⁰ [Google response to the Egress fees working paper](#), Annex response (h); [AWS' response to the Egress fees working paper](#), paragraphs 17-19; [Microsoft's response to the competitive landscape, committed spend agreements and egress fees working papers](#), paragraph 95.

inefficient egress such impacts would materialise, and as set out above, we considered the risk of inefficient egress to be low for switching and unclear for multi-cloud use with mitigating factors.

- 6.579 We would also expect cloud providers to be efficiently investing in network capacity to meet customer demand and providing sufficient network quality, rather than this necessarily being a specific benefit to customers, arising from egress fees.
- 6.580 In addition, any benefits from avoiding unnecessary network investment and/or reducing the risk of reduced network quality do not arise from egress fees alone. All data transfer services generally use network capacity and would be capable of inefficient or excessive usage by customers. Therefore, the impacts identified by cloud providers and any indirect benefits resulting from the cost constraint of fees should be applicable to all data transfer services, and networking services more broadly. It is also not clear that the concerns about significant increases in data traffic and potential DDoS attacks are particularly relevant to cloud-to-cloud egress (the relevant type of egress for switching and multi-cloud), rather than other types of egress to end users.
- 6.581 Cloud providers also appear to use other tools to deter or manage inefficient or excessive egress usage and any resulting impacts. AWS' submission that 'customers may ignore or remove any alerts for volume of network use',¹⁸⁵¹ shows that volume alerts can be used to deter usage. A Microsoft internal document notes that:
- [REDACTED]¹⁸⁵²
- 6.582 As such, to the extent that egress fees might act as a deterrent to inefficient or excessive egress usage, they do not appear to be the sole deterrent and hence not the sole cause of any resulting benefits. The deterrent effect seems likely to benefit customers only against a counterfactual where there is extremely inefficient or excessive egress usage. Given that the inefficient egress risk for switching is likely to be low and for multi-cloud use is unclear and has mitigating factors, in the absence of egress fees for switching and multi-cloud use, the other existing deterrents should continue to provide any associated benefits and we would not expect control limits, eg speed throttling, to need to materially change.
- 6.583 We have not found there to be any rivalry-enhancing efficiency from egress fees potentially deterring inefficient egress usage, given that there is no indication that

¹⁸⁵¹ AWS' submission to the CMA [REDACTED].

¹⁸⁵² Microsoft's response to the CMA's information request [REDACTED].

egress fees themselves are inducing providers to deter inefficient egress usage, in response to actual or expected actions by rivals.¹⁸⁵³

Our assessment of the potential benefits of egress fees

- 6.584 There may be some overall network quality benefits to customers arising from cloud providers' network investment which egress fees could help fund along with the fees charged by cloud providers for their other services. However, we have not found any investment and innovation cited to be dependent on, or solely driven by, the charging of egress fees.
- 6.585 In relation to Microsoft's, AWS' and Google's submissions that egress fee-related cost savings are passed on to customers, we considered that, irrespective of any decrease in prices, there is no clear evidence that there is a customer benefit arising from the charging of egress fees. We also do not consider there to be sufficient evidence of customer benefits arising from egress fees deterring inefficient or excessive egress usage and have found limited evidence that inefficient or excessive egress is a risk, in the absence of egress fees.
- 6.586 Overall, we have not found there to be customer benefits clearly resulting from the charging of egress fees or evidence of rivalry-enhancing efficiencies arising from egress fees funding investment, lowering prices and/or deterring inefficient egress, and therefore do not consider there to be benefits which remove or mitigate egress fees' impact on customer switching and multi-cloud use.

Conclusions on egress fees

- 6.587 The presence and relevance of egress fees to customers' decisions on switching and multi-cloud means that there is a weakened customer response to differences in price, service quality and/or innovation between cloud providers.
- 6.588 This means that customers have less ability and/or incentive to switch away from their current cloud provider or choose a rival cloud provider as part of a multi-cloud architecture, even if the rival is offering a better service for their needs. In other words, the benefits of switching and multi-cloud must be greater than they would need to be in the absence of egress fees. As a result, egress fees, alongside other barriers to switching and multi-cloud, contribute to a degree of 'lock-in', where customers are less able to switch cloud provider, or use multiple cloud providers, once they have made their initial choice for cloud services.
- 6.589 While cloud providers have introduced a range of free switching programmes for egress, these all have scope limitations and restrictions, including that the programmes generally do not cover multi-cloud use and are limited to a 60-day

¹⁸⁵³ CC3 (Revised), paragraph 174.

switching period, as well as cloud providers having a substantial degree of discretion over the application of the programmes. We do not consider the relatively low uptake of the programmes to be evidence that egress fees are not a barrier or concern for customers and note the voluntary nature of the programmes would enable the providers to change or narrow them at their discretion. They do not impact our analysis of egress fees as a barrier to switching and/or multi-cloud use.

- 6.590 We have not found any customer benefits clearly resulting from the charging of egress fees or evidence of rivalry-enhancing efficiencies arising from egress fees as:
- (a) the investment and innovation cited by cloud providers does not appear to be dependent or solely driven by the charging of egress fees;
 - (b) there does not appear to be a clear link between cloud providers charging egress fees and having egress cost savings which are passed on to customers; and
 - (c) there is limited evidence that inefficient or excessive egress is a risk, and there does not seem to be a clear direct benefit to customers arising from egress fees deterring inefficient or excessive egress usage. Any deterrent effect is also not solely caused by egress fees.
- 6.591 Our conclusion is that egress fees reduce the ability of and/or incentives for customers to switch to other cloud providers and/or run a multi-cloud architecture. They also reduce the incentives of providers to compete for the business of their competitors.

Conclusions on barriers to switching and multi-cloud

- 6.592 Cloud providers do not have clear incentives to enable customers to switch and use multiple clouds. Providers may be incentivised to reduce these barriers if that means they are seen to be relatively easy to switch to and from or multi-cloud with and this will lead customers to be more willing to choose them as a provider in order to limit any long-term lock-in effects. However, when reducing the barriers to switching and multi-cloud, customers' ability to switch and multi-cloud increases and so does their bargaining power.
- 6.593 The conclusions of our assessments are set out in the relevant sections above. In summary:
- (a) Technical barriers raise the cost to customers of switching and using multi-cloud; and

- (b) Egress fees are relevant to customers' decisions to switch and multi-cloud and are therefore likely to reduce the ability and/or incentives for customers to switch to other cloud providers or run a multi-cloud architecture.

6.594 Commercial barriers (in the form of egress fees) are barriers to switching and multi-cloud. Together, these barriers restrict customers from exercising effective choices about their cloud providers and affect customers' ability to either switch providers or use multiple providers.

7. Licensing

- We have examined whether Microsoft's licensing practices are influencing customers' choice of cloud provider, and whether this harms competition by preventing AWS and Google from competing as effectively for customers that use Microsoft software on cloud.
- We have assessed: (i) Microsoft's market power in the relevant markets for Windows Server, SQL Server, Windows 10/11, Visual Studio and its productivity suites; (ii) the importance of Microsoft's software as inputs in cloud services; (iii) Microsoft's licensing practices (price and non-price practices); (iv) the impact on rivals' competitive offerings arising from Microsoft's licensing practices and whether this is adversely impacting competition in cloud services.
- Microsoft has significant market power in relation to the relevant software products we have considered. This is because Microsoft has moderate to high market shares in all of these products and customers generally are unable or unwilling to switch away from these products.
- We have considered the importance of the Microsoft software inputs, both in terms of their significance to the cost base of AWS and Google and their potential to shape downstream competition for cloud services. Windows Server alone, and when combined with SQL Server, accounts for a material proportion of costs for all sizes of customer and are therefore important inputs.
- We considered any differences in the quality of Windows Server and SQL Server when used on Azure compared to non-Azure clouds. Most customers are not aware of any differences, and for these customers, the inputs do not shape downstream competition beyond their significance in the cost base. However, for customers that do consider quality factors of Microsoft software, the Microsoft software may be a particularly important input in terms of affecting the overall quality or attractiveness of AWS' and Google's competitive offerings.
- Windows 10/11 and the Microsoft productivity suites are important to the provision of cloud-based VDI services, while Visual Studio is an important input for particular customers. While VDI workloads currently make up a relatively small proportion of public cloud workloads and the evidence is mixed as to whether VDI will increase in importance over time, these three software products do not constitute important inputs on their own but do contribute to the overall importance of Microsoft software products as inputs to cloud services.
- We have examined the differences between the software that Microsoft offers Azure customers and that which it supplies to its rivals in cloud services, both in terms of price and quality factors. The input cost for AWS and Google is higher than the customer-

facing price that Microsoft charges its own customers that qualify for AHB to use Windows Server on Azure. For SQL Server, the input cost for AWS and Google is higher than Microsoft's PAYG customer-facing price.

- Further, Microsoft does not make Windows 10/11 and Microsoft 365 available to AWS and Google through their respective SPLAs, and customers with existing licences cannot bring these to AWS and Google (except for specific Microsoft 365 licences to Amazon Workspaces). Customers without existing licences are unable to purchase Microsoft Office or Visual Studio IDE on Google's cloud services, nor can they purchase Windows 10/11 and Microsoft 365 on Google or AWS.
- Evidence from AWS and Google is consistent with high input costs resulting in them charging a higher price for cloud services using Microsoft software and having a less competitive offering. Despite difficulties with estimating and interpreting profit margins for these services on AWS and Google, we consider that our estimates are consistent with partial foreclosure.
- In addition, Windows Server and SQL Server are used disproportionately on Azure compared to on AWS and Google. We consider that the very large difference is at least partly because some customers' choice of cloud is influenced by Microsoft's licensing practices.
- Microsoft's licensing practices with respect to the productivity suites, Windows 10/11 and Visual Studio are likely to affect competition for VDI workloads. While VDI represents a relatively small proportion of usage of all cloud services, we consider that Microsoft's licensing practices relating to these software products have an adverse effect on the competitiveness of AWS' and Google's offerings.
- Considering all of the evidence in the round, Microsoft's licensing practices are adversely impacting the competitiveness of AWS and Google in the supply of cloud services, particularly in competing for customers that purchase cloud services which use the relevant Microsoft software as an input.
- As a result, Microsoft faces weaker competitive constraints from AWS and Google, its most significant competitors, which is reducing competition in cloud services markets.

Background

- 7.1 We received submissions regarding Microsoft's software licensing practices that allege that these practices may make it less attractive for customers to use licensed software products on the cloud infrastructure of rival providers.
- 7.2 The statutory task of this market investigation is to decide whether any relevant feature, or combination of features, prevents, restricts or distorts competition, in connection with the supply or acquisition of goods or services in the UK or a part

of the UK pursuant to the EA02 and not whether there has been an infringement of the Competition Act 1998 (CA98). The process is 'investigative and inquisitorial, not accusatorial', so involvement in a market investigation or being subject to remedies does not imply that participants have infringed the law or are suspected of wrongdoing.¹⁸⁵⁴

- 7.3 Microsoft provides both software products and the public cloud services that host them and is therefore vertically integrated. Despite the potential of vertical relationships to enhance efficiency and consumer welfare, they can also sometimes lead to an AEC in a market, particularly by allowing firms to foreclose, fully or partially, rivals' access to inputs and customers and/or otherwise have a dampening effect on competition.¹⁸⁵⁵
- 7.4 As customers have migrated software systems to the cloud, licensing arrangements have evolved in different ways. Some customers have been able to use their existing on-premises licences to use the relevant software in the cloud. Other customers have had to procure a new licence to use the software in the cloud.
- 7.5 We have focused on Microsoft's software licensing practices because:
- (a) the majority of the concerns raised in submissions we have received relate to Microsoft; and
 - (b) as set out elsewhere in this report, Microsoft has significant market power in cloud services¹⁸⁵⁶ which raised a concern that its licensing practices could have an adverse effect on competition.
- 7.6 We have assessed five Microsoft software products as potentially relevant to our assessment: Windows Server (which includes Active Directory functionality), SQL Server, Windows 10/11, Visual Studio and Microsoft 365/Office.
- 7.7 We have focused on these five software products because licensing concerns were raised in relation to them. The evidence shows that customers who have workloads involving these software products, which have migrated to the public cloud, are more likely to 'lift and shift' those workloads compared to others, and continue using these software products in the cloud.¹⁸⁵⁷

¹⁸⁵⁴ CC3 (Revised), Guidelines for market investigations: Their role, procedures, assessment and remedies, paragraph 21.

¹⁸⁵⁵ CC3 (Revised), paragraph 267.

¹⁸⁵⁶ See chapter on Competitive Landscape, market power and concentration.

¹⁸⁵⁷ [redacted] submission to the CMA [redacted]; [redacted] submission to the CMA [redacted]; [redacted] response to Ofcom's information request [redacted]; Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), page 29; Lift and shift applications workloads.

Description of the software licensing practices

- 7.8 In this section, we provide some background on Microsoft's software licensing practices. We set out below descriptions of:
- (a) Microsoft's licensing practices including some changes over time; and
 - (b) the ways in which customers use Microsoft's software on the public cloud.

Timeline of licensing practices

- 7.9 While our analysis has focused on Microsoft's current licensing practices, it is useful to explain how these have developed.

Pre-2019

- 7.10 With certain specific exceptions, prior to October 2019, customers with perpetual on-premises licences for certain Microsoft software did not have the right to deploy these licences on the shared hardware (ie public cloud services) of non-Azure cloud providers on a 'bring your own licence' ('**BYOL**') basis.^{1858, 1859} However, in relation to dedicated hardware (ie private cloud services), customers could use their licences on any cloud on a BYOL basis, whether that was on Azure or the cloud infrastructure of a third party.¹⁸⁶⁰

2019 changes

- 7.11 Microsoft modified its licensing terms in 2019 when it created a new category of 'Listed Providers' of cloud infrastructure services.¹⁸⁶¹ These were Microsoft, Alibaba, Amazon, and Google. Customers of Listed Providers could now no longer use their pre-existing licences for Microsoft software on a BYOL basis on dedicated hardware of Listed Providers.^{1862, 1863}

¹⁸⁵⁸ [X] submission to the CMA [X].

¹⁸⁵⁹ For certain specific Microsoft products, such as SQL Server, customers that also purchased Software Assurance subscriptions with eligible on-premises perpetual Microsoft products could BYOL (BYOL is explained further below) their licence to dedicated or shared hardware (ie the public cloud) via Microsoft's 'License Mobility via Software Assurance' policy. See, [License Mobility & Software Assurance - Microsoft Volume Licensing](#), accessed on 18 November 2024. A cloud provider submitted that customers were also able to BYOL specific Microsoft products to the shared hardware of non-between 2013 and s that were part of Microsoft's Qualified Multitenant Hoster Program.

¹⁸⁶⁰ [X] submission to the CMA [X]. A cloud provider submitted that customers who BYOL their licences to non-Azure dedicated infrastructure are not necessarily able to carry over the same rights that they would be able to if they migrated to Azure.

¹⁸⁶¹ [X] submission to the CMA [X]; [Updated Microsoft licensing terms for dedicated hosted cloud services](#), accessed on 18 November 2024.

¹⁸⁶² [X] submission to the CMA [X]; [Updated Microsoft licensing terms for dedicated hosted cloud services](#), accessed on 18 November 2024; [X] submission to the CMA [X].

¹⁸⁶³ Nor could customers use their pre-existing licenses on shared hardware of Listed Providers, however this was the case prior to Microsoft's modification of its licensing terms. For completeness, customers could still use their pre-existing licence for certain products, such as SQL Server, that were eligible for License Mobility via Software Assurance, on a

- 7.12 Microsoft submitted that it updated its licensing terms to correct a ‘licensing loophole’.¹⁸⁶⁴ The ‘licensing loophole’ related to AWS and Google using customers’ outsourcing rights to effectively create a public cloud service on dedicated hardware (ie the private cloud).^{1865, 1866}
- 7.13 Microsoft said that its 2019 changes to its licensing terms were not applied to smaller cloud providers because their offering was closer to operating as a genuine outsourcing partner, and as such, they were a better fit for Microsoft’s outsourcing model.¹⁸⁶⁷
- 7.14 We have received submissions on how cloud providers were selected to be included as a Listed Provider and the likelihood of other providers being added:
- (a) CISPE said that these Listed Providers are unilaterally identified by Microsoft and include Microsoft’s major current competitors. It added that the list can be extended at Microsoft’s own discretion to add new competitors as and when they become a threat.¹⁸⁶⁸
 - (b) Google said that the list was relatively arbitrary,¹⁸⁶⁹ and that the Listed Providers were Microsoft’s largest competitors and likely among the most capable of competitively constraining Microsoft.¹⁸⁷⁰
- 7.15 Microsoft submitted that [REDACTED].¹⁸⁷¹ Microsoft added that other players do not come close to the scale of Listed Providers and [REDACTED].¹⁸⁷² In this regard, Microsoft said that, based on scale, [REDACTED].¹⁸⁷³
- 7.16 Microsoft also said it had [REDACTED] and that the declining importance of legacy software would make the Listed Provider status continue to decline in relevance in overall competitive dynamics between cloud providers.¹⁸⁷⁴

2022 changes

- 7.17 In the summer of 2021, Aruba S.p.A, OVHcloud, and the Danish Cloud Community complained to the European Commission that their customers faced higher prices and more licensing restrictions than Azure customers when trying to use

BYOL basis on either shared or dedicated hardware of a Listed Provider. The changes did not apply to Listed Provider customers with existing software licenses purchased before 1 October 2019. [REDACTED] submission to the CMA [REDACTED]; [REDACTED] submission to the CMA [REDACTED].

¹⁸⁶⁴ Microsoft’s submission to the CMA [REDACTED].

¹⁸⁶⁵ Microsoft’s submission to the CMA [REDACTED].

¹⁸⁶⁶ ‘Outsourcing rights’ refers to rights included in Microsoft’s software licensing terms which allowed customers to run their software on servers built, managed and run by ‘outsourcers’. Microsoft, submission to the CMA [REDACTED].

¹⁸⁶⁷ Microsoft’s submission to the CMA [REDACTED].

¹⁸⁶⁸ CISPE’s response to the Issues Statement [REDACTED].

¹⁸⁶⁹ Note of meeting with Google [REDACTED].

¹⁸⁷⁰ Google’s response to the CMA’s information request [REDACTED].

¹⁸⁷¹ Microsoft’s submission to the CMA [REDACTED].

¹⁸⁷² Microsoft’s submission to the CMA [REDACTED].

¹⁸⁷³ Note of meeting with Microsoft [REDACTED].

¹⁸⁷⁴ Microsoft’s submission to the CMA [REDACTED].

Microsoft's licensed software on their cloud infrastructure, and also that they could not use some versions of Microsoft's products.¹⁸⁷⁵

- 7.18 Microsoft said that, whilst it disagreed with the complaint, in response on 1 October 2022 it introduced licensing changes globally that were designed to enable customers to use subscription licences in any non-Listed Provider cloud free of any additional charge.¹⁸⁷⁶ In a blog post announcing the changes, Microsoft acknowledged that while not all of the European cloud providers' claims were valid, some of them were and that it would make changes soon to address them.¹⁸⁷⁷
- 7.19 One of the changes introduced by Microsoft was the 'Flexible Virtualisation Benefit', which enables customers of non-Listed Providers to use either their existing subscription or perpetual licences with Software Assurance on non-Listed Provider cloud infrastructure, whether dedicated or shared.¹⁸⁷⁸
- 7.20 Microsoft submitted to the CMA that these changes 'comprehensively resolved the concerns of all but the largest hyperscale cloud providers' and submitted that the complaint filed by Aruba, OVHcloud and the Danish Cloud Community had been withdrawn.¹⁸⁷⁹ Microsoft said that the changes amounted to granting like-for-like economics on Microsoft software whether used on Azure or on another non-Listed cloud provider.¹⁸⁸⁰
- 7.21 Another change that was introduced with the 2022 changes was that, from 1 October 2025 onwards, customers will no longer be able to buy and deploy Microsoft licences from independent managed service providers if those providers host their services on Listed Providers' clouds.¹⁸⁸¹

Using Microsoft software products on public cloud

- 7.22 This section explains the ways in which customers can use Microsoft's software on the public cloud, setting out: (i) the providers through which customers can obtain the rights to use Microsoft software on public cloud services; and (ii) the routes by which customers can obtain and deploy their rights to use Microsoft software on non-Azure clouds.

¹⁸⁷⁵ [REDACTED] submission to the CMA [REDACTED].

¹⁸⁷⁶ Microsoft's submission to the CMA [REDACTED]; See Microsoft's announcements: [Microsoft responds to European Cloud Provider feedback with new programs and principles - EU Policy Blog](#); and [New licensing benefits make bringing workloads and licenses to partners' clouds easier](#), accessed on 18 November 2024.

¹⁸⁷⁷ See, [Microsoft responds to European Cloud Provider feedback with new programs and principles - EU Policy Blog](#), accessed on 18 November 2024.

¹⁸⁷⁸ See, [New licensing benefits make bringing workloads and licenses to partners' clouds easier](#); [New options for partner hosted cloud](#), accessed on 8 May 2024.

¹⁸⁷⁹ Microsoft's submission to the CMA [REDACTED].

¹⁸⁸⁰ Microsoft's submission to the CMA [REDACTED].

¹⁸⁸¹ See, [New licensing benefits make bringing workloads and licenses to partners' clouds easier](#), accessed on 18 November 2024; [REDACTED] submission to the CMA [REDACTED].

The providers

7.23 There are a few providers through which customers can obtain the rights to use Microsoft software on the cloud:

- (a) directly from Microsoft;
- (b) through Cloud Solution Provider programme licensors (CSPPs) and CSP-Hosters; and
- (c) through non-Azure cloud providers or independent managed service providers.

Directly from Microsoft

7.24 There are two ways through which customers can obtain rights to use Microsoft software products on the cloud directly from Microsoft. The first is to buy a licence for the Microsoft products through Microsoft's volume licensing programme and to BYOL that licence to the cloud.¹⁸⁸² The second is to purchase cloud services directly on Azure, incorporating the Microsoft products – this option does not require a licence as the products on Azure will be 'license-included'.¹⁸⁸³

Cloud Solution Provider program

7.25 The Cloud Solution Provider program (CSPP) is a reseller programme that enables partners to sell licences to Microsoft cloud solutions on Azure (rather than the Microsoft products themselves).¹⁸⁸⁴ Over time, and subject to partner feedback, Microsoft has enabled Microsoft 365 Apps sold by cloud solution providers (CSPs) via this program to also be run in non-Listed Provider clouds.¹⁸⁸⁵ In addition, Microsoft has enabled the resale of on-premises licences by CSPs for certain Microsoft products, such as Windows Server, which a customer can BYOL onto the cloud.¹⁸⁸⁶

7.26 The CSP-Hoster programme is an expansion of the CSP programme, as it enables partners to pre-build hosted cloud desktop and server solutions that can be sold alongside licences in the CSP programme.¹⁸⁸⁷ These solutions can be licence-included hosted solutions offered through CSP or the opportunity for customers to BYOL their licences to access partner-provided solutions, but either

¹⁸⁸² Microsoft's response to the CMA's information request [38].

¹⁸⁸³ Microsoft's response to the CMA's information request [38].

¹⁸⁸⁴ Microsoft's response to the CMA's information request [38].

¹⁸⁸⁵ Microsoft's response to the CMA's information request [38].

¹⁸⁸⁶ Microsoft's response to the CMA's information request [38].

¹⁸⁸⁷ Microsoft's response to the CMA's information request [38].

way the licences for the Microsoft software products must belong to the end client.¹⁸⁸⁸

Alternative cloud providers or independent managed service providers

- 7.27 Customers that want to use Microsoft's software products on non-Azure clouds may purchase cloud services incorporating the Microsoft products through an alternative cloud provider (whether Listed or non-Listed).¹⁸⁸⁹ The alternative cloud provider acquires the licences to use the Microsoft software in its own cloud services through a Services Provider Licensing Agreement (SPLA).¹⁸⁹⁰
- 7.28 Alternatively, a customer may also purchase cloud services incorporating the Microsoft products in a similar way but through an independent managed service provider, which instead of hosting the cloud services on its own cloud infrastructure, hosts its services on another cloud provider's cloud.¹⁸⁹¹ As noted above, from 1 October 2025, independent managed service providers will not be able to host these licences on Listed Providers' clouds.¹⁸⁹²

The routes to obtaining the right to use the Microsoft software

- 7.29 There are two possible routes to obtaining the right to use Microsoft's software on non-Azure clouds:
- (a) bring your own licence (BYOL); and
 - (b) a SPLA.
- 7.30 The main differences between the two routes are explained below.

BYOL

- 7.31 BYOL is a term used when a customer relies on their existing on-premises Microsoft product licence or subscription licence to deploy the Microsoft product on the cloud (whether Azure, non-Azure, Listed, non-Listed, public or private).
- 7.32 Since 2022, for the majority of Microsoft's products, customers with a relevant software subscription licence or a licence with active Software Assurance, are able

¹⁸⁸⁸ Microsoft's response to the CMA's information request [§].

¹⁸⁸⁹ [§] response to Ofcom's information request [§]; [§] submission to the CMA [§].

¹⁸⁹⁰ Microsoft's response to Ofcom's information request [§] Microsoft's submission to the CMA [§]; Microsoft's response to the CMA's information request [§].

¹⁸⁹¹ [§] submission to the CMA [§].

¹⁸⁹² [§] submission to the CMA [§]; [New licensing benefits make bringing workloads and licenses to partners' clouds easier](#), accessed on 18 November 2024.

to use those on-premises Microsoft product licences on non-Listed Providers' public or private clouds (ie they can make use of BYOL).¹⁸⁹³

- 7.33 BYOL rights are included with certain subscription licences, which includes licences that are either covered by Microsoft Software Assurance or for which the underlying licence right is itself a subscription.¹⁸⁹⁴ Software Assurance is a Microsoft subscription offering that customers can add to their underlying licences for certain benefits, including licence mobility, which allows customers to use their licences via the BYOL route on non-Listed Provider clouds.¹⁸⁹⁵

SPLA

- 7.34 Microsoft's SPLA programme provides cloud providers with the right to integrate certain Microsoft products into their own cloud services and offer those cloud services to their end customers directly.¹⁸⁹⁶ The licence purchased under the SPLA covers the right to use the software on the hardware that the service provider uses to provide their services to their end customers.¹⁸⁹⁷ From Microsoft's perspective, the cloud provider is Microsoft's customer – the cloud provider pays Microsoft for its usage monthly in arrears based on how much Microsoft software the cloud provider actually used. In turn, the cloud provider charges its own end customer. SPLA is not a reseller programme for Microsoft software.¹⁸⁹⁸
- 7.35 The routes available to customers (SPLA and BYOL) to obtain the right to use Microsoft's software on non-Azure clouds (Listed and non-Listed Providers) depend on two variables: (i) the software product; and (ii) the cloud provider the customer wishes to use. The different routes available per relevant software product is discussed where relevant in this chapter.

Differences between using Microsoft software products on Azure compared to on non-Azure clouds via SPLA

- 7.36 We received submissions setting out issues and concerns relating to Microsoft's software licensing practices, including both price and non-price factors.
- 7.37 In relation to the price factors, the submissions set out that it is more expensive to use the relevant Microsoft products on Listed Providers' clouds compared to Azure as a result of BYOL restrictions (whereby customers cannot typically BYOL to Listed Providers' clouds). As such, where a customer is unable to use the BYOL

¹⁸⁹³ Microsoft's response to Ofcom's information request [REDACTED]; Microsoft's submission to the CMA [REDACTED].

¹⁸⁹⁴ Microsoft's submission to the CMA [REDACTED]. Microsoft's response to Ofcom's information request [REDACTED].

¹⁸⁹⁵ Microsoft's submission to the CMA [REDACTED]; See: [Microsoft licensing programs](#) and [License mobility through software assurance](#), accessed on 18 November 2024.

¹⁸⁹⁶ Microsoft's submission to the CMA [REDACTED]; Microsoft's response to Ofcom's information request [REDACTED].

¹⁸⁹⁷ Microsoft's submission to the CMA [REDACTED].

¹⁸⁹⁸ Microsoft's submission to the CMA [REDACTED].

route, the cloud provider will need to purchase a licence via its SPLA and the customer will need to purchase licence included services.

- 7.38 Microsoft includes itself as a Listed Provider, so the conditions for Listed Provider customers also apply to Microsoft. However, we have been told by Google and CISPE that Microsoft applies a different set of rules to itself and its own cloud offering.¹⁸⁹⁹ In particular, we were told that Microsoft has excluded Azure from the same restrictions as other Listed Providers, and markets this exclusion as the 'Azure Hybrid Benefit'.¹⁹⁰⁰
- 7.39 The Azure Hybrid Benefit ('AHB') allows customers with existing on-premises Windows Server or SQL Server core licences with Software Assurance subscriptions to migrate these licences onto Azure at a discount.¹⁹⁰¹
- 7.40 In reality, customers cannot simply 'migrate' their licences to Azure; instead, Microsoft will give customers a discount when purchasing virtual machines on Azure that include Windows Server as the operating system or SQL Server as the database management system.¹⁹⁰² For Windows Server, the discount is generally determined to charge the customer the same amount that Microsoft charges for a virtual machine (VM) with Linux on Azure.¹⁹⁰³ As explained on Microsoft's website, '[t]he license for Windows Server is covered by Azure Hybrid Benefit, so [the customer] only need to pay for the base computer rate of the VM. The base compute rate is equal to the Linux rate for VMs'.¹⁹⁰⁴
- 7.41 The submissions relating to the non-price factors are wide-ranging and include (but are not limited to):
- (a) The non-availability of certain Microsoft products (eg Microsoft 365 and Desktop 10/11) on Listed Providers clouds. These products are not available via the SPLA to use on other cloud providers' public cloud.
 - (b) Limiting extended security updates (ESUs). ESUs are only available for three years on non-Azure clouds, whereas they are available for 4 years on Azure. In addition, ESUs are free on Azure but need to be paid for on other clouds.
 - (c) The non-availability of certain features for Microsoft products that are run in other clouds (eg Windows 10/11 Multisession).

¹⁸⁹⁹ [Listed Providers](#), accessed on 18 November 2024; Google's submission to the CMA [redacted]; CISPE's response to Ofcom's information request [redacted].

¹⁹⁰⁰ [redacted] submission to the CMA [redacted].

¹⁹⁰¹ [Azure Hybrid Benefit - Hybrid Cost Calculator](#), accessed on 18 November 2024; See [Explore Azure Hybrid Benefit for Windows VMs](#) and [Azure Hybrid Benefit - Azure SQL Database & SQL Managed Instance](#), accessed on 18 November 2024.

¹⁹⁰² [redacted] response to Ofcom's information request [redacted].

¹⁹⁰³ Microsoft's submission to the CMA [redacted]; Responses to the CMA's information requests [redacted].

¹⁹⁰⁴ [Azure Hybrid Benefit for Windows Server](#), accessed on 25 June 2025.

Conceptual framework

- 7.42 We have investigated whether Microsoft's licensing practices, as described above, adversely impact competition in the supply of cloud services by inhibiting AWS and Google, its main rivals, from competing effectively for certain customers. The economic framework we have applied to this assessment is partial foreclosure, as set out in our Guidance.¹⁹⁰⁵
- 7.43 When assessing if a vertical relationship has an adverse effect on competition, the CMA will evaluate the overall impact on competition, taking into account rivalry-enhancing, as well as adverse effects. This will normally require an assessment of the impact of the vertical relationship on rivalry at different stages of the supply chain.¹⁹⁰⁶
- 7.44 A vertically integrated firm that has significant market power in the upstream market may have an incentive to refuse access to the input or to raise its price and consequently increase the costs of competing downstream firms. As a result, the firm may be able to maintain high prices and/or increase the prices charged to customers relative to the prices obtained in the absence of vertical integration.¹⁹⁰⁷
- 7.45 Foreclosure can be total (where rivals are forced to exit from the market or are prevented from entering) or partial (where rivals or potential entrants—are materially disadvantaged and consequently compete less effectively).¹⁹⁰⁸
- 7.46 The concern we have investigated is whether Microsoft's licensing practices, as described above, partially foreclose AWS and Google, its two principal downstream competitors, in the supply of cloud services, particularly in competing for customers that purchase cloud services which use certain upstream Microsoft software as an input.
- 7.47 In considering the effect of Microsoft's licensing practices, we have considered whether they are likely to have the effect of reducing competition in cloud services by materially disadvantaging AWS and Google who consequently compete less effectively.¹⁹⁰⁹
- 7.48 Any partial foreclosure of this nature could potentially harm competition in the supply of cloud services. Practices that make software licences more expensive when used with rival clouds may raise rivals' costs of supplying cloud services. One effect of this could be that Microsoft's rivals, acting as profit-maximising firms, are induced to pass on a proportion of these higher costs to their customers, thereby weakening the competition faced by Microsoft. The weakening of

¹⁹⁰⁵ CC3 (Revised), paragraphs 269 and 274.

¹⁹⁰⁶ CC3 (Revised), paragraph 273

¹⁹⁰⁷ CC3 (Revised), paragraph 270.

¹⁹⁰⁸ CC3 (Revised), paragraph 269.

¹⁹⁰⁹ CC3 (Revised), paragraph 269

competitors would allow Microsoft to grow its downstream share and raise price and/or reduce its offering in other ways.

- 7.49 Whilst in our assessment we focus on the above effect, there are other ways in which the raising of rivals' costs can have an impact on competition. Any negative impacts on margins and profitability may weaken ability and incentives to compete for some customers and (over the longer term) investment. Further, any practices that make a significant proportion of customer demand less contestable, may weaken rivals' ability to benefit from economies of scale, learning effects and/or network effects.¹⁹¹⁰
- 7.50 We have examined whether Microsoft has the ability through its licensing practices to partially foreclose its rivals, whether it has the incentive to do so and whether the conduct has an adverse effect on competition. In considering the effect of Microsoft's licensing practices, we have considered whether they are likely to have the effect of reducing competition in cloud services by materially disadvantaging the competitiveness of Microsoft's rivals and consequently harm competition.
- 7.51 We have set out Microsoft's licensing practices, and how they apply to Listed and non-Listed Providers above. We understand that, while some restrictions apply to non-Listed Providers, in general customers can use Microsoft software products on the same, or a similar, basis on non-Listed Provider clouds as on Azure.¹⁹¹¹ As such, our analysis focused on Microsoft's licensing practices for other Listed Providers, and AWS and Google in particular due to their presence in the UK.¹⁹¹²
- 7.52 We have set out evidence relevant to our assessment in the following sections:
- (a) Microsoft's market power in related software markets;¹⁹¹³
 - (b) the importance of Microsoft software inputs to AWS and Google;
 - (c) Microsoft's supply of software: price and non-price differences; and
 - (d) the impact on AWS' and Google's competitiveness from Microsoft's licensing practices.
- 7.53 These sections are interrelated and the evidence provided in them is in some cases relevant to more than one of the limbs of our assessment of ability, incentive and effect. Given this, we have provided summary conclusions for each of these

¹⁹¹⁰ CC3 (Revised), paragraph 268.

¹⁹¹¹ In July 2025, CISPE announced that it had reached an agreement with Microsoft related to its competition complaint filed against Microsoft with the European Commission in November 2022.

¹⁹¹² We do not consider conduct towards Alibaba given it has no UK presence as a cloud provider.

¹⁹¹³ See CC3 (Revised), paragraphs 9, 14, 101 and 178 to 204. The term 'market power' is used to denote the ability of a firm to influence aspects of competition. This investigation does not assess whether an undertaking holds a dominant position which is specific to an investigation under Chapter II of the Competition Act 1998.

sections; we then conclude on ability, incentive and effect in the round at the end of the chapter.

Inclusion of other cloud providers' conduct in this assessment

- 7.54 We received submissions from cloud markets stakeholders in relation to Microsoft's licensing practices which we have investigated.
- 7.55 We have also received some submissions relating to Oracle's software licensing practices. Further, Microsoft submitted that it is unfairly singled out when AWS and Google do not license their popular proprietary products (such as Aurora and Big Query).¹⁹¹⁴
- 7.56 We have not investigated the above submissions for the following reasons:
- (a) We have received far fewer submissions raising concerns about Oracle's licensing practices compared to the number of submissions about Microsoft's licensing practices. We have not received any submissions (apart from Microsoft's submission) about AWS' and Google's practices in relation to their cloud services.¹⁹¹⁵
 - (b) Oracle's market share of cloud is relatively small ([X] % for IaaS and [X] % for PaaS, as set out above in the Competitive landscape chapter) and it is not seen by large customers we spoke to as a suitable alternative to their main cloud providers. In addition, the Jigsaw report notes that none of the customers used Oracle as their sole cloud provider, and their main use was as a secondary cloud, for example, for supporting legacy systems.¹⁹¹⁶ As such, the potential for Oracle's licensing practices to have an adverse effect on competition for cloud services is far less than is the case for Microsoft.
 - (c) As already set out above, and in relation to AWS' and Google's cloud services listed by Microsoft, we note that it is not sufficient for a product or service to be proprietary to adversely impact competition. This also applies to our investigation of Microsoft's licensing practices where we have applied the framework to a specific set of Microsoft's software products rather than all of Microsoft's software products or cloud services (for example, Microsoft also offers cloud services, such as Azure Cosmos DB, for which it does not

¹⁹¹⁴ Microsoft has submitted that Amazon has S3 (Simple Storage Service), Aurora (Relational Database Management System), Dynamo DB (serverless document database) and more, whilst Google has BigQuery (data warehouse), Google Analytics and Looker (business intelligence and data analytics platform). [Microsoft's response to the PDR](#), paragraphs 42 and 43. Microsoft also submitted that the same way that Google does not share its internet search SaaS service to power Bing, Microsoft should not be forced to power Google's VDI solution [Microsoft's response to the PDR](#), paragraph 261.

¹⁹¹⁵ The examples identified by Microsoft are cloud services, while the licensing practices that we have investigated relate to software products that have historically been distributed to, and licensed for use by, customers and ISVs.

¹⁹¹⁶ [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), paragraph 3.4.8.

license an equivalent software product).¹⁹¹⁷ We have, however, considered whether AWS and/or Google having a stronger, differentiated offering as part of our assessment of their counterstrategies below.

- (d) In the absence of any indications that licensing practices in relation to other proprietary software products (of Microsoft or other providers) are having an adverse effect on competition in the supply of cloud services, we did not investigate these further.

Additional cloud provider submissions

Cloud provider submissions

- 7.57 AWS submitted that the raising of AWS' and Google's costs makes their offering less attractive and decreases competition for Microsoft workloads with Azure, which directly undermines competition and harms customers.¹⁹¹⁸
- 7.58 Google submitted that Microsoft's conduct raises rivals' costs and diminishes its rivals' ability to offer a compelling cloud offering.¹⁹¹⁹ Google also said that it considered that Microsoft's licensing practices harm competition and restrict customer choice, which ultimately leads to higher prices and prevents UK companies from realising the full benefits of public cloud.¹⁹²⁰
- 7.59 In response to our Licensing working paper, Microsoft submitted that it was 'spot on' to frame our assessment around whether the practice of making software licenses more expensive for rivals' customers may serve to raise rivals' costs.¹⁹²¹
- 7.60 In response to the PDR, Microsoft submitted that it wins some workloads that include a component – such as Microsoft software – for which it is offering a discount. It submitted that this is good evidence of price competition.¹⁹²² Microsoft submitted that the standing AHB discount is not foreclosure but the essence of competition, and that it benefits UK customers.¹⁹²³
- 7.61 Microsoft further submitted that for suppliers, discounts are at the deal level and the customer cares about the overall price for the deal they have negotiated rather than the pricing for components of a wider deal. It submitted that a discount can notionally be applied to one input and not the other to achieve a workload price, or a deal price could apply across aggregated workloads. Microsoft gave the example that it could remove the AHB discount on Windows Server VMs usage

¹⁹¹⁷ 'Azure Cosmos DB – Microsoft Azure', accessed on 9 June 2025

¹⁹¹⁸ AWS' submission to the CMA [38].

¹⁹¹⁹ [Google's response to the Issues Statement dated 17 October 2023](#), paragraph 45.

¹⁹²⁰ [Summary of hearing with Google](#), 17 March 2025, paragraph 21

¹⁹²¹ [Microsoft's response to the Licensing working paper](#), dated 10 July 2024, paragraph 2.15.

¹⁹²² [Microsoft's response to the PDR](#), paragraph 229.

¹⁹²³ [Microsoft's response to the PDR](#), paragraph 41.

and apply an equivalent discount to the 'storage' component. This would have no impact on the customers' overall price given that storage is always purchased with compute.¹⁹²⁴

- 7.62 In addition, Microsoft submitted that, in its view, Amazon (the market number one) cannot be a foreclosure 'victim' and that the implication of the PDR is that AWS would/should have yet more market share and profitability.¹⁹²⁵

Our assessment

- 7.63 With respect to Microsoft's submission that its licensing practices should be interpreted as competition through discounting rather than partial foreclosure, we note that a firm seeking to raise a specific set of rivals' costs could set a high initial price for all firms and then discount for all firms except for those rivals. The presence of a 'discount' therefore cannot, of itself, imply that a price disparity is driven by lower prices and competition (rather than higher prices to rivals that are consistent with partial foreclosure). Therefore, we consider that it is relevant to assess the likelihood of partial foreclosure even where price differences are framed as discounts.
- 7.64 A price disparity (in this case, between the cost of software licences when used with rival clouds and the cost when used with Microsoft's Azure service) is not sufficient evidence of foreclosure or an adverse effect on competition. As described above, a concern arises only if the pricing behaviour materially weakens the competitiveness of rivals and causes them to compete less effectively. We have therefore considered the specific features of Microsoft's licensing practices as well as its impact on AWS and Google to determine whether in this case the price disparity is likely to have an adverse effect on competition.
- 7.65 Further, we have focused our investigation on assessing Microsoft's current conduct rather than any hypothetical alternative conduct. In relation to Microsoft's argument that it could in theory apply a different type of discount (for example on storage), we note that any conduct that raises rivals' costs at the upstream level with the result of reducing competition downstream would potentially be of concern. This would also be the case if Microsoft were to create the final price disparity by applying the discount to another workload component.
- 7.66 Finally, we disagree that AWS cannot be partially foreclosed due to its size. AWS being a market leader in cloud infrastructure services is not inconsistent with AWS also being partially foreclosed in relation to a subset of customers or a subset of their spend. Our concern is about whether rivals compete materially less

¹⁹²⁴ [Microsoft's response to the PDR](#), paragraph 205.

¹⁹²⁵ [Microsoft's response to the PDR](#), paragraph 177

effectively than they otherwise would and whether this harms competition in the relevant markets.

Microsoft's market power in related software markets

Introduction

- 7.67 For a vertically integrated firm to foreclose its rivals, it must have significant market power in one or more markets along the supply chain.¹⁹²⁶
- 7.68 If Microsoft has limited or no market power in the software products relevant to the licensing concerns, cloud providers would be able to provide equivalent services with alternative software, mitigating any effect of the licensing practices in distorting customer choice towards Azure. We have therefore assessed the extent of any market power held by Microsoft in the relevant software products.
- 7.69 This section sets out:
- (a) the frameworks and types of evidence we used to assess the extent of any market power held by Microsoft in the supply of software;
 - (b) a description of each of the relevant Microsoft software products and their use cases on the public cloud; and
 - (c) our assessment of market definition and market power in relation to each of the five relevant software markets.

Framework

- 7.70 In this section we set out our framework for assessing the extent of any market power held by Microsoft in the supply of the relevant software products.

Approach to market definition and assessing market power

- 7.71 As regards market definition, the principles of our approach, including the hypothetical monopolist test, are set out in chapter 3. We have considered the definition of the relevant software markets here as a useful tool to inform our subsequent market power assessments.¹⁹²⁷
- 7.72 Market power refers to the ability of a firm to act independently of its competitors, consumers and customers, enabling it to prevent effective competition in the relevant market. It can be understood as a firm being able to profitably sustain

¹⁹²⁶ CC3 (Revised), paragraph 274.

¹⁹²⁷ One additional consideration that is relevant to our assessment of market definition in software markets is the concept known as the 'cellophane fallacy'. CC3 (Revised), paragraph 178.

prices that are above competitive levels, or output or quality levels that are below competitive levels.

- 7.73 The firm or firms holding market power will have the ability and incentive to influence market outcomes and other important aspects of competition¹⁹²⁸ which may include slowing innovation, reducing the variety and quality of the goods and services, raising entry barriers as well as other parameters of competition. This leads to harm in the process of competition.
- 7.74 Market power can arise for a range of reasons, including high market concentration, capacity constraints, lack of product substitutability, and an absence of supply-side constraints.¹⁹²⁹
- 7.75 In our assessment, we have focused on whether Microsoft has unilateral market power in relation to the relevant software products, and the degree of market power it has in each of those markets.
- 7.76 In doing so, we carefully considered the strength of any competitive constraints that would be likely to prevent Microsoft from profitably sustaining prices above competitive levels. This includes within-market and out-of-market constraints.
- 7.77 We have also considered whether there is any cumulative effect of Microsoft's position due to technical or commercial links across all of the relevant software markets primarily as part of individual market power assessments, and also cumulatively.

Evidence used in each assessment

- 7.78 In our assessment of market definition and market power for each of the relevant software products, we have relied on evidence on product characteristics, market shares, and evidence from customers and providers. Below we also consider some cross-cutting considerations relevant to interpreting these sources of evidence.

Market shares evidence

- 7.79 In general, a highly concentrated market might be an indicator that one or more firms hold unilateral market power. If a firm has a high market share it might have less incentive to compete vigorously with its rivals (particularly if there are barriers to entry).¹⁹³⁰

¹⁹²⁸ CC3 (Revised), paragraph 178.

¹⁹²⁹ CC3 (Revised), paragraph 185.

¹⁹³⁰ CC3 (Revised), paragraph 187.

- 7.80 Observed changes in market shares over time are also important. When market shares have been stable over time, especially in the face of historical changes in prices or costs, high concentration may indicate that competition within the market is weak. However, a highly concentrated market may be competitive if market shares fluctuate over short periods in response to changes in competitive offers; such volatility may indicate the existence of effective competitive constraints, such as successful entry and innovative developments.¹⁹³¹

Product characteristics evidence

- 7.81 Substitutability between two products can be assessed by reference to the characteristics and purpose of the relevant products. Where two products are designed to meet the same customer requirement, they are more likely to be substitutable for each other. Where two products satisfy quite different needs, they may be weaker substitutes or not substitutable at all. Nevertheless, we recognise that products with different characteristics can still be substitutes, while similar ones may not meet customer needs equally.¹⁹³²

Customer and provider evidence

- 7.82 Most of our evidence gathering has been focused on demand side substitution factors. We have also gathered some evidence on supply side substitution such as views from Microsoft and other providers on barriers to entry into the relevant markets and about competitive constraints they face in the relevant markets.
- 7.83 For each of the relevant software products, we asked Microsoft customers that use the products on the public cloud about the likelihood that they would switch away from Microsoft products if Microsoft was to raise its prices by 5% from the price customers are being charged today.¹⁹³³
- 7.84 As set out above, the customer evidence we have collected is qualitative and so we have given a narrative summary of the key points from the evidence.

Evidence from public and non-public cloud customers

- 7.85 We asked customers about their use of the products on the public cloud. Most customers answered with this in mind, but some also or only use the product on non-public cloud. In summarising customer evidence, we make the distinction

¹⁹³¹ CC3 (Revised), paragraph 192.

¹⁹³² An example of this may be seen in the advent of new, efficient technologies. While both technologies aim to meet the same requirement in principle, if the new technology does so with considerably greater efficiency, the older technology may no longer represent a good substitute once price and quality factors are taken into account. Another example is where high switching costs or a weak customer response may mean that substitutability is limited even in the presence of similar product characteristics.

¹⁹³³ In evidence gathering, we prioritised asking customers about their use of products on the public cloud. We also have some limited evidence from customers who used the products on premises and have highlighted where this is the case.

between customers who use these products on the public cloud and customers who use these products on non-public cloud.¹⁹³⁴ Depending on the particular product, evidence from users of the product on non-public cloud has been included for the following reasons:

- (a) Customer evidence on use of the product on non-public cloud and public cloud suggests customers may find it hard to switch away from the Microsoft product;
- (b) Some customers' reasons for choosing Microsoft products relate to their historical use of the product (for example, staff have developed skills in it or other applications rely on it); and
- (c) Customer evidence on alternatives did not highlight that there was a different range of alternatives on the public cloud versus non-public cloud.

Product background

- 7.86 In this section, we describe each of the relevant Microsoft software products and outline their use cases on the public cloud.
- 7.87 Windows Server and SQL Server are run on a server on premises. For these products when used in the public cloud they are installed directly on a virtual machine or can be used as part of a PaaS solution.
- 7.88 Windows 10/11, Office and Visual Studio run on the desktop or are 'client side' products.¹⁹³⁵ Their use case in the public cloud is primarily when they are installed as part of a virtual desktop infrastructure (VDI) service.

Microsoft Windows Server

- 7.89 Server operating system (OS) software is designed to run a server's hardware and provide a platform for the use of application software. This is similar to how a desktop OS is used to run applications on a personal computer. For example, in a typical corporate use case, Microsoft Windows Server (Windows Server) can be installed on a central computer to coordinate and manage employees' access to shared storage, printers, or other devices.¹⁹³⁶
- 7.90 Microsoft Windows Server is one of the most popular types of server OS. Other types of server OS provide the same basic functionality and include variants of

¹⁹³⁴ VDI allows customers to use Windows Desktop, Visual Studio and Office/Microsoft 365 in the public cloud. Where relevant we separate customer evidence where customers use VDI. Further detail about how we consider the customer evidence for Windows Desktop, Visual Studio and Office/Microsoft 365 can be found in appendix R.

¹⁹³⁵ The client-side relevant Microsoft software products are those usually accessed on through the desktop, for which the use case on the cloud is for use through a virtual desktop.

¹⁹³⁶ Panek C (2019), [Windows Server Administration Fundamentals](#), Wiley.

Linux and UNIX OSs. Windows Server is proprietary, and Linux distributions (ie operating systems built around the Linux kernel), are open-source.¹⁹³⁷ Customers using on-premises versions of server OSs may install it on a physical computer to which other devices on the same network connect. Customers using a cloud version of a server OS may install it on a virtual machine hosted in the cloud.¹⁹³⁸

- 7.91 There are various ways in which customers use Windows Server on a virtual machine ‘in the cloud’.¹⁹³⁹
- 7.92 Active Directory is software that is included in Windows Server. It can be used to set up a so-called ‘directory service’. In a network of Windows PCs or servers, a directory service can be thought of as a list of objects—for example names, users, company locations, printers, and lists—that describe who has access to what.
- 7.93 Windows Server integrates well with Windows 10/11. For example, one cloud provider submitted that Windows 10/11 has superior interoperability (including in speed) with Windows Server.¹⁹⁴⁰

Microsoft SQL Server

- 7.94 Microsoft SQL Server (SQL Server) is a Relational Database Management System (RDBMS). A RDBMS is a type of Database Management System (DBMS) which manages and stores data in separate tables and defines relationships between those tables.
- 7.95 Microsoft SQL Server and alternative RDBMS software products can be used on-premises, in the cloud on a virtual machine, or as a managed service. RDBMS products can be open-source, for example PostgreSQL or MySQL, or they can be proprietary like SQL Server or IBM DB2. Open-source software variants can be accessed as a self-hosted variant (through a deployment type of the organisations’ choice), or in a supported version provided by a supplier (typically in the cloud as a

¹⁹³⁷ Dalheimer, MK and Welsh, M (2005), [Running Linux, 5th Edition](#), O’Reilly.

Individual distributions of Linux or UNIX may be free or paid for. Paid for Linux: Ubuntu (Ubuntu Pro for more than 5 machines); SUSE Linux Enterprise Server; Red Hat Enterprise Linux; Amazon Linux (included within Amazon EC2 and AWS charges). Free Linux: Debian; CentOS Linux. Paid for UNIX: IBM AIX; Oracle Solaris.

Source: company websites ([‘Virtual Machines \(VMs\) for Linux and Windows’](#), accessed on 18 November 2024; [Ubuntu Pro: security, compliance and support](#), accessed on 18 November 2024; [SUSE Linux Enterprise Server](#), accessed on 18 November 2024; [Red Hat Enterprise Linux Server](#), accessed on 18 November 2024; [Amazon Linux 2 FAQs](#), accessed on 18 November 2024; [Amazon Linux 2 FAQs](#), accessed on 18 November 2024; [IBM Power AIX](#), accessed on 12 December 2024; Free with Oracle cloud product ([The Oracle Solaris 11.4 operating system](#), accessed on 18 November 2024) and paid for on non-Oracle hardware ([Oracle Solaris](#), accessed on 18 November 2024); [About Debian](#), accessed on 18 November 2024; [About CentOS](#), accessed on 18 November 2024) plus CMA analysis.

¹⁹³⁸ [‘Virtual Machines \(VMs\) for Linux and Windows’](#), accessed on 18 November 2024.

¹⁹³⁹ Ways Windows Server is used in the cloud include: (1) as a ‘work group server’, meaning an OS that runs on a central network computer that provides services to office workers in their day-to-day work, such as file and printer sharing, security, and user identity management; (2) to set up Virtual Desktop Infrastructure (VDI) services; (3) to host customers’ custom software applications, such as custom web applications; and (4) installed on an organisation’s server to host off-the-shelf enterprise applications, such as Customer Relationship Management (CRM) software. References: [Microsoft Corp. v Commission of the European Communities](#); [Recommended settings for VDI desktops](#), accessed on 18 November 2024; [§] submission to the CMA [§]; Responses to the CMA’s information requests [§].

¹⁹⁴⁰ [§] submission to the CMA [§].

PaaS service). Proprietary RDBMS can also be provided in a managed service (such as Microsoft's Azure SQL¹⁹⁴¹).

- 7.96 While SQL Server can also be run on Linux server OSs, Microsoft submitted that it is designed to integrate with Windows Server¹⁹⁴² and a cloud provider submitted that customers are most likely to run this workload on Windows Server.¹⁹⁴³ We also understand that SQL Server works well with other Microsoft products and technologies,¹⁹⁴⁴ and Microsoft has also submitted that it is designed to integrate well with other Microsoft products that customers use together with SQL Server, such as Windows Server, Azure and Microsoft Dynamics.¹⁹⁴⁵ In addition, Microsoft's technical documentation suggests that either SQL Server or Azure SQL Managed Instance (the 'halfway to PaaS' service) is required for some SharePoint software (referred to as SharePoint Server).¹⁹⁴⁶

Microsoft Windows 10/11

- 7.97 Desktop OS software is designed to run a personal computer's hardware and provides a platform for the use of application software.¹⁹⁴⁷
- 7.98 Microsoft Windows 10/11 (Windows 10/11) is the most popular desktop OS. Other types of personal computer (PC) OS provide this same basic functionality and include macOS, ChromeOS and deployments of Linux OS eg Ubuntu. We understand that there are also a variety of use cases in which customers use Windows 10/11 in combination with public cloud infrastructure services. For example, we have seen evidence that:
- (a) a customer who uses a VDI cloud service, such as AWS' WorkSpaces, can allow its staff to access a Windows 10/11-based virtual desktop;¹⁹⁴⁸ and
 - (b) a version of Windows 10/11 can be installed on a cloud-hosted virtual machine in a 'multi-session' configuration, meaning that multiple users can concurrently use a single instance of the OS.¹⁹⁴⁹

¹⁹⁴¹ Microsoft provides two managed services which share a common code base with SQL Server, Azure SQL MI and Azure SQL DB, for which most of the standard SQL language, query processing, and database management features are identical to SQL Server. Details about the functionalities of these services are explained here ([Compare SQL database engine features - Azure SQL Database & Azure SQL Managed Instance](#), accessed on 16 October 2024).

¹⁹⁴² Microsoft's response to the CMA's information request [§].

¹⁹⁴³ [§] response to the CMA's information request [§].

¹⁹⁴⁴ For example, see [PostgreSQL vs. SQL Server: What's the difference?](#) accessed on 18 November 2023. This is discussed further by some customers in SQL Server market power assessment, customer responses.

¹⁹⁴⁵ Microsoft's response to the CMA's information request [§].

¹⁹⁴⁶ [Software requirements for Database Servers for SharePoint Server Subscription Edition](#), accessed on 3 October 2024.

¹⁹⁴⁷ [§] response to the CMA's information request [§].

¹⁹⁴⁸ 'What is Amazon WorkSpaces?', AWS online documentation, accessed on 18 November 2024.

¹⁹⁴⁹ Submissions to the CMA [§]; [Windows 10 or Windows 11 Enterprise multi-session remote desktops](#), accessed on 15 April 2024.

Microsoft's productivity suites

- 7.99 Microsoft has a number of packages which provide productivity functionality. The main categories of packages we have considered as part of this investigation are:
- (a) Microsoft Office (Office): includes desktop versions of Word, Excel, PowerPoint, and Outlook (the Microsoft Apps) at a minimum.¹⁹⁵⁰ This package is available on a one-time purchase basis for one PC or Mac.
 - (b) Office 365: includes all of the products included in Office as desktop versions, as well as additional apps which provide productivity and other functionality.¹⁹⁵¹ These additional apps vary dependent on specification.
 - (c) Microsoft 365: includes all of the products included in Office as desktop versions, as well as additional apps which provide productivity and other functionality.¹⁹⁵² The additional apps vary dependent on specification. The popular Microsoft 365 for Enterprise packages (E3/5) include Windows 10/11 (and the ability to deploy this in the cloud on Azure), security capabilities and advanced identity and access management (including Entra ID).
- 7.100 Office 365 and Microsoft 365 are purchased through a subscription service where updates are provided as new versions are released. Both packages include desktop installed versions of the Microsoft Apps, as well as access to the software through a SaaS solution in the browser. In the following analysis, we do not distinguish between Office 365 and Microsoft 365, in particular with respect to customer responses, as reported use of Office 365 was minimal, and it is listed by Microsoft as a Microsoft 365 package.¹⁹⁵³
- 7.101 On the cloud a customer can access Office installed as part of a virtual desktop solution. A customer can also access Microsoft 365 functionality through a virtual desktop solution provided by Microsoft,¹⁹⁵⁴ and there is some limited Microsoft 365 functionality available on Amazon Workspaces.¹⁹⁵⁵

Microsoft Visual Studio

- 7.102 Microsoft Visual Studio (Visual Studio) is a type of Integrated Development Environment (IDE). IDEs are a type of software containing a range of tools that

¹⁹⁵⁰ For example, see Office Home & Business 2021 ([Buy Office Home & Business 2021 \(PC or Mac\) - Download & Pricing - Microsoft Store](#)) and Office 2021 Professional ([Buy Microsoft Office Professional 2021 - Download Key & Pricing](#)), accessed on 18 September 2024).

¹⁹⁵¹ See [Compare Office 365 Enterprise Pricing and Plans](#), accessed on 18 November 2024.

¹⁹⁵² See [Compare Microsoft 365 Enterprise Plans](#), accessed on 18 November 2024.

¹⁹⁵³ See [Compare Office 365 Enterprise Pricing and Plans](#), accessed on 11 October 2024.

¹⁹⁵⁴ Note, this functionality is only available for those using Microsoft 365 in Azure.

¹⁹⁵⁵ Recently, Microsoft has made some functionalities of Microsoft 365 (the Microsoft 365 Apps for enterprise [namely Word, Excel, PowerPoint, Outlook and OneDrive]) available through Amazon Workspaces. [Microsoft 365 Apps for enterprise now available on Amazon WorkSpaces services](#), accessed on 23 May 2024. This is discussed in more detail below.

software engineers use to build applications, web pages or services. IDEs typically include a code editor (a text editor designed for editing source code). They may also have additional features such as intelligent code completion, a compiler/interpreter, build automation tools, debugger, testing or project management tools and AI integration.

7.103 Visual Studio is no longer supported for Mac OS,¹⁹⁵⁶ and is not supported for Linux.¹⁹⁵⁷ Therefore Mac OS or Linux customers would need to use Visual Studio Code,¹⁹⁵⁸ or Visual Studio on Windows 10/11 or choose an alternative IDE. Visual Studio is also commonly used to develop business applications which run on Windows Server and Windows 10/11.

7.104 We understand that customers either use Visual Studio:

- (a) on-premises;¹⁹⁵⁹ or
- (b) as part of a VDI solution, for example by installing Visual Studio on a virtual machine, using a cloud infrastructure service such as AWS EC2.¹⁹⁶⁰

How the Microsoft Software products are purchased

7.105 As outlined above, customers can purchase Microsoft software products through various routes. When customers purchase licences directly from Microsoft, we understand that one of the available routes for purchase for some organisations is through an Enterprise Agreement (EA).

7.106 Many of the customers whose responses we consider below reported having an EA with Microsoft.¹⁹⁶¹ By entering into an EA with Microsoft, customers can purchase cloud services and software together, and simplify licensing by using only per user licensing, rather than per device licensing.¹⁹⁶² Microsoft submitted that EAs are available to organisations with more than 500 users/devices (or 250 users/devices for public sector entities). EAs are typically entered into for a period of three years [REDACTED].¹⁹⁶³

¹⁹⁵⁶ In August 2023, Microsoft announced that it would deprecate Visual Studio for Mac, from 31 August 2024 '[Visual Studio for Mac Retirement Announcement](#)', accessed on 18 November 2024.

¹⁹⁵⁷ [Visual Studio 2022 Downloads](#), accessed on 18 November 2024.

¹⁹⁵⁸ [Visual Studio 2022 Downloads](#), accessed on 18 November 2024.

¹⁹⁵⁹ [REDACTED] response to the CMA's information request [REDACTED].

¹⁹⁶⁰ [REDACTED] response to the CMA's information request [REDACTED].

¹⁹⁶¹ Responses to the CMA's information requests [REDACTED].

¹⁹⁶² [Enterprise Agreement | Microsoft Volume Licensing](#), accessed on 18 September 2024.

¹⁹⁶³ Microsoft's response to the CMA's information request [REDACTED].

Stakeholder views

- 7.107 Some stakeholders agreed with our assessment that Microsoft has market power in some software products:
- (a) Google said ‘customer feedback overwhelmingly supports the CMA’s emerging view that Microsoft has a “significant degree of market power in relation to its supply of [the five relevant products]”’.¹⁹⁶⁴
 - (b) AWS said ‘its customers [have continued] concerns over use of Microsoft’s numerous critical software products using the IT infrastructure of their choice’.¹⁹⁶⁵
 - (c) The SMF said ‘as a legacy provider with key legacy software products, Microsoft is believed to substantially undermine competition and limit consumer choice in the cloud services sector by the foreclosure of alternatives’.¹⁹⁶⁶
- 7.108 An academic disagreed with our approach to market definition suggesting a single global market for all software products for enterprises.¹⁹⁶⁷ We note this submission was not supported by accompanying evidence. In defining markets we have started with the narrowest plausible market and then consider whether the market should be widened with reference to demand and supply side substitution.¹⁹⁶⁸
- 7.109 Both Microsoft and rival cloud providers commented on the relevance of on-premises workloads.
- (a) Microsoft submitted that much or most of the CMA’s analysis on alleged market power in various on-premises software markets is irrelevant to cloud competition.¹⁹⁶⁹
 - (b) Google said that on-premises workloads matter. It said: ‘the proportion of existing cloud customers that are locked into Microsoft’s dominant software ecosystem is “significant” and the proportion of addressable cloud customers

¹⁹⁶⁴ [Google’s response to the Licensing working paper](#), 24 July 2024, paragraph 5.

¹⁹⁶⁵ [AWS’ response to the Licensing working paper](#), 1 July 2024, page 23-24.

¹⁹⁶⁶ [Social Market Foundation’s response](#) to the Licensing working paper, page 22.

¹⁹⁶⁷ [Dr George R Barker’s response to the Licensing and Potential remedies working papers and Updated Issues statement](#) dated 06 June 2024, page 36. He suggested instead of five individual market definitions, the market definition should be the global market for software products for enterprises. Further, they submitted that a market defined on this basis is competitive because barriers to entry and expansion are low, existing suppliers face competition from open source providers and piracy, and consumers face low switching costs. We note that Dr George Barker is a member of the Oxford Cross Disciplinary Machine Learning Research Cluster (OXML), which is supported by Microsoft (see [Dr George R Barker’s response to the Competitive landscape, Egress fees and CSAs working papers](#) dated 23 May 2024, page 1).

¹⁹⁶⁸ See [CC3 \(Revised\)](#), paragraph 130; Further, CMA guidance states that market definition is a useful tool, but not an end in itself, and that identifying the relevant market involves an element of judgement. The boundaries of the market do not determine the outcome of our competitive assessment in any mechanistic way. The competitive assessment will take into account any relevant constraints from outside the market, segmentation within it, or other ways in which some constraints are more important than others. See [CC3 \(Revised\)](#), paragraph 133.

¹⁹⁶⁹ [Microsoft response to the CMA’s Licensing working paper](#), 10 July 2024, paragraph 5.1.

(ie existing cloud customers plus customers whose workloads are currently on-premises) that are locked into Microsoft's dominant software ecosystem is even more significant.¹⁹⁷⁰ It submitted that in its view, shares of on-premises workloads provide a strong indication of the competitive impact that Microsoft's licensing practices will have as the migration from on-premises to cloud continues.¹⁹⁷¹

- (c) A cloud provider outlined that most customers 'lift and shift' in the first instance as part of their transition to cloud. In support of this it submitted that based on [REDACTED] internal data, [90-100]% of Windows Server VMs run on IaaS (versus running on non-IaaS), as opposed to [50-60]% of non-Windows Server VMs.¹⁹⁷²
- (d) This could provide additional evidence that customers may be sticky when they migrate to the cloud in terms of which server OS they choose to deploy their workloads on. However, there may be other reasons for this trend, for example if there are fewer managed services which are run on/with Windows Server.
- (e) This point was supported by AWS, which said 'customers expect to be able to run the same Microsoft software they can run on-premises, such as Windows Server and SQL Server, on EC2 [AWS' compute service]'.¹⁹⁷³

Market definition and market power assessments

7.110 In this section we:

- (a) consider geographic market definition, which is common to all of the software markets considered;
- (b) set out the evidence on product market definition and market power for each of the following Microsoft software products in turn: Windows Server, SQL Server, Windows 10/11, Visual Studio and Microsoft 365; and
- (c) set out additional evidence which relates to the cumulative effect of market power across markets where this is not specific to one product, and which may strengthen our views on the finding of market power in individual product markets.

¹⁹⁷⁰ Google's submission to the CMA [REDACTED].

¹⁹⁷¹ Google's submission to the CMA, [REDACTED].

¹⁹⁷² [REDACTED] submission to the CMA [REDACTED].

¹⁹⁷³ AWS' submission to the CMA [REDACTED].

Geographic market

- 7.111 Geographic markets may be based on the location of either suppliers or customers. In the case of the former, the geographic market is an area covering a set of firms or outlets which compete closely because enough customers consider them to be substitutes (as in the case of retail markets and some industrial markets). In the latter case, a geographic market is an aggregation of customers paying individually negotiated prices but enjoying sufficiently similar purchasing options (ie in effect many customers in industrial markets).¹⁹⁷⁴ The geographic market may be local, regional, national or wider.¹⁹⁷⁵
- 7.112 The key to defining both supplier-based and customer-based geographic markets, as to defining the product market, is the degree of substitutability, ie the extent to which suppliers can switch their areas of supply and the extent to which customers in one area may be served in another area.¹⁹⁷⁶
- 7.113 Market characteristics in the relevant markets point towards a global market definition because:
- (a) the same product is sold internationally for each of the Microsoft products (with language differences);
 - (b) consumers can use the same Microsoft product licence across multiple countries; and
 - (c) barriers to the flow of goods are minimal as the Microsoft software can be downloaded anywhere.
- 7.114 On this basis, our view is that there is a global geographic market for each of the relevant products.¹⁹⁷⁷

Microsoft Windows Server

- 7.115 For the purposes of this investigation, server OSs is the narrowest product market within which Windows Server sits. In the following section, we consider whether the market should be widened to include desktop OSs. We then consider the extent of any market power held by Microsoft in relation to the relevant market.

¹⁹⁷⁴ CC3 (Revised), paragraph 145.

¹⁹⁷⁵ CC3 (Revised), paragraph 146.

¹⁹⁷⁶ CC3 (Revised), paragraph 147.

¹⁹⁷⁷ In response to the Licensing working paper, only one party commented on the geographic market. They agreed with this approach. Dr George R Barker, Comment on The UK Competitive Market Authority's (CMAs) Cloud Services Market Investigation Updated Issues Paper and Working Papers 4-6 on The Supply of Public Cloud Infrastructure Services In the UK Covering The CMA's 1. Updated issues statement on Public cloud infrastructure services market investigation: 2. Licensing Practises Working Paper; 3. Technical Barriers; and 4. Potential Remedies, page 35.

Product market definition

Provider submissions

- 7.116 We asked Microsoft and software competitors whether there were certain use cases where a desktop OS could be used as a substitute for a server OS.
- (a) Microsoft said this was possible in theory but believed it would not be a common scenario. It submitted that it was unclear why any customer would install a desktop OS to control a much more powerful server. Microsoft also said both server OSs and desktop OSs can be used to provide desktop as a service offerings.¹⁹⁷⁸
 - (b) AWS and IBM are other providers of server OSs. AWS said desktop OSs are generally not substitutable for server OSs because server OSs are built for multiple users logging in at the same time while desktop OSs are not.¹⁹⁷⁹
 - (c) IBM said the degree of substitutability depends on the application and whether the application will sufficiently and effectively run on the desktop OSs, and considered the opposite is more common (server OSs can be used as a desktop OS).¹⁹⁸⁰
- 7.117 Views from providers suggest that the relevant market should not be expanded to include desktop OSs.

Customer evidence

- 7.118 We asked customers that use Windows Server on the public cloud to identify any alternatives to Windows Server that they could use for the same purpose. Most customers that responded to this question identified other server OSs.¹⁹⁸¹ No customers we heard from suggested that a desktop OS would serve as an alternative. This suggests that customers do not view them as substitutes on the demand side.
- 7.119 We asked the same customers (that use Windows Server) which other server OSs they used, if any. All customers that responded to this question submitted that they currently used both Windows Server and various Linux distributions (ie versions), suggesting that customers may value differentiated functionality of Linux.¹⁹⁸² We considered whether each server OS could be used for different purposes, and that evidence is explained below.

¹⁹⁷⁸ Microsoft's response to the CMA's information request [X].

¹⁹⁷⁹ AWS' response to the CMA's information request [X].

¹⁹⁸⁰ IBM's response to the CMA's information request [X].

¹⁹⁸¹ Responses to the CMA's information requests [X].

¹⁹⁸² Responses to the CMA's information requests [X].

The product market for Windows Server

- 7.120 Our view is that the relevant product market for Windows Server is the market for server OSs and that it should not be widened to include desktop OSs, meaning Linux/UNIX server OS distributions would be included within the market.

Market power

Market shares

- 7.121 Large and persistent market shares can indicate barriers to entry/expansion and market power. However, shares should be interpreted carefully because they can understate market power if substitutability is limited by products being very differentiated in terms of their use cases, and/or because of switching costs in moving away from an incumbent supplier.
- 7.122 The table below summarises the variety of different measures we have to understand Microsoft's share of server OSs. Full explanation with relevant caveats is in appendix D.

Table 7.1 Microsoft's market share in server OS, variety of measures

<i>Measure</i>	<i>Deployment type</i>	<i>Source</i>	<i>Description</i>	<i>Microsoft's market share</i>
Revenue	All deployments	CMA analysis of IDC data	Shares of supply by revenue are typically the most direct distribution of customer demand as they take into account the differences in prices and quality of firm's offerings.	2023: [70-80]%
Revenue	Public cloud			2023: [60-70]%
Installed base	Cloud and 'non-cloud'	CMA analysis of data submitted by Microsoft	A metric of OS units on server hosts.	2022: Combined: [40-50]% Non-cloud: [70-80]% Cloud: [20-30]%
Shipments	Virtualised and physical	CMA analysis of data submitted by Microsoft	A flow measure of shares. These shares illustrate the direction of travel of shares of supply. Virtualized shipments include guest OS instances associated with an OS licence as well as OS instances deployed in the public cloud.	2022: Combined: [20-30]% Physical: [30-40]% Virtualised: [10-20]%
Azure's share of Windows Server compute	Public cloud	Keystone submission from Microsoft	A measure of use of Windows Server on the cloud by count of Azure VMs running Windows Server.	2023: [20-30]%
Windows Server share of Azure VM usage	Public cloud	CMA analysis of data submitted by Microsoft	A measure of use of Windows Server based on total annual vcore hours of usage of Azure VMs running Windows Server.	2023: [40-50]%

- 7.123 Microsoft has a substantial share of server OS market on several different measures. Based on some measures (particularly on forward-looking 'shipment' shares on the cloud), there is at least the possibility of shares declining in future (even though they appear stable on Azure in the last two years¹⁹⁸³ and the rate of

¹⁹⁸³ See Azure share of Windows Server compute analysis in appendix R.

any future decline is not known). We have not seen evidence on the drivers of declining shares – for example, whether this is caused by substitution to Linux directly (ie the constraint becoming stronger), or simply because of faster growth in the kinds of workloads that are suited to Linux (which would not necessarily imply a stronger/closer substitution).

- 7.124 The data suggests Windows Server represents a significant share of on-premises deployments. This could contribute to market power for Windows Server deployments on the cloud if customers find it hard to switch away when they migrate to the cloud.

Product characteristics

- 7.125 We considered the evidence we received from customers, software providers and academic research to understand the product characteristics of Windows Server:

- (a) Some customers said that they chose Windows Server over alternatives for reasons that relate to functionality of Windows Server.¹⁹⁸⁴ These customers generally said other software or applications require it or integrate with it.¹⁹⁸⁵
- (b) There is also evidence of product characteristics that may make other OSs weak substitutes, such as associations with different application building languages and frameworks.¹⁹⁸⁶ Some customers told us that applications designed for one server OS would be difficult or not possible to run on another.¹⁹⁸⁷
- (c) [REDACTED].¹⁹⁸⁸ Similarly, AWS noted that Linux cannot be used as an alternative to Windows Server in all cases, and that Linux is unable to run any of Microsoft's productivity software.¹⁹⁸⁹

Customer evidence on alternatives

- 7.126 In terms of alternatives to Windows Server, customers we heard from generally listed server OSs from the Linux family, and some said there were none.¹⁹⁹⁰

¹⁹⁸⁴ Other reasons are discussed below section 'Customers' views' and appendix R.

¹⁹⁸⁵ Responses to the CMA's information requests [REDACTED].

¹⁹⁸⁶ For example, historically, .NET apps only ran on Windows and not on Linux. The release of .NET Core made it possible for Linux to become a fully supported platform for the Microsoft development stack, but this has some restrictions and it may not be feasible to port existing apps from .NET Framework to .NET Core. ([Deploying .NET apps on Google Cloud](#), accessed on 18 November 2024).

¹⁹⁸⁷ Responses to the CMA's information requests [REDACTED].

¹⁹⁸⁸ [REDACTED] submission to the CMA [REDACTED].

¹⁹⁸⁹ AWS' response to the CMA's information request [REDACTED].

¹⁹⁹⁰ This included a few customers who only uses the product on non-public cloud. Responses to the CMA's information requests [REDACTED]. Further customer evidence is presented in appendix R.

7.127 We considered the evidence we received from customers who responded to our question about whether they would move away from Windows Server in response to a 5% price rise. This evidence showed that:

- (a) Many of these customers said they would be unlikely to move away. Reasons included that it is required for some software, the cost to re-build custom applications, the requirement to re-train staff, loss of functionality, and integrations with other Microsoft products.¹⁹⁹¹
- (b) Some customers indicated that they would be likely to move away. For some of those who were likely to move away in response to the price rise, it was in the context of a general preference to reduce reliance on Windows Server and one customer explicitly mentioned this could take years. Reasons included wanting to move individual workloads to Linux to improve portability, or that Linux was their preferred type of server OS.¹⁹⁹²
- (c) Reasons customers gave for why they chose Windows Server rather than alternatives included: staff skills, technical requirements, required by third party software providers or other software, integrations with other Microsoft software, support provided by Microsoft, required to run a legacy code base and 'market dominance'.¹⁹⁹³

Provider submissions on alternatives

7.128 In relation to demand-side substitution, key points from submissions received from providers are summarised below:¹⁹⁹⁴

- (a) Providers consider other server OS as alternatives to Windows Server.
- (b) One provider said there are some use cases for which Linux cannot be a substitute for Windows Server.

7.129 In relation to supply-side substitution, providers had mixed views on whether there are barriers to entry and expansion.¹⁹⁹⁵

7.130 Microsoft said that, in relation to cloud OSs, there are no material barriers to entry. It said historically the primary barrier to supplying an OS in an on-premises

¹⁹⁹¹ This included a few customers who only uses the product on non-public cloud. Responses to the CMA's information requests [X].

¹⁹⁹² This included some customers who only uses the product on non-public cloud. Responses to the CMA's information requests [X].

¹⁹⁹³ Further customer evidence is presented in appendix R.

¹⁹⁹⁴ Further provider evidence is presented in appendix R.

¹⁹⁹⁵ Further provider evidence is presented in appendix R.

environment was ensuring that there are sufficient applications to run on the OS to meet the customer's needs.¹⁹⁹⁶

- 7.131 AWS said Windows Server is viewed by customers as a critical software and seen by many as a 'must have' service, and that the ability to access this software often affects a customer's decision whether or not to use a particular IT provider.¹⁹⁹⁷ IBM said barriers exist to an extent in terms of regulatory requirements, cost, economies of scale and distribution.¹⁹⁹⁸

Conclusions on Windows Server

- 7.132 Overall, the evidence we received suggests Windows Server is differentiated and some customers choose it because of its specific product functionalities.
- 7.133 Microsoft has a significant degree of market power in relation to Windows Server because Microsoft has a high share of the server OS market and many customers we heard from would be unlikely to switch away from Windows Server in response to a small but significant price increase. We consider that desktop OSs do not pose a constraint on server OSs because they fulfil different use cases and customers did not view them as substitutes.
- 7.134 Some customers also highlighted the links between Windows Server and other Microsoft software, which would increase the difficulty of switching away. This reinforces Microsoft's market power in relation to Windows Server.

Microsoft SQL Server

- 7.135 For the purposes of this investigation, we consider that RDBMS is narrowest plausible candidate market which SQL Server sits within.
- 7.136 In the following section, we consider whether we should widen the market to include other types of database management systems such as non-relational database management systems (NRDBMS). We then consider the extent of any market power held by Microsoft in relation to the relevant market.

¹⁹⁹⁶ Microsoft's response to the CMA's information request [redacted].

¹⁹⁹⁷ AWS' response to the CMA's information request [redacted].

¹⁹⁹⁸ [redacted] response to the CMA's information request [redacted].

Product market definition

Provider submissions

- 7.137 We asked software providers to explain whether they would consider RDBMS and NRDBMS as substitutes, and whether they would consider any other types of DBMSs as substitutes for RDBMSs.
- (a) Microsoft submitted that a RDBMS such as SQL Server can be substituted with a NRDBMS depending on the specific requirements and characteristics of the customer seeking to switch DBMS.¹⁹⁹⁹
 - (b) Oracle submitted that it considers RDBMS and NRDBMS as substitutes, although they may have historically had advantages over one another. It submitted that other forms of DBMS have also started to serve as substitutes for RDBMS in recent years, for example non-schematic (also called NoSQL) DBMS, multi-model DBMS and in-memory DBMS ('IMDBMS').²⁰⁰⁰
 - (c) IBM submitted that it does not consider that RDBMS and NRDBMS are substitutes, as each type of system has unique features and areas in which they would provide a better service.²⁰⁰¹ It explained that the applications for which an RDBMS or NRDBMS would be best suited would differ on a case-by case basis, substituting one with the other is usually not possible without major effort and requires significant changes to the application(s). IBM explained that there is a large ecosystem of existing systems that depend on RDBMS [REDACTED]²⁰⁰²

Customer evidence

- 7.138 We asked customers that use SQL Server on the public cloud to identify alternative products which they could use for the same purpose as SQL Server. Customers we received evidence from set out a number of alternatives, mostly listing alternative RDBMS solutions.²⁰⁰³
- 7.139 We also asked those customers to what extent, if at all, they would consider other types of DBMS (relational, non-relational, data analytics services or any other

¹⁹⁹⁹ Microsoft's response to the CMA's information request [REDACTED].

²⁰⁰⁰ Oracle's response to the CMA's information request [REDACTED].

²⁰⁰¹ It explained that for example, in an account system, an RDBMS would be better suited because atomicity and consistency is a paramount priority, while at the same time relations between entities need to be tracked, and that for a social network, a non-RDBMS would be better suited because availability and scalability is of higher importance than eg returning the correct order and amount of comments for a picture that was posted. IBM's response to the CMA's information request [REDACTED].

²⁰⁰² IBM's response to the CMA's information request [REDACTED].

²⁰⁰³ The most frequently listed alternatives were 'Oracle' or Oracle Database, MySQL and PostgreSQL, which are RDBMS solutions.

types of DBMS) as alternatives for SQL Server, considering their organisational use case for SQL Server.²⁰⁰⁴

- 7.140 In response to this question, some customers said that there would be no substitute for their use of SQL Server.²⁰⁰⁵
- 7.141 **RDBMS:** Most customers that responded to this question reported that they would consider other RDBMS as an alternative to SQL Server.²⁰⁰⁶ Some of these customers highlighted that use of alternatives may depend on use case,²⁰⁰⁷ business or application requirements,²⁰⁰⁸ or that there may be considerable inertia when moving away from SQL Server.²⁰⁰⁹
- 7.142 **NRDBMS:** Most customers that assessed whether NRDBMS could be used for SQL workloads submitted that they would not consider NRDBMS as an alternative to SQL Server.²⁰¹⁰ In a handful of cases, customers did consider NRDBMS as alternatives to their use of SQL Server, albeit with the qualification that it would be difficult to replace one with the other.²⁰¹¹
- 7.143 **Data analytics services:** Some customers responded that they did not consider data analytics services as alternatives to their use of SQL Server.²⁰¹² A few customers said that they did consider data analytics services as alternatives to their use of SQL Server.²⁰¹³ In these instances customers explained that, as above, their use of alternatives would depend on the use case or workload.²⁰¹⁴
- 7.144 **Other types of RDBMS:** Most customers that responded to this question said that they did not consider any other types of DBMS as an alternative to SQL Server, did not answer the question or were not aware of any other types of DBMS.²⁰¹⁵ Some customers responded that they did consider data analytics services as alternatives to their use of SQL Server.²⁰¹⁶
- 7.145 Some customers detailed their use of different database management systems in addition to their use of SQL Server.²⁰¹⁷ The use of multiple solutions across a range of customers suggests that customers view different database solutions as

²⁰⁰⁴ Some customers did not answer this question in their response and instead listed the different types of DBMS solutions their organisation uses. Further customer evidence is presented in appendix R.

²⁰⁰⁵ Responses to the CMA's information requests [REDACTED].

²⁰⁰⁶ This included one customer who only uses the product on non-public cloud. Responses to the CMA's information requests [REDACTED].

²⁰⁰⁷ [REDACTED] response to the CMA's information request [REDACTED].

²⁰⁰⁸ Responses to the CMA's information requests [REDACTED].

²⁰⁰⁹ [REDACTED] response to the CMA's information request [REDACTED].

²⁰¹⁰ Responses to the CMA's information requests [REDACTED].

²⁰¹¹ Responses to the CMA's information requests [REDACTED].

²⁰¹² Responses to the CMA's information requests [REDACTED].

²⁰¹³ Responses to the CMA's information requests [REDACTED].

²⁰¹⁴ Responses to the CMA's information requests [REDACTED].

²⁰¹⁵ Responses to the CMA's information requests [REDACTED].

²⁰¹⁶ This included one customer who only uses the product on non-public cloud. Responses to the CMA's information requests [REDACTED].

²⁰¹⁷ Responses to the CMA's information requests [REDACTED].

somewhat suited for specific use cases and/or applications and therefore may be complements rather than substitutes in some circumstances.

The product market for SQL Server

- 7.146 Our view is that alternative DBMS are not effective demand-side substitutes for RDBMS. This is because most customers who provided evidence on this would not generally consider other DBMS solutions (such as NRDBMS) as alternatives to their use of Microsoft SQL Server. Therefore, the relevant product market within which SQL Server sits is no wider than RDBMS.

Market power

Market shares

- 7.147 The table below shows the shares of supply for RDBMS based on our of IDC data (see appendix R) on a global basis across all deployment types from 2019-2023.²⁰¹⁸

Table 7.2: Market shares for RDBMS, all deployment types, 2019-2023

	2019	2020	2021	2022	%
					2023
Microsoft	[20-30]	[30-40]	[30-40]	[30-40]	[30-40]
Oracle	[30-40]	[30-40]	[30-40]	[30-40]	[20-30]
AWS	[0-5]	[5-10]	[5-10]	[5-10]	[5-10]
IBM	[5-10]	[5-10]	[5-10]	[5-10]	[5-10]
SAP	[5-10]	[5-10]	[5-10]	[5-10]	[5-10]
Other	[10-20]	[10-20]	[10-20]	[10-20]	[10-20]

Source: CMA analysis of IDC data [§]. Shares do not sum to 100 due to rounding.

- 7.148 The evidence illustrates that, although there are several alternative suppliers for RDBMS, Microsoft has the largest share considering RDBMS in 2023, and considering all measures its share has remained at between [20-30]% and [30-40]% over the period 2019-2023.²⁰¹⁹
- 7.149 We note that for 2020-2021 Oracle had the largest share in RDBMS. While shares have remained relatively stable over time, Microsoft has slowly gained share over the past few years.

Product characteristics

- 7.150 RDBMSs provide a system to store and retrieve data stored in separate tables.

²⁰¹⁸ See appendix R for market shares submitted by Microsoft.

²⁰¹⁹ As outlined in appendix D.

- 7.151 We consider that there are some limited characteristics with respect to different RDBMSs which might affect how substitutable these products are for some customers. Products within the market for RDBMS offer different levels of scalability, performance and security, usability and compliance with SQL standards which might affect how substitutable they are for individual customers.²⁰²⁰ This may in turn depend on the customer's workload or specific requirements.

Customer evidence on alternatives

- 7.152 In terms of alternatives to SQL Server, customers we heard from generally listed a variety of alternative products including PostgreSQL, Oracle RDBMS, MySQL and others.
- 7.153 We considered the evidence we received from customers who responded to our question about whether they would move away from SQL in response to a 5% price rise:
- (a) Most customers we heard from on this mentioned being unlikely to or having a very small chance of switching away from SQL Server. Reasons included monetary considerations, some software only runs on SQL Server, functionality and skills. Some customers mentioned that there would be a low likelihood of switching away for existing workloads or switching away in the short term, but a higher likelihood of switching away for new workloads or products or in the long run; and a minority of customers outlined that switching away from SQL server would be likely.
 - (b) Reasons customers gave for choosing SQL Server included: internal application landscape or requirements, integration with other Microsoft software or the Microsoft application/server stack, skills, software engineering preferences and functionality.²⁰²¹
- 7.154 The evidence we received from customers shows that few customers were willing and able to switch away from SQL Server.

Provider submissions on alternatives

- 7.155 Key points from submissions received from providers are summarised below.
- (a) Providers submitted that other forms of DBMS were competitors to SQL Server, with some suggesting that individual customers' choice of type of

²⁰²⁰ For example, see [Oracle vs. SQL Server: Head-to-Head Comparison](#), accessed on 18 November 2024, [Comparing Database Management Systems: MySQL, PostgreSQL, MSSQL Server, MongoDB, Elasticsearch, and others](#), accessed on 18 November 2024.

²⁰²¹ Further customer evidence is presented in appendix R.

DBMS depended on factors such as their price or features they offered, and that this could influence the competitor set for specific customers or use cases.

- (b) Providers submitted that barriers to switching would depend on customer-specific and workload-specific factors. One provider suggested that customers can have difficulty switching DBMS; however, another suggested that there are tools available to make this easier.
- (c) Providers submitted that there were limited barriers to entry and expansion.²⁰²²

Conclusions on SQL Server

- 7.156 Microsoft is the largest provider with a moderate but significant stable market share and, although a number of alternatives exist in the market for RDBMS, customers are generally unwilling to switch away from SQL Server to alternative products in response to a small but significant price increase.
- 7.157 While the evidence is mixed, considering the evidence in the round, our view is that Microsoft has a significant degree of market power in relation to SQL Server. This view is strengthened by the links between SQL Server and the other Microsoft software products, outlined by customers and discussed in Product background above.

Microsoft Windows 10/11

- 7.158 For the purposes of this investigation, desktop OSs is the narrowest product market within which Windows 10/11 sits. In the following section, we consider whether the market should be widened to include OSs for servers or OSs for mobile devices. We then consider the extent of any market power held by Microsoft in relation to the relevant market.

Product market definition

Providers' submissions

- 7.159 We asked Microsoft whether there were certain use cases where a server or mobile OS could be used as a substitute for a desktop OS. Microsoft said server OSs can be used to provide 'Desktop-as-a-Service' offerings (ie virtual desktops).²⁰²³ Microsoft also said mobile OSs could be seen as a substitute for

²⁰²² Further provider evidence is presented in appendix R.

²⁰²³ Microsoft's response to the CMA's information request [36].

desktop OSs, eg by a developer of a web browser because web browsing can be done on both types of OSs.²⁰²⁴

- 7.160 In relation to the functionality and intended use for each of the types of OSs compared to a desktop OS, Microsoft said Windows 10/11 (a desktop OS) is used to manage the hardware of a PC device and allows applications to run on it, and Windows Server (a server OS) to manage the hardware of a server device and allows applications to run on it.²⁰²⁵
- 7.161 IBM said it was possible for a server OS to be a substitute for a desktop OS, and gave the example of Windows Server providing virtual desktops to many users, however it said that the degree of substitutability depends on the application and whether the application will sufficiently and effectively run on the desktop operating system.²⁰²⁶

Customer evidence

- 7.162 We asked customers that use Windows 10/11 on the public cloud to identify any alternatives to Windows 10/11 that they could use for the same purpose. Customers we heard from identified other desktop OSs (Linux, MacOS and ChromeOS).²⁰²⁷ MacOS was the most frequently mentioned alternative. Some customers said there were no alternatives.²⁰²⁸
- 7.163 No customer suggested that a server OS would be an alternative. One customer suggested Android would be an alternative, although we understand this was for a specific use case for a particular segment of staff.²⁰²⁹ This suggests that customers do not generally view server and mobile OSs as substitutes for desktop OSs on the demand side.
- 7.164 Overall, these customer responses show that customers are unlikely to switch from a desktop OS to a server or mobile OS, suggesting there is limited demand side substitution from desktop OSs to server or mobile OSs.

The product market for Windows 10/11

- 7.165 Our view is that the relevant product market for Windows 10/11 is the market for desktop OSs and should not be further widened to include server or mobile OSs.

²⁰²⁴ Microsoft's response to the CMA's information request [REDACTED]. We note that Microsoft viewed its PC OS as distinct from its other OSs in the European Commission [Digital Markets Act designation decision](#).

²⁰²⁵ Microsoft's response to the CMA's information request [REDACTED].

²⁰²⁶ IBM's response to the CMA's information request [REDACTED].

²⁰²⁷ This included a few customers who only uses the product on non-public cloud. Responses to the CMA's information requests [REDACTED].

²⁰²⁸ This included a customer who only uses the product on non-public cloud. Responses to the CMA's information requests [REDACTED].

²⁰²⁹ Response to the CMA's information request [REDACTED].

Market power

Market shares

- 7.166 As outlined in appendix R, the evidence illustrates that Microsoft has had a high and stable share of desktop OS:
- (a) For the period 2019 to 2023, Microsoft's share of the global market for desktop OSs ranged between [80-90]% and [90-100]% . In 2023 it was [90-100]%.
 - (b) Its share is much higher than the next-closest competitor, Google, which had [5-10]% in 2023.

Product characteristics

- 7.167 We have considered evidence on the product characteristics of Windows 10/11 which suggests there is some lack of product substitutability between different types of desktop OS:
- (a) Apple's MacOS is differentiated from Windows 10/11 because it comes preinstalled on Apple hardware, is proprietary, and may be harder to run on non-Apple products.²⁰³⁰
 - (b) A Microsoft internal document described points of difference between Windows 11 and MacOS, such as [REDACTED].²⁰³¹
 - (c) Google's ChromeOS is differentiated from Windows 10/11 because it comes pre-installed on Chromebooks and generally minimal data can be stored on the hardware. Applications are accessed through the web browser, and not all applications can be run.²⁰³²
 - (d) A Microsoft internal document described points of difference between Windows 11 and ChromeOS, such as [REDACTED]. It did not compare Windows 11 to any Linux distributions.²⁰³³
 - (e) A few customers told us using MacOS or ChromeOS would require specific types of hardware.²⁰³⁴
 - (f) Various Linux distributions also provide desktop OS. These are open-source, can be installed on a variety of hardware and may be free or contain

²⁰³⁰ Lakka, S et al. (2012), 'Competitive dynamics in the operating systems market: Modelling and policy implications', Technological Forecasting & Social Change, pages 88-105.

²⁰³¹ Microsoft's response to the CMA's information request [REDACTED].

²⁰³² Miller, MR (2019), My Google Chromebook, 4th Edition, Que Publishing, page 12.

²⁰³³ Microsoft's response to the CMA's information request [REDACTED].

²⁰³⁴ Responses to the CMA's information requests [REDACTED].

proprietary software.²⁰³⁵ A customer said they may need to re-architect applications when moving to a Linux desktop OS from Windows 10/11, if they could run on a non-Windows OS at all.²⁰³⁶

- (g) A Microsoft internal document did not compare Windows 11 to any Linux distributions.²⁰³⁷ One interpretation of this could be that Microsoft considers Windows 11 to be substantially differentiated from them.

7.168 In order to interpret whether lack of product substitutability by functionality may act as a source of market power, we used customers' reasons for choosing the Microsoft products to understand whether customers value the differences, ie whether these differences are relevant drivers of choice.

7.169 We asked customers to explain the reasons they chose Windows 10/11 rather than the alternatives they listed:

- (a) Many customers that responded said a reason was support/compatibility/integration with a large range of applications, including integration into other Microsoft software such as Microsoft 365 and applications they use are Windows dependent or supported.²⁰³⁸
- (b) Several customers said a reason was usability/ user familiarity.²⁰³⁹
- (c) A few customers said a reason was security capabilities,²⁰⁴⁰ one customer said a reason was using the virtualised and hosted desktop capability,²⁰⁴¹ and another said the features suit its intended use.²⁰⁴²

7.170 This evidence suggests that there are some indications that Windows 10/11 is used for different purposes to the next-closest products, MacOS and ChromeOS, for example to run, support or integrate with Windows dependent applications. This suggests that Windows 10/11 is differentiated and alternative desktop OSs may therefore be less close substitutes.

Customer evidence on alternatives

7.171 Evidence we received from customers showed that:

²⁰³⁵ Dalheimer, MK and Welsh, M (2005), [Running Linux, 5th Edition](#), O'Reilly.

²⁰³⁶ [REDACTED] response to the CMA's information request [REDACTED].

²⁰³⁷ Microsoft's response to the CMA's information request [REDACTED].

²⁰³⁸ This included a few customers who use the product on the non-public cloud. Responses to the CMA's information requests [REDACTED].

²⁰³⁹ This included a few customers who use the product on the non-public cloud. Responses to the CMA's information requests [REDACTED].

²⁰⁴⁰ This included a few customers who use the product on the non-public cloud. [REDACTED] response to the CMA's information request [REDACTED].

²⁰⁴¹ [REDACTED] response to the CMA's information request [REDACTED].

²⁰⁴² [REDACTED] response to the CMA's information request [REDACTED].

- (a) Most customers we heard from on this said they would be unlikely to move away from Windows 10/11 in response to a 5% price rise. Their reasons included effort and cost of re-architecture, re-training, loss of functionality and lack of support for some applications.
- (b) In terms of alternatives to Windows 10/11, most customers that responded on this listed Linux and MacOS, and some said there were no alternatives.
- (c) Reasons these customers gave for choosing Windows 10/11 included: staff preference and skillset, support for required applications, compatible with a wide range of hardware, required by other applications and significant cost to move to an alternative.²⁰⁴³

Provider submissions on alternatives

7.172 Key points from submissions received from providers are summarised below.

- (a) Microsoft considers other forms of desktop OS as competitors to Windows 10/11.
- (b) Microsoft internal documents showed evidence of limited substitutability between Windows 10/11 and ChromeOS.
- (c) There is some evidence of network effects with respect to Windows 10/11, as Microsoft submitted that a key consideration for a customer choosing a desktop OS was whether there were a significant number of third party applications available to run on it.²⁰⁴⁴

Conclusions on Windows 10/11

7.173 Microsoft has a significant degree of market power in relation to Windows 10/11. This is because Windows 10/11 is differentiated from the next-closest products, has a very large share of the desktop OS market and customers are unwilling or unable to switch away in response to a small but significant price increase. This view is strengthened by the links between Windows 10/11 and other Microsoft software outlined by customers.

Microsoft's productivity suites

7.174 Microsoft has various packages of products which provide some productivity functionality. For the purposes of this investigation, we consider solutions only for enterprise consumers.

²⁰⁴³ Further customer evidence is presented in appendix R.

²⁰⁴⁴ Further evidence is presented in appendix R.

- 7.175 Customers use a variety of different packages under the ‘Microsoft 365’ label, including Office 365, Microsoft 365 Apps for business and various enterprise Microsoft 365 packages.
- 7.176 For the purposes of this investigation, we have treated productivity suites²⁰⁴⁵ for enterprise as the narrowest product market which the Microsoft productivity products sit within.
- 7.177 In the following section, we consider whether the relevant market is wider than productivity suites. Considering product functionality, the next-closest alternative which would perform the functionality of a productivity suite is a ‘mix and match’ approach considering different applications which, combined, would perform similar functionality to the Microsoft suites of products. We then consider the extent of any market power held by Microsoft in relation to the relevant market.

Product market definition

Customer evidence

- 7.178 Almost all of the customers we contacted had some usage of Microsoft 365 for Enterprise.²⁰⁴⁶ A few customers reported use of Office.²⁰⁴⁷
- 7.179 We asked customers for alternatives to the Microsoft products that they could use for the same purpose.²⁰⁴⁸ These customers almost exclusively listed only alternative productivity suites, rather than individual products.
- (a) Almost all customers that responded to this listed Google Workspace as an alternative which they could use for the same purpose as Microsoft 365.²⁰⁴⁹ Most of these customers listed no other alternatives for Microsoft 365.²⁰⁵⁰
 - (b) A few customers listed a component of Google Workspace (Google Docs) as an alternative.²⁰⁵¹ In addition, one customer mentioned Microsoft Office on-

²⁰⁴⁵ We consider productivity suites at a minimum to cover word processing, presentation and spreadsheet functionalities, however we note that most suites include a number of applications beyond these core functionalities.

²⁰⁴⁶ Customers reported using the E5, E3, E1, F3 and A3 Microsoft 365 for Enterprise packages. The Microsoft packages are outlined in [Microsoft's package comparison pages](#). This included a few customers who only use the product on non-public cloud. Responses to the CMA's information requests [§].

²⁰⁴⁷ We consider customers to use the Office package where this is the use of the desktop applications not as part of Microsoft 365 packages. This included one customer who only uses the product on non-public cloud. Responses to the CMA's information requests [§].

²⁰⁴⁸ CMA's information request [§].

²⁰⁴⁹ This included a few customers who only use the product on non-public cloud. Responses to the CMA's information requests [§].

²⁰⁵⁰ This included a few customers who only use the product on non-public cloud. Responses to the CMA's information requests [§].

²⁰⁵¹ This included a few customers who only use the product on non-public cloud. Responses to the CMA's information requests [§].

premises (desktop installed apps) as an alternative.²⁰⁵² Some customers also listed open-source productivity suites as alternatives to Microsoft 365.²⁰⁵³

- (c) One customer also listed substitutes for component elements of Microsoft 365 including substitutes for security services and eDiscovery services.²⁰⁵⁴
- (d) One customer also listed an alternative productivity application (Click Up).^{2055,2056}
- (e) No customers listed a complete mix and match solution which included individual applications which would fully cover Microsoft 365 functionality.

7.180 When we asked customers for alternatives to Office that they could use for the same purpose, they generally listed a subset of the above responses.

The product market for productivity suites

7.181 The relevant market is no wider than productivity suites for enterprise. This is because customer responses do not support widening the market to include other productivity applications which cover only some functionality of the Microsoft packages.

Market power

Market shares

7.182 We have considered three measures of share of supply for productivity suites which indicate Microsoft's share.²⁰⁵⁷ All measures illustrate that, for the time periods available, Microsoft has held an [80-90]% share of supply in each year in the global market for productivity suites.²⁰⁵⁸

²⁰⁵² [REDACTED] response to the CMA's information request [REDACTED].

²⁰⁵³ This included a few customers who only use the product on non-public cloud. Responses to the CMA's information requests [REDACTED].

²⁰⁵⁴ [REDACTED] response to the CMA's information request [REDACTED].

²⁰⁵⁵ We do not consider productivity applications which are included in a suite and not available for purchase separately as alternative productivity applications, as an enterprise customer would purchase this as part of a productivity suite.

²⁰⁵⁶ ClickUp is an application which combines multiple functionalities such as documents, project management and communication into one application. [REDACTED] response to the CMA's information request [REDACTED].

²⁰⁵⁷ The different measures of share of supply for productivity suites on a global basis are: (i) shares by revenue calculated from an industry report submitted by Microsoft; (ii) shares by revenue calculated from data submitted by Microsoft and Google; (iii) shares by count of users based on data submitted by Microsoft and Google;

²⁰⁵⁸ The CMA also considered shares by revenue submitted by Google based on [REDACTED] which calculates Google's market share across SMB, Corporate and Enterprise and Select segments. This analysis estimates that Google Workspace holds a share of (a) [10-20]% for the [REDACTED] (b) [0-5]% for the [REDACTED]; and (c) [0-5]% for the [REDACTED]. See appendix R for more detail.

Product characteristics

- 7.183 We considered evidence we received from customers and software providers to understand the product characteristics of Microsoft 365.
- 7.184 Microsoft 365 is somewhat differentiated from its next-closest competitor, Google Workspace. For example:
- (a) there are different products included in the Microsoft packages (including additional applications, security and advanced identity and access management functionality with Microsoft 365); and
 - (b) Google Workspace is only available through a browser.
- 7.185 To interpret whether lack of product substitutability by functionality may act a source of market power, we used customers' reasons for choosing the Microsoft products to understand whether customers value the differences, ie whether these differences are relevant drivers of choice.
- 7.186 We asked customers the reasons for choosing Microsoft 365 over alternatives they had listed.²⁰⁵⁹ Some customers provided reasons that relate to functionality or capabilities of Microsoft 365. Alternatives that lack this functionality may therefore be less substitutable.
- (a) Some customers mentioned they value functionality of Microsoft software and/or the large variety of applications.²⁰⁶⁰
 - (b) Some customers mentioned they value collaboration functionality (for example Microsoft Teams).²⁰⁶¹
 - (c) A few customers mentioned valuing security capabilities.²⁰⁶²
 - (d) One customer mentioned directly valuing device management capabilities,²⁰⁶³ however several more mentioned valuing integration, or having to re-integrate systems as a reason for choosing Microsoft 365.²⁰⁶⁴

²⁰⁵⁹ CMA's information request [§].

²⁰⁶⁰ This included a customer who only use the product on non-public cloud. Responses to the CMA's information requests [§].

²⁰⁶¹ We note since these responses were received [Microsoft has announced that there will no longer be sale of enterprise suites to new subscribers including Teams. Realigning global licensing for Microsoft 365](#), accessed on 18 November 2024. This included a customer who only uses the product on non-public cloud. Responses to the CMA's information requests [§].

²⁰⁶² This includes customer who only uses the product on non-public cloud. Responses to the CMA's information requests [§].

²⁰⁶³ This customer only uses the product on non-public cloud. [§] response to the CMA's information request [§].

²⁰⁶⁴ This included a few customers who only use the product on non-public cloud. Responses to the CMA's information requests [§].

7.187 The above evidence suggests that there are some indications that Microsoft 365 is used for different purposes to the next-closest products, for example to use the Microsoft suite of products and the variety of functionality those products include.

Customer evidence on alternatives

7.188 Evidence received from customers²⁰⁶⁵ showed that:

- (a) Almost all customers we contacted who use Microsoft 365 said they were unlikely or had a very small chance of switching away.
- (b) Reasons included that alternatives had reduced functionality, it would involve switching costs (for example, re-training, loss of compatibility or integration with other software, cost of re-architecture), and because they wanted to use the same solution as other companies.

Providers' submissions on alternatives

7.189 Key points from submissions received from providers are summarised below:

- (a) [REDACTED].
- (b) Microsoft said that it does not believe there are any barriers to switching to an alternative product, with the reason that all the file formats relied upon by Microsoft 365 Apps are documented and supported such that other productivity solutions can open the files and use them in their applications. We note that Microsoft's response contrasts with our own customer evidence outlined above and in appendix R.²⁰⁶⁶
- (c) Providers submitted that factors such as regulatory requirements, development cost (sunk and ongoing) economies of scale do not act as barriers to entry or expansion in the supply of the Microsoft products or competitors.²⁰⁶⁷

Conclusions on productivity suites

7.190 Microsoft has a significant degree of market power in relation to its productivity suites. This is because Microsoft has a very high and stable market share and there are limited competitive alternatives to the Microsoft packages and customer evidence suggests that customers are unwilling or unable to switch away in response to a small but significant price increase.

²⁰⁶⁵ Further customer evidence is presented in appendix R.

²⁰⁶⁶ Further provider evidence is presented in appendix R.

²⁰⁶⁷ Responses to the CMA's information requests [REDACTED].

Microsoft Visual Studio

- 7.191 For the purposes of this investigation, we have treated IDEs specialised in Windows development as the narrowest plausible candidate product market, because customer evidence (see below) suggested that one reason customers choose to use Visual Studio is because they want to develop applications to run in the Windows environment. In the following section we consider whether the market should be widened to consider all IDEs. We then consider the extent of any market power held by Microsoft in relation to the relevant market.

Product market definition

Provider and stakeholder submissions

- 7.192 Microsoft explained that Visual Studio can be used for building applications to run on non-Windows environments (in addition to Windows environments).²⁰⁶⁸
- 7.193 In response to our Licensing working paper, only one stakeholder commented on this market definition and agreed with our view.²⁰⁶⁹

Customer evidence

- 7.194 We asked customers that use Visual Studio on the public cloud to identify any alternatives to Visual Studio that they could use for the same purpose.
- (a) Some customers we contacted listed one other IDE focused on Windows development: Visual Studio Code.²⁰⁷⁰ As this is a Microsoft product, is provided for free, and includes less functionality as it is a lightweight rather than heavyweight IDE, we consider it does not present a strong competitive constraint to Visual Studio, and we will not consider it separately from Visual Studio.
 - (b) Some customers listed other IDEs for non-Windows development or cross platform IDEs: JetBrains' Rider, JetBrains' IntelliJ IDEA, JetBrains' PyCharm, Eclipse, Apache NetBeans, Xcode, Android Studio and Github's Atom.²⁰⁷¹ Some of these are specialised eg Xcode is specialised in Apple development.

²⁰⁶⁸ Microsoft's response to the CMA's information request [🔗].

²⁰⁶⁹ [Dr George R Barker, Comment on The UK Competitive Market Authority's \(CMAs\) Cloud Services Market Investigation Updated Issues Paper and Working Papers 4-6 on The Supply of Public Cloud Infrastructure Services In the UK Covering The CMA's 1. Updated issues statement on Public cloud infrastructure services market investigation: 2. Licensing Practises Working Paper; 3. Technical Barriers; and 4. Potential Remedies, page 44.](#)

²⁰⁷⁰ This included one customer who only used the product on non-public cloud. Responses to the CMA's information requests [🔗].

²⁰⁷¹ Responses to the CMA's information requests [🔗].

(c) Other customers that responded listed examples of software that are slightly different to IDEs, eg VIM (a lightweight text editor) and GitHub (a developer platform).²⁰⁷²

(d) A few customers said there were no alternatives.²⁰⁷³

7.195 Overall, this suggests there could be a wide range of alternative products to Visual Studio.

7.196 We asked customers for more information to understand their views on alternatives to Visual Studio. We asked customers to tell us whether and to what extent they would consider an IDE tailored to other types of software development, eg Java, to be a substitute for Visual Studio.

(a) Some customers that responded said they would not, or would be unlikely to, consider another IDE to be a substitute for Visual Studio.²⁰⁷⁴ Reasons included the effort of re-training, impact on staff recruitment/hiring, effort to re-integrate with other software, that Microsoft provides good support, Visual Studio integrates well, strategic alignment with Microsoft and other IDEs are less functional.

(b) Some customers said they would consider other IDEs to be a substitute for Visual Studio, for example Eclipse and IntelliJ were mentioned as alternatives.²⁰⁷⁵ One customer said Visual Studio is still required for some purposes.²⁰⁷⁶

(c) Some other customers did not express a strong opinion, explaining they tend to use the IDE most suited to each task or let the developer choose their preferred tool.²⁰⁷⁷

7.197 This suggests that customers have mixed views on whether an IDE tailored for non-Windows development would be a good substitute for Visual Studio.

7.198 We asked customers to tell us whether they can use an IDE tailored to other types of software development, eg Java, to build applications for the Windows environment. The purpose of this was to explore if Visual Studio is the only IDE that has capability to build Windows applications.

²⁰⁷² This included a few customers who only use the product on non-public cloud. Responses to the CMA's information requests [X].

²⁰⁷³ This included one customer who only use the product on non-public cloud. Responses to the CMA's information requests [X].

²⁰⁷⁴ This included a few customers who only use the product on non-public cloud. Responses to the CMA's information requests [X].

²⁰⁷⁵ This included a few customers who only use the product on non-public cloud. Responses to the CMA's information requests [X].

²⁰⁷⁶ This customer only uses the product on non-public cloud. [X] response to the CMA's information request [X].

²⁰⁷⁷ Responses to the CMA's information requests [X].

- (a) Some customers said this was possible;²⁰⁷⁸ some spoke positively of it; one said it was common;²⁰⁷⁹ one said that it gives them access to more solutions;²⁰⁸⁰ one was less positive, saying there would be fewer features.²⁰⁸¹
- (b) Some indicated it may be possible to an extent, but they do not do this;²⁰⁸² one of these explained that the efficiency of the developer would be reduced.²⁰⁸³ Another customer said they cannot use another IDE for development of Windows applications built in C#/VB.net which we understand are programming languages developed by Microsoft and closely associated with Microsoft technology.^{2084, 2085}

7.199 This suggests that other IDEs can be used for Windows development, but Visual Studio may be the most well suited for this.

7.200 Customers view IDEs not specialised in Windows development as alternatives to Visual Studio. IDEs that are not specialised in Windows development can still be used for Windows development, and customers have mixed views on whether they would consider an IDE tailored for non-Windows development to be a good substitute for Visual Studio. In addition, Microsoft explained that Visual Studio can also be used for building applications to run on non-Windows environments. Therefore, there does not seem to be a good reason to draw a line between IDEs specialised in Windows development, and those that do not.

The product market for Visual Studio

7.201 The relevant product market is the market for IDEs.

Market power

Market shares

7.202 We have not been able to gather reliable data on Microsoft Visual Studio's market share in the market for IDEs. However, we have two sources of related²⁰⁸⁶ evidence:

- (a) The first is Visual Studio's market share in a market for Development Languages, Environments, and Tools (DLET). This market is much broader

²⁰⁷⁸ This included one customer who only uses the product on non-public cloud. Responses to the CMA's information requests [X].

²⁰⁷⁹ [X] response to the CMA's information request [X].

²⁰⁸⁰ [X] response to the CMA's information request [X].

²⁰⁸¹ [X] response to the CMA's information request [X].

²⁰⁸² This included a few customers who only use the product on non-public cloud. Responses to the CMA's information requests [X].

²⁰⁸³ [X] response to the CMA's information request [X].

²⁰⁸⁴ This customer only uses the product on non-public cloud. [X] response to the CMA's information request [X].

²⁰⁸⁵ 'The history and legacy of Visual Basic - Ryan Lucas', accessed on 18 November 2024.

²⁰⁸⁶ We explain the caveats associated with each of these in the appendix R.

than just IDEs; therefore this is likely to be a very considerable underestimate of Microsoft's market share in relation to the supply of IDEs;

- (b) For the second metric we refined the DLET data to construct a global market for IDEs. This showed Microsoft Visual Studio had the largest share in 2023, with [30-40]%.

Product characteristics

- 7.203 We considered the evidence we received from customers and software providers to understand the product characteristics of Visual Studio. Evidence we have seen suggests that there are multiple substitute products for Visual Studio. At a product level, IDEs perform the same basic functionality, enabling software developers to build applications, web pages or services.
- 7.204 To interpret whether lack of product substitutability by functionality may act as a source of market power, we use customer reasons for choosing the Microsoft products to understand whether customers value the differences, ie whether these differences are relevant drivers of choice.
- 7.205 We asked customers to explain the reasons they chose Visual Studio rather than the alternatives they listed. Many customers that responded said that they value the additional functionality of Visual Studio.²⁰⁸⁷ Alternatives that lack this functionality may therefore be less substitutable. Examples of specific functionality that customers valued included wanting to develop .Net code (for Windows environments),²⁰⁸⁸ and integrations with other Microsoft products.²⁰⁸⁹
- 7.206 The above evidence suggests that there are some indications that Visual Studio is used for different purposes to the next-closest products, for example developing applications for Windows environments. Therefore, this suggests Visual Studio is differentiated and alternative IDEs may not be close substitutes.

Customer views on alternatives

- 7.207 Evidence received from customers showed the following (a detailed discussion of the customer evidence is presented in appendix R):
 - (a) Some customers that responded said they would be unlikely to move away from Visual Studio (or Visual Studio Code) in response to a 5% price rise. Reasons given included: cost of change, integrations with other software

²⁰⁸⁷ This included a few customers who only use the product on non-public cloud. Responses to the CMA's information requests [X].

²⁰⁸⁸ This included one customer who only uses the product on non-public cloud. Responses to the CMA's information requests [X].

²⁰⁸⁹ This included one customer who only uses the product on non-public cloud. Responses to the CMA's information requests [X].

including Microsoft software, more or desired functionality, cost of re-training staff, little perceived benefit, existing investment, additional licensing cost, codes would need to be re-written and Visual Studio is best for Windows development.

- (b) Reasons customers gave for choosing Visual Studio included: wanting to build in the Windows environment, integrations with other Azure and other Microsoft products, staff skills, functionality, legacy use, it gets updates and optimisations for .NET development faster than other IDEs, it is bundled with other Microsoft licences, and it is an embedded product.
- (c) Some customers that responded said there would be significant or some switching costs when switching from Visual Studio to an alternative IDE including: retraining, impact on developer efficiency, staff recruitment issues and others. Some customers said there would be no or low switching costs.

Provider submissions on alternatives

7.208 Key points from submissions received from providers are summarised below.

- (a) Microsoft has submitted that its main competitors in supplying Visual Studio are other IDEs.
- (b) Microsoft has submitted that it does not believe there are any major barriers to switching IDE.
- (c) Microsoft has submitted that it does not believe there are any material barriers to entry or expansion for developer tools.²⁰⁹⁰

Conclusions on Visual Studio

7.209 While there are a wide variety of alternative IDEs in the market, Visual Studio is differentiated, many customers are unwilling or unable to switch away in response to a small but significant price increase for some use cases, there are various barriers to switching including cost and staff re-training, and Microsoft has a moderate share in the IDE market (and the largest of all providers).

7.210 Microsoft has significant market power in relation to Visual Studio. This view is strengthened by the links between Visual Studio and other Microsoft software outlined by customers.

²⁰⁹⁰ Further provider evidence is presented in appendix R.

Assessment of any cumulative effect of Microsoft's market power across markets

- 7.211 The key factors we explored when considering any additional cumulative effect of any market power are:
- (a) how the Microsoft products are sold or purchased; and
 - (b) actual or perceived technical benefits or limitations to using the Microsoft products together.
- 7.212 These factors may make customers more likely to use multiple Microsoft products by impacting customer decision making in two ways: when a customer is selecting a software product for the first time; and by increasing barriers to switching.²⁰⁹¹
- 7.213 Cumulative effects from links between the products and evidence which may impact customer switching relating to individual products have already been taken into account in the individual market power assessments above.
- 7.214 This section builds on those views and sets out additional evidence which relates to the cumulative effect of market power across markets where this is not specific to one product, and which may strengthen our views on the finding of market power in individual product markets. This evidence includes relevant customer and provider evidence in particular on Enterprise Agreements (EAs).
- 7.215 We asked customers how Microsoft products were purchased. For example, we asked whether products being included together in packages/suites or purchased under framework agreements (with Microsoft or other third parties), influenced their consumption decisions at the point of purchase.
- 7.216 Customers typically responded with reference to Microsoft 365 containing a variety of different applications, or with reference to their EA with Microsoft.
- (a) Most customers we heard from on this said their purchasing decisions were impacted by the way the Microsoft products are sold or purchased.²⁰⁹² For example, one customer said non-Microsoft alternatives to Microsoft 365 (which includes productivity, compliance and security capabilities) were less attractive because of bundling, but non-Microsoft alternatives to Windows

²⁰⁹¹ For example, if a customer is looking to purchase a new software product (eg productivity software), and it already uses a different Microsoft product (eg desktop OS), it may be more likely to select the Microsoft productivity product if it is easier and/or cheaper to buy them together. Having chosen the Microsoft desktop and productivity products, a customer may be less willing to consider alternative desktop products in future if it thinks the functionality of the Microsoft productivity product could be reduced by switching away from Microsoft's desktop product.

²⁰⁹² Responses to the CMA's information requests [X].

Server, SQL Server and Visual Studio were not impacted because capability is more discrete.²⁰⁹³

- (b) Some customers said that their purchasing decisions were not impacted but did not explain why.²⁰⁹⁴
- (c) Others said their purchasing decisions were not impacted because they do not have an Enterprise Agreement.²⁰⁹⁵

7.217 We asked a similar question regarding whether the same factors act as barriers to switching away from the Microsoft products.

7.218 Most customers that responded to this said these factors could act as a barrier to switching,²⁰⁹⁶ however, most of these considered other factors to be more important.²⁰⁹⁷

7.219 As outlined above, where we asked customers their reasons for choosing the Microsoft products, and whether they were likely or unlikely to switch to an alternative they had listed, some customers proactively raised technical benefits or limitations in response. In addition, we asked all customers directly whether there are any actual or perceived technical benefits when using more than one of Microsoft's software products, or limitations when using one of Microsoft's software products without one another (or with alternatives).

- (a) All customers we heard from on this said there were technical benefits to using Microsoft products together, for example customers expected a high level of interoperability,²⁰⁹⁸ skills and staff training can be transferred between tools,²⁰⁹⁹ and industry-wide adoption of Microsoft makes document sharing easier.²¹⁰⁰
- (b) Most customers that responded to this thought there were technical limitations when using one of Microsoft's software products without one another (or with alternatives).²¹⁰¹ Examples of limitations included extra work and infrastructure being required to integrate non-Microsoft software,²¹⁰² functionality of some Microsoft products may be reduced when using non-Microsoft products with it (eg using Zoom with Outlook),²¹⁰³ or using M365

²⁰⁹³ [REDACTED] response to the CMA's information request [REDACTED].

²⁰⁹⁴ Responses to the CMA's information requests [REDACTED].

²⁰⁹⁵ Responses to the CMA's information requests [REDACTED].

²⁰⁹⁶ Responses to the CMA's information requests [REDACTED].

²⁰⁹⁷ Responses to the CMA's information requests [REDACTED].

²⁰⁹⁸ Responses to the CMA's information requests [REDACTED].

²⁰⁹⁹ Responses to the CMA's information requests [REDACTED].

²¹⁰⁰ [REDACTED] response to the CMA's information request [REDACTED].

²¹⁰¹ Responses to the CMA's information requests [REDACTED].

²¹⁰² Responses to the CMA's information requests [REDACTED].

²¹⁰³ [REDACTED] response to the CMA's information request [REDACTED].

without Windows 10/11,²¹⁰⁴ security vulnerabilities and ineffective collaboration.²¹⁰⁵

- (c) Some customers said there were no or minimal technical limitations to using Microsoft products without one another or with alternatives.²¹⁰⁶ For example, one customer said it expected most non-Microsoft software would integrate with Active Directory/ Entra ID.²¹⁰⁷

Provider submissions

- 7.220 In response to our Licensing working paper, some parties we contacted commented on the cumulative effect of market power across markets.
- 7.221 An academic said that economies of scope in consumption are clearly benefits for consumers.²¹⁰⁸
- 7.222 Google supported the idea of Microsoft having some level of cumulative effect of market power across markets: ‘We also agree with the CMA that customer feedback relating to individual Microsoft software products may “understate the overall extent of [Microsoft’s] market power”’. Google considered SQL Server, Windows 10/11, Active Directory and Microsoft 365 are tightly integrated with Windows Server and reinforce its market power.²¹⁰⁹
- 7.223 A cloud provider submitted that there are technical difficulties in fully replacing Microsoft products with a fully functional alternative, especially where those products integrate across multiple different workloads. The integration of these ‘horizontal’ services often results in customers encountering technical difficulties given that, by their very nature, such services will usually have numerous dependencies on other parts of a customer’s IT set-up (eg Active Directory, .NET framework).²¹¹⁰
- 7.224 The SMF said that there may be an issue with the way Microsoft sells its products. It submitted that another way in which a cloud software provider may harm competition is through the bundling of different services into a package. It is alleged that providers with a dominant market position offer their services at a lower total price than competitors that do not have a full range of services or operate across the full cloud stack, giving them an advantage.²¹¹¹

²¹⁰⁴ [redacted] to the CMA’s information request [redacted].

²¹⁰⁵ [redacted] response to the CMA’s information request [redacted].

²¹⁰⁶ Responses to the CMA’s information requests [redacted].

²¹⁰⁷ [redacted] response to the CMA’s information request [redacted].

²¹⁰⁸ [Dr George R Barker’s response to the Licensing and Potential remedies working papers and Updated Issues statement](#) dated 06 June 2024, page 52.

²¹⁰⁹ Google’s submission to the CMA [redacted].

²¹¹⁰ [redacted] response to the CMA’s information request [redacted].

²¹¹¹ [SMF response to the Licensing working paper](#), page 30.

Enterprise Agreements

- 7.225 Microsoft submitted that EAs provide customers with a range of benefits, including: (i) access to lower prices, with discounts typically of between 15% and 45% available depending on the size of the agreement; (ii) greater certainty, by locking in prices over the course of the agreement; and (iii) greater flexibility of payment terms (eg equal monthly payments, deferred payments, or ramped payments) and the ability to scale use up or down depending on circumstances.²¹¹²
- 7.226 In this section we consider additional evidence relating to Enterprise Agreements, and whether this affects customer choice of Microsoft software products.

Customer evidence

- 7.227 The Jigsaw report sets out various factors influencing customer purchasing decisions around choice of cloud, one of which is Enterprise Agreements.²¹¹³ The report highlights that these agreements are perceived by some as good value compared to the cost of other software licences and services; however, some perceive that Enterprise Agreements might create a dependence on the use of a wide range of services, long-term legacy use and the in-house skills and experience that they had developed.²¹¹⁴
- 7.228 In response to our Licensing working paper, SMF highlighted the role of contracts and agreements. It submitted that, because its software is embedded across most UK offices, one participant in semi-structured interviews conducted by SMF said companies like Microsoft capitalise on their positions in order to impose ‘umbrella agreements’ with highly integrated contracts across entities and services and this points to the embedded nature of legacy providers within IT ecosystems, which can result in challenges.²¹¹⁵
- 7.229 We gathered customer evidence to assess the importance of Enterprise Agreements in customer purchasing decisions.
- 7.230 We heard from customers who had an Enterprise Agreement.²¹¹⁶ Some Enterprise Agreements lasted five years²¹¹⁷ and some three years.²¹¹⁸ Most of these customers had at least two of the relevant Microsoft software products included as part of their Enterprise Agreement.²¹¹⁹

²¹¹² Microsoft’s response to the CMA’s information request [redacted].

²¹¹³ [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), page 29.

²¹¹⁴ [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), page 29.

²¹¹⁵ [SMF response to the Licensing working paper](#), page 28.

²¹¹⁶ Responses to the CMA’s information requests [redacted].

²¹¹⁷ Responses to the CMA’s information requests [redacted].

²¹¹⁸ Responses to the CMA’s information requests [redacted].

²¹¹⁹ Responses to the CMA’s information requests [redacted].

7.231 We asked customers about the benefits and disadvantages of Enterprise Agreements:

- (a) Customers we heard from generally said that price certainty,²¹²⁰ flexibility²¹²¹ and discounts compared to list prices²¹²² were the key benefits of Enterprise Agreements.
- (b) Customers that responded generally said that the disadvantages of Enterprise Agreements included restricting the ability to discontinue use or fall below minimum purchase requirements²¹²³ and a high commitment level over the term.²¹²⁴

7.232 We asked customers whether their Enterprise Agreement or the prospect of an Enterprise Agreement influenced their software purchasing decision initially (by influencing the number of licences or products purchased) or whether their Enterprise Agreement acts as a barrier to switching to alternative software products for some of their demand.

- (a) All customers that responded to this said that their Enterprise Agreement did not influence their consumption at the point of purchase or was not a major factor in their decision.²¹²⁵
- (b) Most customers we heard from on this outlined that their Enterprise Agreement would not act as a barrier to switching to alternative software products.²¹²⁶ Some customers said that it might to the extent that they would not purchase another product to fulfil the same business need as a product included in their Enterprise Agreement.²¹²⁷ One customer outlined that its Enterprise Agreement does act as a barrier to moving to alternative products.²¹²⁸

7.233 The customer evidence suggests that while Enterprise Agreements might prevent some customers from considering alternative products during the term, customers also recognise the benefits of these agreements, as well as their flexibility.

7.234 The evidence shows that Enterprise Agreements are not a key source of market power, and that, absent Enterprise Agreements customers would likely make the same purchasing decisions.

²¹²⁰ Responses to the CMA's information requests [X].

²¹²¹ Responses to the CMA's information requests [X].

²¹²² Responses to the CMA's information requests [X].

²¹²³ Responses to the CMA's information requests [X].

²¹²⁴ Responses to the CMA's information requests [X].

²¹²⁵ Responses to the CMA's information requests [X].

²¹²⁶ Responses to the CMA's information requests [X].

²¹²⁷ Responses to the CMA's information requests [X].

²¹²⁸ [X] response to the CMA's information requests [X].

Provider submissions

- 7.235 Google submitted that Enterprise Agreements create sticky demand, and that Microsoft's frequent use of discounting structures across cloud and non-cloud products – in particular those offered under its Enterprise Agreements – create sticky demand and remove the ability for other cloud providers to compete once a customer has already migrated to Azure.²¹²⁹

Our assessment

- 7.236 Selling products together can have benefits for customers, such as reducing search costs, and service quality and ease of integration are parameters of competition.²¹³⁰ However, we consider that how the Microsoft products are purchased, technical benefits when using Microsoft products with other Microsoft products, and technical limitations in using alternative products with the Microsoft products, may act as sources of market power. This is not accounted for in measures like shares of supply and therefore, considered in isolation, market shares for the individual software products may understate Microsoft's market power. We consider that customer switching may be impacted by these factors, and that our evidence relating to market power with respect to individual products would already capture these additional cumulative effects.
- 7.237 Most customers who responded to this highlighted that the way the Microsoft products are sold or purchased influenced their decisions around consumption of the software, and some reported that this might prevent switching to alternatives. All that responded said there were technical benefits to using Microsoft products together. Most customers who responded to this thought there were technical limitations when using one of the Microsoft's software products without one another.
- 7.238 In addition to the cumulative effects due to links between the Microsoft software products (which are already accounted for in the market power assessments set out by product), the evidence presented in this section does not provide a strong basis for any additional cumulative effects arising from Enterprise Agreements (which is not already accounted for in the individual market power assessments for each product).

View on cumulative effect of market power

- 7.239 Our view is that Microsoft has significant market power in relation to each of the relevant software products.

²¹²⁹ [Google response to the Licensing working paper](#), paragraph 23.

²¹³⁰ [CC3 \(Revised\), Guidelines for market investigations: Their role, procedures, assessment and remedies](#), paragraph 288. As discussed in appendix B: Parameters of Competition.

Modernisation

- 7.240 In response to our PDR, Microsoft told us that modernisation (and AI) is acutely relevant to understanding competition when it comes to the licensing of infrastructure (Windows Server VM) components in a legacy workload.²¹³¹
- 7.241 Modernisation can refer to various forms of transformation of workloads in the cloud including re-writing applications, adopting PaaS services and containerising applications.²¹³² The ability and willingness to switch away from Microsoft software – whether as part of a modernisation programme or otherwise – is relevant to the question of market power. There are various types of modernisation of workloads, some of which do not involve switching away from Microsoft software. For example:
- (a) Containerisation involves the application running the same operating system albeit in a container, and therefore would not involve switching away from Microsoft software.
 - (b) When customers switch to PaaS services, they do not have control of the underlying infrastructure, which may or may not be running on Microsoft software, and they may or may not choose PaaS services which rely on Microsoft Software (for example Cloud SQL for SQL Server²¹³³).
 - (c) When a customer re-writes an application to run on alternative server OS or with an alternative database management system, this is modernisation which switches away from Microsoft software.
- 7.242 This underlines how the term ‘modernise’ is a broad term that includes processes that would not necessarily be expected to dilute market power in software. As such, the aspect of modernisation that we have considered for the purposes of our assessment is that which involves switching away from Windows Server and SQL Server.
- 7.243 We have gathered further evidence relating to ‘modernisation’ in order to assess how this may affect our assessment of market power. In particular, we gathered evidence, including internal documents, from AWS and Google and evidence from customers which use the Microsoft software products on the public cloud. However, we note that some of the evidence set out above may already have reflected that there is a possibility to ‘modernise’ workloads by switching away from the Microsoft software, for example to the extent that customers are aware of

²¹³¹ [Summary of a hearing with Microsoft](#), 05 March 2025, page 6, paragraph 34.

²¹³² TechTarget, [What is cloud modernization? Benefits and best practices](#), accessed on 16 June 2025.

²¹³³ Google, [Cloud SQL for SQL Server features | Google Cloud](#), accessed on 24 June 2025.

the option to modernise when answering questions about switching this would already be taken into account in the assessment above.

Cloud providers' views

- 7.244 In this section we summarise the submissions and documents we have received from Microsoft, AWS and Google relating to modernisation.²¹³⁴
- 7.245 In response to the PDR, Microsoft submitted that Amazon and Google are constantly moving customers off Windows Server, and that customers get the greatest benefit from the cloud when they modernise workloads to use more cloud-native technologies. It submitted that AWS and Google have high incentives to help accelerate customers' cloud journeys including by modernising workloads before bringing them to the cloud and that AWS and Google both advertise these offers.²¹³⁵
- 7.246 Microsoft submitted that AI is a driver for customers to modernise workloads. It submitted that AI is central and relevant to speeding up migration and modernisation and is relevant for the use of legacy software in the cloud. It submitted that Microsoft itself is encouraging customers to migrate and modernize Windows Server, SQL Server, Linux, Oracle, VMware and SAP on-premises workloads to the cloud, to create an AI-ready infrastructure.²¹³⁶
- 7.247 Microsoft also submitted that AI is used by cloud customers to make the process of modernising applications to deploy in the cloud easier.²¹³⁷
- 7.248 AWS and Google submitted that modernisation does not necessarily mean switching away from Microsoft software. In particular:
- (a) Google submitted that Microsoft conflates migration and modernisation initiatives, with the former aimed at encouraging migration by reducing licensing costs for customers' Windows workloads on Google Cloud not switching away from Microsoft software.²¹³⁸ Google submitted that re-writing

²¹³⁴ We asked AWS and Google to set out their programmes and services that enable customers to switch from Microsoft software to alternatives, asking providers to explain the degree to which such programmes have helped manage the cost of licensing Microsoft software, and to explain (including supporting documentation, estimates and/or data) how providers monitor the success of these initiatives. We also requested providers internal documents relating to any counterstrategies AWS and Google adopt due to the licensing practices and received some relevant submissions.

²¹³⁵ [Microsoft's response to the PDR](#), paragraph 230. In support of this statement, Microsoft provided links to public webpages of AWS and Google in which they advertise how they assist customers migrating to the cloud. It is unclear from these whether the modernisation advertised involves switching away from Microsoft (or other proprietary) software. (see Google, [Pricing Overview | Google Cloud](#), accessed on 29 May 2025; AWS, [App Migration Tool - AWS Application Migration Service - AWS](#), accessed on 29 May 2025.)

²¹³⁶ Microsoft's submission to the CMA [§]. In support of this, Microsoft has also submitted some marketing materials which support the fact that it is encouraging customers to migrate and modernise workloads, although it is unclear whether this modernisation involves switching away from Microsoft (or other proprietary) software. Microsoft's submission to the CMA [§].

²¹³⁷ Microsoft's submission to the CMA [§].

²¹³⁸ Google's response to the CMA's information request [§]. We present evidence from Google on migration initiatives below in this chapter.

an application to run on a different type of OS is a particularly difficult type of modernisation which takes an extremely significant investment of cost and time – far more than, for example, containerisation.²¹³⁹ It submitted that containerisation requires the workloads to run the original operating system (albeit in a container).²¹⁴⁰

- (b) AWS submitted that modernisation does not remove customers' dependency on Microsoft software.²¹⁴¹ AWS submitted that most of customers' cost savings in the cloud are derived from right-sizing underutilised resources to make sure customers are not paying for licences they aren't using instead of switching away from Microsoft software.²¹⁴² We address cost savings for customers derived from sources other than modernisation below in other counterstrategies.

7.249 AWS and Google have modernisation initiatives which involve switching away from Microsoft software. AWS and Google provided submissions and internal documents in relation to these modernisation initiatives and have suggested that they have had limited success.

7.250 [X] submitted that customers find it extremely time consuming and complex to modernise their workloads to run on a different OS environment. It said [X].²¹⁴³

7.251 [X] submitted some customer examples where it has been able to successfully modernise workloads away from Windows Server and SQL Server. It explained that these examples had certain commonalities including:

- (a) Time: modernisation of limited workloads took years to achieve [X];
- (b) Cost: all [X] projects entailed significant cost [X];²¹⁴⁴
- (c) Dedicated internal resource;
- (d) Relatively limited/discrete exposure to Microsoft software; and

²¹³⁹ Google submission to the CMA [X].

²¹⁴⁰ Google's response to the CMA's information request [X].

²¹⁴¹ AWS submitted that modernisation refers to steps taken after customers migrate their workloads to a cloud provider, where customers can transform their on-premises workloads into more cost effective, resilient and efficient solutions. It explained that modernisation involves various things such as porting applications from .NET Framework to .NET Core, decomposing monoliths to microservices, implementing DevOps techniques with container and serverless technologies, and using more fit-to-purpose software solutions. AWS' submission to the CMA [X].

²¹⁴² AWS provided a link to a webpage where AWS advertises its assistance for migrating and modernising Microsoft workloads (see AWS, [AWS for Microsoft workloads | Amazon Web Services | AWS](#), accessed on 29 May 2025). This page does advertise that 'by replacing your Microsoft-licensed products and applications with open source alternatives and cloud based technologies, you can lower costs...' It also provided an example business case of a modernisation and migration plan (See AWS, [AWS_Migration Business Case_V250_042920241.pdf](#), accessed on 29 May 2025). This migration plan refers to optimising licenses for Microsoft software but does not directly refer to switching away from Microsoft software. AWS' submission to the CMA [X].

²¹⁴³ [X] submission to the CMA [X].

²¹⁴⁴ [X] response to the CMA's information request [X].

(e) Modernisation was not at the point of first migration.²¹⁴⁵

- 7.252 The other submitted that customers that were dependent on Microsoft software before migrating to [X] remain dependent on Microsoft productivity software after modernisation due to the lack of real alternatives to Microsoft's productivity software. This is because many applications are operating system specific commercial off the shelf solutions.²¹⁴⁶ It submitted that as these applications are typically provided by a third party and are not typically Linux compatible, they therefore cannot be re-written by customers and remain tied to specific operating systems. This means that it is impossible to migrate customers who require the functionality of these products away from these proprietary systems to open-source solutions. It submitted that, despite the fact that it has maintained an average team of [X] dedicated to supporting just Microsoft modernisation initiatives over the past five years, between 2023 and 2024, [X] Windows usage was eliminated as a result of modernisation efforts. This represents [X] decrease in its customers' estimated [X] spend on Microsoft licensing costs on [X].²¹⁴⁷
- 7.253 Google submitted that customers have a strong preference to 'lift and shift' workloads when they migrate to the cloud, meaning that they are likely to continue to rely on Microsoft software.²¹⁴⁸ To support this claim, [X].²¹⁴⁹
- 7.254 Additionally, evidence submitted by Microsoft shows that, for customers that joined between May 2022 and December 2023, [70-80]% of new customers (accounting for [80-90]% of spend) use Windows Server during their first year on Azure.²¹⁵⁰ This suggests that a majority of customers continue to rely on Microsoft software as they migrate to the Azure. While we do not have equivalent evidence for customers new to AWS or GCP, given Azure's market share this suggests that a large proportion of new customers still rely on Microsoft software when they first migrate.
- 7.255 We have also received some internal documents which provide some detail on providers' strategies with respect to modernisation: ²¹⁵¹
- (a) [X] provided one relevant internal document on modernisation which is supportive of its submission that its ability to compete with Microsoft is clearly affected by Microsoft's licensing practices and there has been limited success in customers moving away from Windows Server through

²¹⁴⁵ [X] submission to the CMA [X].

²¹⁴⁶ [X] submitted that according to an [X] survey these applications make up to [X] of customers' Microsoft licensed workload usage.

²¹⁴⁷ [X] submission to the CMA [X]. See below for more details on this internal document. [X].

²¹⁴⁸ Google's submission to the CMA [X].

²¹⁴⁹ Google's submission to the CMA [X].

²¹⁵⁰ Microsoft's submission to the CMA [X].

²¹⁵¹ We asked AWS and Google to provide internal documents that discuss AWS'/GCP's ability to compete with Azure and/or any counterstrategies that AWS/Google could adopt (or has already adopted) in relation to their ability to compete.

modernisation programmes. [X].²¹⁵² However, the document also suggests that customers are shifting from [X] to [X]. The document gave the example [X].²¹⁵³

- (b) [X] provided some [X] internal documents that indicate it uses or plans to use modernisation [X].²¹⁵⁴ For example, one document states that [X] offer is to [X] by identifying customers with [X] that are willing to [X]. [X] credits for these customers.²¹⁵⁵ Another document indicates that, while [X] perceives itself as [X] Azure and that Microsoft is strengthening its licensing stronghold on customers, it sees significant [X] in helping customers modernise [X].²¹⁵⁶

7.256 Google and AWS have made submissions on the success of their initiatives which use AI to help customers switch away from Microsoft software.

7.257 Google submitted that it is using [X] to help customers' [X].²¹⁵⁷ With respect to this initiative, it submitted that:

- (a) The programme has had no material impact to date on [X];
- (b) It is only [X];
- (c) To make use of the programme the customer must [X] which means it can't be used for [X] which it submitted often make up a large proportion [X] of [X].²¹⁵⁸ and
- (d) [X].²¹⁵⁹

7.258 A cloud provider [X] has submitted some customer examples which outline difficulties faced by customers who have attempted modernisation of Windows Server workloads with [X], even with the availability of AI coding tools. These examples include:

- (a) One customer that has a large number of Windows Server based applications and would like to migrate away from the Microsoft ecosystem, but given that no open source alternatives are available for their commercial off the shelf applications they are unable to re-write these applications to run on Linux; and

²¹⁵² [X] response to the CMA's information request [X].

²¹⁵³ [X] response to the CMA's information request [X].

²¹⁵⁴ [X] response to the CMA's information request [X].

²¹⁵⁵ [X] response to the CMA's information request [X].

²¹⁵⁶ [X] response to the CMA's information request [X].

²¹⁵⁷ Google explained that [X].

²¹⁵⁸ Google's response to the CMA's information request [X]; Google's submission to the CMA [X].

²¹⁵⁹ Google's submission to the CMA [X].

- (b) Another customer that has [REDACTED] lines of code to translate from a Windows Server coding framework to one which supports open source, and the dependencies on Windows libraries make this more complex. [REDACTED].²¹⁶⁰

7.259 Google submitted that the process of modernising databases and linked applications away from SQL Server is highly complex, and there is limited scope for AI coding tools to reduce the resulting time and cost burden for customers. It submitted that Google Cloud's Database Migration Service [REDACTED] allows customers to switch from SQL Server to an alternative. It submitted that of the databases scanned for prospective migration, [REDACTED] have been or are currently being migrated.²¹⁶¹ It also submitted that for [REDACTED] [REDACTED] customers [REDACTED] there are additional obstacles which cannot be addressed by AI coding tools including where features are not replicated, skills barriers and operational challenges in adopting alternatives.²¹⁶²

7.260 Google also submitted some customer examples of the [REDACTED]:²¹⁶³

(a) [REDACTED].²¹⁶⁴

(b) [REDACTED].²¹⁶⁵

7.261 AWS submitted that the transformer function of Amazon Q Developer can be used by customers to update their applications and move them to another IT provider at great speed and is cloud/on-premises agnostic. Using the transformation capability of Amazon Q Developer, modernisation teams can deliver large and complex projects accelerating various tasks including .NET porting. It submitted specifically that Amazon Q Developer has been found to accelerate the process of .NET application porting from Windows to Linux up to 4 times faster and reduces operational costs by up to 40%.²¹⁶⁶ It submitted an example of a customer who used Amazon Q Developer to migrate their data validation and business logic from SQL Server with .NET to AWS Lambda. It submitted that the customer was able to do this in less than a day whereas it would have previously taken them weeks to migrate manually.²¹⁶⁷ It is unclear whether AWS' AI modernisation programme's efforts are included in the statistics provided above [REDACTED].

²¹⁶⁰ Google's submission to the CMA [REDACTED].

²¹⁶¹ Google submitted that since June 2023, its Database Migration tool has scanned [REDACTED] SQL Server databases across [REDACTED] SQL Server instances in efforts to migrate the customers database to [REDACTED].

²¹⁶² Google's submission to the CMA [REDACTED].

²¹⁶³ Google's submission to the CMA [REDACTED].

²¹⁶⁴ Google submitted that [REDACTED].

²¹⁶⁵ Google's submission to the CMA [REDACTED].

²¹⁶⁶ See AWS, [AWS Transform for .NET | AWS](#), accessed on 29 May 2025.

²¹⁶⁷ AWS' submission to the CMA [REDACTED].

Customers' views

- 7.262 In this section we set out evidence gathered from calls with customers that use Microsoft software on the public cloud.²¹⁶⁸ We asked customers whether they had considered modernisation and, if so, the extent to which they have modernised workloads to switch away from Microsoft software. If they had modernised their workloads, we asked them to give a rough estimate of how easy or difficult this was (in terms of time and/or engineering effort for example). We also asked some customers if there were any workloads which they would not be able to switch away from Microsoft software. In line with the qualitative nature of the evidence we gathered, we have given a narrative summary of the key points that emerged from the evidence. A more detailed summary of this evidence can be found in appendix R: Additional market power evidence.
- 7.263 Some customers have undertaken little to no modernisation, are exploring or assessing their suitability for modernisation or explained that they have an ad-hoc approach to modernisation.²¹⁶⁹
- (a) Where these customers explained the mode of modernisation they were adopting, customers generally referred to adopting PaaS or SaaS solutions in the cloud.²¹⁷⁰
 - (b) In these instances, customers either did not mention whether this would involve re-engineering their estate to work with alternatives to the Microsoft software or explicitly explain that they would not consider that as an option.²¹⁷¹
 - (c) For example, one customer explained that the extent to which it has been able to modernise depends on the application. Its focus is moving services from on premises with the existing architecture and modernisation will come with time. It added that when it moves a workload out of on-premises it will either move it to SaaS if it is available for the type of workload or it will move the workload to the cloud as it is, due to time constraints.²¹⁷²
- 7.264 Other customers mentioned adopting larger amounts of PaaS and/or SaaS in the cloud.²¹⁷³
- 7.265 Some customers have undertaken more complex modernisation efforts, some of which include re-writing applications to run with alternatives to the Microsoft

²¹⁶⁸ We held calls with 12 customers which responded to our initial requests for information that they used Windows Server and SQL Server on the public cloud.

²¹⁶⁹ Notes of meetings with [REDACTED].

²¹⁷⁰ Notes of meetings with [REDACTED].

²¹⁷¹ We understand for customers adopting PaaS solutions in the cloud, this may not be switching away from the use of Microsoft software. Notes of meetings with [REDACTED].

²¹⁷² Note of meeting with [REDACTED].

²¹⁷³ Notes of meetings with [REDACTED].

software applications.²¹⁷⁴ For example, one customer said it has modernised in a number of ways, including rewriting in-house applications to run on Linux, containerising, moving to SaaS or PaaS. It now provides a hybrid style technology stack that is based on containerisation paths, public cloud and on-premises capability as well. It said that the most significant chunk of its modernisation is containerisation.²¹⁷⁵

7.266 Some customers who had undertaken modernisation were able to estimate the time or effort it took to modernise applications, for example:²¹⁷⁶

- (a) One customer said that it had a team of 8-16 people dedicated to database migration. It said simple migrations could happen as quickly as a few weeks and at the other extreme more complex apps would take many months.²¹⁷⁷
- (b) One customer stated that its containerisation projects take three to six months or up to three to four years depending on complexity.²¹⁷⁸

7.267 Other customers explained that even with modernisation there would remain some dependency on Windows Server.²¹⁷⁹ For example, two customers said that they require Active Directory which needs to run on Windows Server OS.²¹⁸⁰

7.268 Some customers said that one factor preventing modernisation of certain applications was the application requirements and support of commercial off the shelf software.²¹⁸¹

7.269 Although we did not directly ask customers about the impact of AI on modernisation, generally customers did not mention AI as a driver for modernisation or as aiding modernisation efforts. One customer suggested that AI could change the way it views its IT estate and that it is mindful AI may change everything.²¹⁸² It is unclear whether this customer was referring to AI as a motivation for migration as discussed below, or as aiding modernisation of applications.

7.270 In the Jigsaw report,²¹⁸³ some customers mentioned technical challenges associated with re-writing software applications²¹⁸⁴ which can be a necessary step when modernising (as discussed above). Some of these challenges related to re-

²¹⁷⁴ Notes of meetings with [REDACTED].

²¹⁷⁵ Note of meeting with [REDACTED].

²¹⁷⁶ Notes of meetings with [REDACTED].

²¹⁷⁷ Note of meeting with [REDACTED].

²¹⁷⁸ Note of meeting with [REDACTED].

²¹⁷⁹ Notes of meetings with [REDACTED].

²¹⁸⁰ Notes of meetings with [REDACTED].

²¹⁸¹ Notes of meetings with [REDACTED].

²¹⁸² Note of meeting with [REDACTED].

²¹⁸³ [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024).

²¹⁸⁴ [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), pages 46, 47, 61.

writing and modernising applications written for .NET, which is closely associated with Microsoft technology.

Our assessment

- 7.271 Evidence from AWS and Google shows that, while they are attempting to encourage customers to switch away from Microsoft software, their modernisation programmes have had limited impact to date. Google highlights in particular the time and cost of modernising to switch away from Microsoft software. Evidence from providers shows that AI services so far have a limited effect on making modernisation away from Microsoft software significantly easier.
- 7.272 While evidence from customers suggests that some have been able to ‘modernise’ their architecture to switch away from Microsoft software, some customers outlined that they were switching to PaaS or SaaS solutions (some of which may not be directly switching away from Microsoft software), and others mentioned that they would not be able to modernise applications in a way that would allow them to cease their use of Microsoft software. Customers did not proactively mention AI as a driver for modernisation or as aiding modernisation efforts. Taken in the round, the customer evidence shows that modernisation to move away from Microsoft software is not widespread amongst customers.
- 7.273 Overall, the evidence suggests that while some customers are modernising workloads, modernisation attempts by AWS and Google have had a limited effect on moving customers away from Microsoft software products.

Conclusions on Microsoft’s market power in related software markets

- 7.274 Microsoft has a significant degree of market power in relation to Windows Server, SQL Server, Windows 10/11, Visual Studio and its productivity suites. In particular, Microsoft has moderate to high market shares in all of these products, all of which have a degree of differentiation in terms of characteristics and features. Customer evidence generally suggests that customers would be unlikely to switch to alternatives in response to a 5% price rise. Our findings about Microsoft’s market power in these products is also generally supported by provider evidence. This is consistent with Microsoft having the ability and incentive to foreclose rivals.
- 7.275 In addition, AWS’ and Google’s efforts to encourage customers to modernise away from Microsoft software products have had limited success.
- 7.276 These conclusions are robust to variations of the precise product market definition used, meaning that they would not change based on the precise boundaries of the relevant product markets. If the relevant markets were wider, the additional constraints that would then fall within the market would be relatively weak compared to those we have assessed as in-market constraints. If the relevant

markets were narrower, the constraints that would then fall outside the market would be relatively strong out-of-market constraints (albeit weaker than the in-market constraints). As a result, any difference in how we define the relevant product markets would not affect our decision.

The importance of Microsoft software inputs to AWS and Google

7.277 In this section, we assess the importance of the Microsoft software inputs in shaping both the prices and quality of the downstream cloud offerings of AWS and Google. We structure our assessment as follows:

- (a) First, we consider the importance of the cloud services that make use of each software product relative to overall cloud spend. To inform our assessment we have calculated the proportion of total revenues that spend on Microsoft software workloads accounts for on Azure.
- (b) Second, we consider the significance of the relevant Microsoft software inputs by reference to the proportion of AWS' and Google's cost base that they account for. This section focuses on Windows Server and SQL Server as these products are available via AWS' and Google's respective SPLAs. In particular, we consider the following evidence:
 - (i) Analysis submitted by Microsoft which calculates AWS' and Google's actual software licensing spend as a proportion of their cloud revenues from customers that use either Windows Server or SQL Server;
 - (ii) Our own analysis which estimates what Microsoft's software licensing spend would be as a proportion of its cloud revenues if Microsoft were subject to the same licensing terms as AWS and Google; and
 - (iii) Analysis submitted by Microsoft which estimates the proportion of total cloud spend accounted for by spend on Windows Server VMs.
- (c) Third, we consider other factors that may affect the significance of Microsoft software inputs to AWS and Google regardless of the proportion of the cost base they account for. This section focuses on Windows Server, SQL Server, Windows 10/11, Microsoft's productivity suites, and Visual Studio as these are subject to non-price related licensing practices. In particular, we consider:
 - (i) the importance of other non-price factors including Extended Security Updates (ESUs);

- (ii) the importance of the Microsoft software products to VDI and the relative size of cloud-based VDI and future growth.²¹⁸⁵
- (d) Finally, we consider evidence on whether the importance of Microsoft's software inputs will change in the foreseeable future.

Distribution of cloud use by spend

7.278 We have assessed the importance of the cloud services that make use of each software product relative to overall cloud spend. We have calculated the proportion of total revenues that spend on Microsoft software workloads accounts for. The greater the proportion of cloud spend accounted for by Microsoft software products, the greater scope there is for any price or non-price differences in the Microsoft input to impact downstream competition.

7.279 We used Azure UK revenue data for 2024. To the extent that Microsoft's licensing practices cause customers to deploy their Microsoft software workloads on Azure disproportionately, this analysis overstates the importance of Microsoft software workloads relative to overall cloud revenues.

7.280 Figure 7.1 below shows the distribution of cloud use on Azure by spend on each service, and the further distribution of spend on particular services by the spend to usage that includes each of the Microsoft software products.

Figure 7.1: [REDACTED]

[REDACTED]

Source: CMA analysis of Microsoft's response to the CMA's information request [REDACTED].

7.281 Our analysis shows that:

- (a) Virtual Machines is [REDACTED] cloud service on Azure, accounting for [20-30]% of total UK revenues. Of that, spend on Windows Server and Windows 10/11 VMs account for [40-50]% and [10-20]% respectively.
- (b) Storage is the [REDACTED] cloud service, accounting for [20-30]% of total Azure spend.
- (c) SQL Database is the [REDACTED] cloud service, accounting for [5-10]% of total Azure spend. SQL Database is Azure's primary SQL Server PaaS offering.
- (d) Virtual Machines Licences includes sales of all the additional software licences (beyond the OS) that customers use with their VMs. These revenues

²¹⁸⁵ As we outline above, VDI is what allows customers to use Windows 10/11, Visual Studio and Office/Microsoft 365 in the public cloud.

account for [0-5]% of total Azure revenues, of which SQL Server licences account for [60-70]%.

- (e) Microsoft Defender, Power BI and Entra ID are constituent parts of certain Microsoft 365 suites. Together, these account for [0-5]% of spend on Azure.
- (f) Azure Virtual Desktop is Microsoft's primary first-party VDI offering. [REDACTED].²¹⁸⁶ Together, therefore, VDI revenues account for [0-5]% of total Azure revenues.
- (g) All other services ([REDACTED]) account for less-than [0-5]% of Azure UK revenues each.

7.282 Overall, Microsoft software workloads account for [20-30]% of total Azure revenues. Windows Server Virtual Machines and SQL Database account for [10-20]% and [5-10]% respectively. VDI accounts for [0-5]% which makes it [REDACTED].

7.283 Although this analysis overstates the importance of Microsoft software workloads relative to overall cloud revenues, we have considered this analysis in the round with the other analysis and evidence presented to assess the importance of the Microsoft software products.

Significance in the cost base of AWS and Google

Introduction

7.284 To assess whether Microsoft's licensing practices are likely having the effect of reducing competition in cloud services, this part of our assessment of the importance of the inputs focuses on the significance of the input in the cost base of AWS and Google. The greater the proportion of total costs that a particular input accounts for, the more likely it is that a higher price of that input will lead to a higher price for the downstream product.

7.285 For example, if an input accounts for 50% of the cost base, then (in a case where there is 100% cost pass-through) a 10% increase in the price of the input will translate into a 5% increase in the price of the downstream product. For comparison, if an input accounted for 10% of the cost base, the same 10% increase in the price of the input would translate into a 1% price increase downstream.

7.286 The licensing practices we have considered relate to the input prices that Microsoft charges to AWS and Google via their respective SPLA contracts for the right to resell Windows Server and SQL Server as part of their own cloud solutions.

²¹⁸⁶ Note of meeting with [REDACTED].

Accordingly, this section considers various measures of the significance of the Microsoft software IP input costs.

Aggregation of workloads

- 7.287 Before setting out evidence on the indicators of Microsoft software's significance in the cost base, it is necessary to consider an overarching issue relating to how competition works in cloud services, which affects the analysis that follows.
- 7.288 In some industries, when calculating the input cost as a proportion of the total costs of supplying customers, we would use the unit cost of the downstream product as the comparator.²¹⁸⁷ However, cloud providers do not compete to supply a single downstream product with a single unit cost. Rather, they compete to supply bundles of services that vary in composition and value. Consequently, when considering the significance of the cost of Microsoft software relative to the cost base, it is necessary to consider what set of services form part of the relevant cost base.
- 7.289 To illustrate why the relevant cost base matters, it is useful to consider some possibilities. One possibility would be to compare spend on the Microsoft software to the individual product or service that uses the software as an input (for example comparing the spend on the Windows Server software to the cost to provide a Windows Server virtual machine (VM)). Another possibility would be to compare the Microsoft software costs not only to the spend on the service that uses the software, but also any other services used to achieve the customer's objective, eg running a particular workload on the cloud. Another possibility would be to compare it to groups of workloads that are purchased together (which could be considered a project or a deal) and another would be simply to compare it to the total of customer's spend on cloud services.
- 7.290 The broader the set of services used in the comparator, the less significant the spend on Microsoft software will tend to appear in relative terms.
- 7.291 In our PDR, we compared the licensing input costs to (i) customer spend on Windows Server VMs (and the SQL Server licensing IP where appropriate) and (ii) total customer spend on Azure services, which represent narrow and broad purchasing decisions respectively, and so provided an upper and lower bound estimate of the significance of the Microsoft software licensing cost to rivals.
- 7.292 We set out in the PDR that neither of these upper and lower bound estimates are likely to represent the significance of the licensing input costs in competing for any

²¹⁸⁷ For example, one cloud provider [REDACTED] submitted that the Windows Server licence typically comprises around [40-50]% of the on-demand list price for [REDACTED] that include Windows Server. If the downstream market was simply for Windows Server VMs, then this would suggest that the Microsoft software input forms a significant part of the cost base. [REDACTED] response to Ofcom's information request [REDACTED].

single customer. That is, unless customers make fully aggregated decisions (ie, one decision that covers all of their cloud services), comparing the licensing input costs to total customer spend would underestimate the importance of the software. Similarly, unless customers make fully disaggregated decisions (ie, separate decisions about each individual service or product such as a Windows server VM), comparing the licensing spend to customer spend on Windows Server VMs will overstate the importance of the software.

7.293 As such, we considered that values within the range were likely to be more reliable as indicators of the importance of the Microsoft software as an input. Following the PDR, we collected evidence that has allowed us to consider where customer behaviour is likely to sit within this range (ie the level of aggregation of workloads by customers).

7.294 In light of the above, we have structured this section as follows:

- (a) First, we set out submissions and evidence relating to the extent to which customers aggregate the purchases of cloud services into a single decision (or a smaller set of decisions that strongly influences each other). On the basis of this evidence, we set out our assessment of the appropriate set of services against which to compare the significance of the Microsoft software.
- (b) Second, we set out a range of indicators of the importance of Microsoft software inputs to the cost base, including indicators submitted by Microsoft as well as some of the results of our own analysis.²¹⁸⁸ This assessment incorporates our conclusions on customers' aggregation of their purchasing decisions relating to cloud services.

Aggregation of customers' cloud spend

7.295 In this section, we consider evidence on how customers approach purchasing cloud services and the extent to which they tend to aggregate their purchases of cloud services into a single decision (or a set of decisions that strongly influence each other) versus the extent to which they disaggregate, or consider disaggregating, their purchases.

7.296 The structure of this subsection is as follows:

- (a) We consider overarching party submissions on aggregation and the role of Committed Spend Agreements (CSAs) in the level of aggregation.
- (b) We present evidence on the prevalence and coverage of CSAs.

²¹⁸⁸ For the full results, please see appendix T.

- (c) We set out evidence from customer calls.
- (d) We provide an assessment of the cloud provider submissions in light of the evidence presented.
- (e) We consider evidence on the minimum bundle of services that Microsoft software can serve as an input for.
- (f) We set out our overall assessment of aggregation and the relevant comparator.

Cloud providers overarching views on aggregation

- 7.297 Microsoft submitted that the individual workload is the narrowest possible contestable unit of competition, but that most commonly customers make deployment decisions for multiple workloads at one time.²¹⁸⁹ Microsoft said that discounts are given at the deal level, cash is fungible, and customers care about the overall deal price rather than the price of any individual component.²¹⁹⁰
- 7.298 Microsoft submitted that the prevalence of CSAs is higher the more customers spend on cloud and, therefore, larger customers are more likely to aggregate workloads.²¹⁹¹
- 7.299 Microsoft submitted that AWS and Google would have the incentive to encourage customers to structure the aggregation of workloads in ways that dilute the significance of Microsoft IP costs in the overall opportunity.²¹⁹²
- 7.300 Microsoft also submitted that we did not explain why the most sophisticated (that is, highest spending) customers would disaggregate workloads in a way that would make it harder for AWS and Google to compete or that would sacrifice discounts that could be earned from aggregating workloads into a CSA.²¹⁹³
- 7.301 Microsoft submitted that we presented a hypothesis in the PDR that large customers ‘hive off’ workloads composed of one or two Microsoft IP-based IaaS components and place these on Azure without supporting evidence.²¹⁹⁴ It said that this lack of evidence is consistent with its own experience that large customers do not disaggregate workloads in this way.²¹⁹⁵
- 7.302 In relation to this point, Microsoft submitted:

²¹⁸⁹ Microsoft’s response to the PDR, paragraph 162

²¹⁹⁰ Microsoft’s response to the PDR, paragraph 204.

²¹⁹¹ Microsoft’s response to the PDR, paragraph 143.

²¹⁹² Microsoft’s response to the PDR, paragraph 171.

²¹⁹³ Microsoft’s response to the PDR, paragraph 149.

²¹⁹⁴ Microsoft’s response to the PDR, paragraphs 146-148.

²¹⁹⁵ Microsoft’s response to the PDR, paragraph 169.

- (a) Azure does not market itself as a narrow destination for Microsoft workloads rather than a broad destination for the full suite of Azure services (such as the [X] services that customers in the >\$20M bracket consume on average).²¹⁹⁶
- (b) Microsoft submitted analysis to show that multi-cloud customers consume a [X] number of services to customers that do not multi-cloud,²¹⁹⁷ and have a [X] relative spend on Windows Server.²¹⁹⁸

7.303 Microsoft has also submitted that it competes for each workload.²¹⁹⁹

7.304 Microsoft submitted that there is inconsistency between defining a wide relevant market downstream and focusing the analysis on a subset of workloads.²²⁰⁰ Microsoft said that AWS and Google's submissions treat Microsoft workloads as part of the market which they cannot compete effectively for, but this an artificially narrow lens. Microsoft added that IaaS components (Windows Server VMs and/or SQL Server) are not workloads and still less are they a market, let alone the PDR's defined market.²²⁰¹ Microsoft said that the CMA has no basis for finding an AEC based on a different market definition to that specified in the PDR.²²⁰²

7.305 [X] submitted that, as competition takes place at a workload level, [X] should never have the incentive to use the margin of one workload to compensate for the loss made on another workload, and that it would always be better off supplying only the unaffected workloads (ie the workloads that do not include the Microsoft software).²²⁰³ [X] submitted that, even where [X] if a customer runs a significant number of other workloads on [X], which is a major uncertainty.²²⁰⁴ [X] also referred to the average attainable margins for [X] based on an AEC test, and said that because customer needs are highly variable, the low margins observed implies that margins will likely be negative for many workloads and customers and suggests that it is not commercially viable for Microsoft's rivals to compete on a significant share of workloads.²²⁰⁵

7.306 [X] submitted that its CSAs cover [X] submitted that CSAs are entirely voluntary, determined by each customer, and generally driven by each customer's specific business needs, rather than being aimed at aggregating customers'

²¹⁹⁶ [Microsoft's response to the PDR](#), paragraph 169.

²¹⁹⁷ Microsoft's submission to the CMA [X].

²¹⁹⁸ Microsoft's submission to the CMA [X].

²¹⁹⁹ [Microsoft's response to the Technical barriers working paper](#), paragraph 23.

Similarly, Microsoft submitted that competition is characterised by intense and persistent competition to win the next workload, irrespective of whether customers are new or existing to Azure ([Microsoft's response to the Issues Statement dated 09 November 2023](#), paragraph 11).

²²⁰⁰ Microsoft's submission to the CMA [X].

²²⁰¹ Microsoft's submission to the CMA [X].

²²⁰² Microsoft's submission to the CMA [X].

²²⁰³ [X] submission to the CMA [X].

²²⁰⁴ [X] submission to the CMA [X].

²²⁰⁵ [X] submission to the CMA [X].

workloads/spend to increase their overall commitment. [X] said that competing for a larger portion of customer spend (ie what the CMA frames as ‘aggregating workloads’) is the definition of competition, but these CSAs are driven by customers. [X] and customers have complete flexibility to choose how much spend they place with [X].²²⁰⁶

- 7.307 [X] submitted that the discounts offered depend on the commitment made by the [X]. In this sense, [X] said that it cannot effectively minimise the costs of Microsoft’s licensing practices by ‘encouraging the aggregation of workloads’ as different workloads (ie services) will have different margins, and [X] does not dictate the type of workloads that a customer might choose to run on [X], nor [X].²²⁰⁷
- 7.308 [X] submitted that it is usually in its interest that customers use more of [its] services and it therefore encourages aggregation of workloads to an extent. [X] added that it is important to ‘right-size’ a particular deal and ensure any discounts are linked to the customer’s actual existing or forecasted consumption.²²⁰⁸
- 7.309 [X] submitted that it primarily uses CSAs [X] to win incremental non-Windows workloads and in circumstances where customers typically already have a CSA in place with their primary cloud provider, [X].²²⁰⁹
- 7.310 [X] said that CSAs do not help them [X].²²¹⁰
- 7.311 [X] also said that while CSAs can give some indication of expected customer spend, the projections are not set in stone and the customer may renegotiate their agreement. [X] further submitted that customers are willing to disaggregate workloads when it suits their purchasing decisions and can do so easily due to the flexible nature of CSAs.²²¹¹
- 7.312 [X] submitted that the variability in the types of cloud services and the amount of other cloud services (ie those that do not include the Microsoft software) purchased means that it is highly difficult and impracticable to predict consumption, and the profitability of those other sales is therefore uncertain. [X] said that it would not be a rational strategy to use Windows Server VMs as a ‘loss leader’ with the uncertain prospect of earning profits from higher margin cloud services at an unknown time in future.²²¹²
- 7.313 [X] also submitted that this variability in purchasing patterns means that cloud competition is fundamentally different from other markets where there is genuine

²²⁰⁶ [X] response to the CMA’s information request [X].

²²⁰⁷ [X] response to the CMA’s information request [X].

²²⁰⁸ [X] response to the CMA’s information request [X].

²²⁰⁹ [X] response to the CMA’s information request [X].

²²¹⁰ [X] response to the CMA’s information request [X].

²²¹¹ [X] response to the CMA’s information request [X].

²²¹² [X] response to the CMA’s information request [X].

bundle-on-bundle competition. [X] said that such circumstances often involve customers purchasing bundles that comprise a limited number of core components and where those components are generally purchased at the same time or at predictable intervals, which allows firms to compete by offering discounts on particular components or the entire bundle. [X] added that, however, the variability in cloud customer purchasing patterns makes it difficult to predict what other services a customer will purchase when it migrates Windows Server-based workloads to the cloud.²²¹³

- 7.314 One of a cloud provider's [X] internal documents states that Windows has a pull-through effect, with the average Windows customer on its cloud platform bringing in [X] times non-Windows workloads.²²¹⁴

Customer evidence

- 7.315 We asked cloud customers about the extent to which they aggregate their cloud services and to understand their decision-making processes. We also asked them how considerations such as their multi-cloud strategy and CSAs influence them to make aggregated or disaggregated decisions.²²¹⁵
- 7.316 Some of the customers we spoke to made one largely aggregated decision to allocate most or all of their cloud spend to a single cloud provider:
- (a) One customer underwent a large IT transformation project after it was sold by its parent company, which involved migrating all of its workloads to the cloud. It held a formal procurement process with two providers, and negotiated a five-year CSA with the provider it chose based on its total predicted cloud spend over the period.²²¹⁶
 - (b) Another customer took one decision to move all of its workloads to the cloud as quickly as possible, and predominantly uses Azure although a small portion of its stack is on GCP. Although the employees we spoke to were not present at the initial time of migration, they said that as nearly all their workloads are on Azure, their team is experienced in Azure. They are underspending on their CSA with Azure, the only decisions they now make are over which product to use, rather than which cloud.²²¹⁷
 - (c) A further customer said that it migrated from on-premises to the cloud in one go, and that staying with Microsoft was a natural progression. It said that it negotiates CSAs based on its expected spend over the three-year term.

²²¹³ [X] submission to the CMA [X].

²²¹⁴ [X] response to the CMA's information request [X].

²²¹⁵ We had in-depth calls with customers, which involved asking open-ended questions about how they approach their purchasing decisions.

²²¹⁶ Note of meeting with [X].

²²¹⁷ Note of meeting with [X].

However, it said that it was currently overspending against its CSA target and that certain aspects of its demand were harder to predict. For example, it mentioned that some companies that it acquired already had an existing cloud provider but that it migrates their workloads to Azure once acquired. The speed at which new projects are integrated can be unpredictable; however due to regulatory requirements it aims to integrate newly acquired businesses at fast pace. In addition, it recently had some big projects come to fruition that it did not anticipate when it negotiated its CSA.²²¹⁸

- 7.317 One customer migrated workloads to the cloud more incrementally and described it becoming an Azure-only customer as ‘accidental’. It said that it initially chose Azure based on the skills of its IT team and had considered deploying workloads on alternative clouds but decided against doing so because of the CSA discounts that it would forego.²²¹⁹ This customer made incremental but not independent allocation decisions and its CSA encouraged the customer to aggregate workloads.
- 7.318 One customer said that it wanted to migrate within a set timeframe and aimed to avoid lift-and-shift but in some cases this was unavoidable. It said that while it uses AWS and GCP, it mostly uses Azure, which it chose because it considered itself to be a ‘Microsoft shop’ and it was better to use Microsoft technologies given their prior knowledge. It added that it tried to avoid PaaS due to the lock-in associated with it. The customer said that it had a CSA at one point but the target was far too high and it did not meet it, so it no longer has a CSA.²²²⁰ This customer appeared to make a largely aggregated decision during its initial migration, but now uses multiple clouds and actively aims to avoid being locked-in to a single cloud.
- 7.319 Some customers said that they use multiple clouds, but their spend on each cloud was largely dependent on the business area (and therefore aggregated to an extent):
- (a) One customer said that it assessed each application for migration, but within a broader pattern of certain business areas preferring Azure or AWS and typically having a critical mass on each. It added that it would not typically split components of a business stack across clouds if it had the opportunity to keep them together. For example, it said that it makes sense to keep the database and analytics services that support an application on the same cloud as that application. That customer said [REDACTED]. However, it said [REDACTED]. It added that [REDACTED].²²²¹

²²¹⁸ Note of meeting with [REDACTED].

²²¹⁹ Note of meeting with [REDACTED].

²²²⁰ Note of meeting with [REDACTED].

²²²¹ Note of meeting with [REDACTED].

- (b) Another customer adopted a multi-cloud strategy from the outset to avoid lock-in, putting workloads for certain parts of its business on Google and workloads for other parts of its business on AWS. This allowed key applications to be colocated and helped its engineers to focus on specialising on a particular cloud in the early stages of its cloud adoption. However, the customer said that it has now built up its in-house capabilities and has a workload placement strategy that considers things like concentration risk, operational resilience, and cost efficiencies, and is increasingly aiming to create portable solutions that can run across multiple clouds.²²²²

7.320 The customers referred to above disaggregated their total cloud demand by splitting it between clouds; however, their spend on each cloud was allocated in an aggregated way initially. However, both customers suggested that they now make less aggregated decisions, allocating workloads based on the merits of the constituent services.

7.321 Some customers said that they make decisions about where to allocate particular projects or groups of workloads, based on business need or best fit:

- (a) One customer initially moved a portion of its workloads to Azure in one go but has since 2018 migrated workloads incrementally to Azure and other clouds. It said that it now evaluates cloud solutions on a case-by-case basis, according to its enterprise architecture standards. It added that it has a CSA with Azure that covers the workloads it initially migrated, but this would not influence decisions over where to place additional workloads, which it makes based on business need.²²²³
- (b) Another customer migrated one customer-facing product area at a time to AWS over a 7-8 year period. The customer explained that its moves were not part of a coordinated plan; rather it had learned from each project and adjusted its migration plan as it went along. Since then, its decisions have been made on an application-by-application basis, in line with its multi-cloud strategy that aims to keep workloads relatively balanced across AWS and Azure, having negotiated CSAs with both. It added that having relationships with two cloud providers makes it easier to compare pricing and gave it greater flexibility. The customer said that around 80% of its deployments on the cloud were driven by forecasts, and 20% were opportunistic or driven by external factors. In allocating the 20%, it considers whether either provider possesses any unique technologies from which it could benefit.²²²⁴

²²²² Note of meeting with [REDACTED].

²²²³ Note of meeting with [REDACTED].

²²²⁴ Note of meeting with [REDACTED].

- (c) A further customer said that it historically migrated to the cloud on an ad hoc, project-by-project basis, with business areas taking largely decentralised decisions about where to deploy workloads. It said that it has typically migrated existing workloads to Azure and developed new workloads in AWS, but that no decisions are made ahead of time where a workload will be deployed because of the difficulties to determine the appropriate Cloud venue. The customer added that it makes decisions about ‘chunks’ of workloads based on business needs and the benefits of the technology, as well as cost structures.²²²⁵
- (d) Another customer said that it has individually evaluated each of the workloads it has moved to the cloud. It said that it tried a mass migration to one provider at one point, but this was unsuccessful, and even in that case it evaluated each workload in the proposed move. The customer said that it made individual application deals at first but now has CSAs with Azure and AWS. It said that it agreed its spend targets with AWS based on projected growth of existing applications rather than a commitment to move a percentage of its portfolio to AWS. It does not feel compelled to deploy workloads in Azure or AWS to meet spend thresholds as the discounts are similar across each cloud. However, it said that data gravity was a factor in its decision-making, ie workloads that interact with each other would need to be hosted together to reduce latency.²²²⁶
- (e) Another customer said that it adopted cloud incrementally, starting with a couple of pilots. It said that it could not have anticipated all of the workloads that it decided to build in or move to the cloud and has moved to the cloud in chunks as opportunities came up through projects or programmes. It added that its adoption was often triggered by new services and not always part of a coordinated plan. It said that it aims to put workloads where they best fit, independently of the cloud provider, and considers factors such as functionality, risk, people’s skills, and cost.²²²⁷

7.322 These customers all made decisions about subsets of their cloud demand, either at the workload or project level, and have allocated spend based on the merits of the constituent services rather than their choices being determined by where their other workloads are placed, or to draw down their CSA discounts. We therefore consider that these customers make largely (though not completely) disaggregated decisions.

7.323 Overall, this customer evidence provides an insight into the varying factors that drive customers’ decisions about how and where they allocated cloud spend.

²²²⁵ Note of meeting with [REDACTED].

²²²⁶ Note of meeting with [REDACTED].

²²²⁷ Note of meeting with [REDACTED].

CSA coverage

- 7.324 Below we set out some evidence relating to the coverage of CSAs across customers in the UK. This provides evidence on the level of demand for cloud services that is not impacted by the presence of a CSA.
- 7.325 Table 7.3 below shows the share of Azure UK customers that had active CSAs in 2024 in each revenue bracket and the share of revenue in that bracket that they account for, as submitted by Microsoft.

Table 7.3: Share of Azure UK customers with CSAs active in 2024 by revenue bracket

Revenue bracket (\$)	Share of customers with a CSA (%)	Revenue share for customers with a CSA
10k-1M	[5-10]	[5-10]
1M-5M	[20-30]	[20-30]
5M-10M	[60-70]	[60-70]
10M-20M	[80-90]	[80-90]
>20M	[90-100]	[90-100]

Source: Microsoft's submission to the CMA [REDACTED]

- 7.326 The data shows that the prevalence of CSAs among Azure UK customers increases with customer spend and that the vast majority of customers in the top two revenue brackets have CSAs.
- 7.327 The prevalence of CSAs is low among customers in the bottom two revenue brackets, which together account for [40-50%] of Azure revenue.²²²⁸ This shows that a large proportion of customers (weighted by spend) are not incentivised by CSAs to aggregate workloads.
- 7.328 As set out in the CSAs chapter, roughly [REDACTED]% of Azure and AWS' revenues come from customers that have CSAs. Additional analysis submitted by [REDACTED] suggests that the spend commitments of [REDACTED] customers for [REDACTED] CSAs covered around a [REDACTED] of [REDACTED] revenues in 2022.²²²⁹ Taken together, these figures suggest that while most large customers have CSAs, a substantial proportion of their spend is not covered by CSAs. This could be explained by customers seeing some benefit in the option to potentially disaggregate their workloads. It would be consistent with evidence from our customer calls (see above) that some customers have CSAs that cover a discrete set of workloads rather than their total spend or total anticipated spend.

²²²⁸ CMA analysis of Microsoft's submission to the CMA [REDACTED].

²²²⁹ [REDACTED] estimates \$[REDACTED] of spend on [REDACTED] cloud infrastructure services, of which \$[REDACTED] of spend on [REDACTED] services under CSDs. Therefore $\$[REDACTED] / \$[REDACTED] = [REDACTED]\%$. [REDACTED]. As explained in the CSA working paper, we consider that this analysis may underestimate the coverage of CSAs on [REDACTED] because the CMA understands that where a CSA is not held by the UK part of a business that has a non-UK parent entity, the UK revenues may be counted in [REDACTED] total UK revenues (the denominator) but excluded from the UK revenues that are subject to CSAs (the numerator).

Our assessment

- 7.329 At either extreme, customers could aggregate their entire demand for cloud services into a single decision or make separate decisions about each service; the evidence we have set out above speaks to how near or far we are from either extreme. We set out our views on this evidence below.
- 7.330 Microsoft's submissions in relation to large customers and multi-cloud to an extent mischaracterised the reasoning we set out in our PDR.
- 7.331 In the context of the aggregation of customers' purchasing of cloud services, we made two observations in relation to customers that multi-cloud to which Microsoft's submissions relate:
- (a) First, we observed that large customers are more likely to multi-cloud. This observation is supported by a range of evidence, including the evidence on the prevalence of multi-cloud across customer groups in Competitive Landscape, section 'prevalence of multi-cloud'.
 - (b) Second, we observed that multi-cloud is consistent with a willingness and ability to disaggregate workloads. The fact that a customer uses multiple clouds shows that they have considered splitting their demand across clouds and, indeed, have already done so. We consider that this is a relevant consideration that goes against treating competition for multi-cloud customers as generally being for all of their cloud spend at once.
- 7.332 Importantly, we did not suggest that multi-cloud customers would necessarily consume fewer services on Azure or have a higher relative spend on workloads that rely on Microsoft software. Rather, that multi-cloud customers are more likely to make disaggregated decisions such that it is reasonable to compare the licensing input costs to a subset of their spend.
- 7.333 We agree with Microsoft that providers have a clear incentive to encourage customers to allocate more of their demand to them, and that the more customers aggregate demand, the easier it is for AWS and Google to dilute the significance of the Microsoft software input.
- 7.334 However, as evidence from customers shows, customers also have other incentives when deploying workloads to the cloud: they may weigh the discounts on offer for aggregating workloads against the flexibility to choose the best-in-class services for each workload or group of workloads, or to avoid lock-in.
- 7.335 The evidence shows that the extent to which customers aggregate their workloads varies. While some customers made largely aggregated decisions, others made more disaggregated decisions, either by splitting their demand across multiple

clouds, or making separate, independent decisions about groups of workloads or on a project-by-project basis.

- 7.336 We considered a range of evidence on the extent to which CSAs encourage aggregation, as well as the prevalence and coverage of CSAs. This evidence suggests that while CSAs do encourage customers to aggregate their workloads, the extent of this effect varies significantly by customer. In particular, our evidence shows that:
- (a) CSAs encourage some customers to aggregate all or most of their workloads.
 - (b) Customers that have CSAs with multiple cloud providers may benefit from similar discounts on each cloud, allowing them to make more independent decisions about particular workloads.
 - (c) Some customers appear to have CSAs based on growth in demand for a set of existing workloads, giving them freedom to make disaggregated decisions about additional workloads.
 - (d) Some customers appear to form coordinated migration plans and negotiate CSAs only once they have matured in the cloud, suggesting that their earlier allocation decisions are more disaggregated.
 - (e) Some customers have CSAs which cover all or most of their spend, while others have CSAs which cover only a subset of their spend, giving them freedom to make disaggregated decisions about spend not covered by the CSA.
- 7.337 This is consistent with evidence set out elsewhere in this report: in particular, most large customers we spoke to said that CSAs were either somewhat likely or very likely to impact which cloud they allocate additional workloads to.
- 7.338 Additionally, while the evidence on the prevalence of CSAs shows that most larger customers have CSAs, a substantial portion of their spend is not covered by CSAs, and smaller customers that do not have CSAs account for a substantial portion of revenue. This is consistent with our understanding that most larger customers can qualify for CSAs that cover just a portion of their spend, and our customer evidence shows that some customers have CSAs that cover a discrete set of workloads rather than their total actual or anticipated spend.

Conclusions on aggregation of workloads

- 7.339 The extent to which customers make aggregated decisions varies significantly. Given this variation, it is not possible, nor desirable to identify a specific set of services that comprises the relevant cost base for an assessment of the

importance of the input. As such, it is appropriate to consider values within the range, as these are likely to be more reliable as indicators of the importance of the Microsoft software as an input. Most customers are not at either extreme (that is, the relevant unit of their demand is not the costs of supplying all of the services they consume on Azure, nor is it individual products or services such as a Windows Server VM). Therefore, we have placed greater weight on values towards the middle of the range for the cost base.

- 7.340 As noted above, there is no bright-line division between those customers with Microsoft workloads that AWS and Google can compete effectively for and those for which they cannot. Whilst there are some customers with Microsoft workloads that AWS and Google can compete effectively for, there is a significant group of customers for which AWS and Google cannot compete effectively. Our guidelines recognise that one set of customers may be more affected than others by a particular feature and therefore, contrary to Microsoft's submission set out above, it is not correct to say that we would not be able to find that Microsoft's licensing practices are having an adverse effect on competition in the cloud services markets. We consider this group of customers to be significant such that the licensing practices could have an impact on competition in the cloud services markets.²²³⁰

The minimum bundle

- 7.341 We have also received submissions from Microsoft and [X] about the 'minimum bundle' – ie the narrowest set of services that Microsoft's software can serve as an input for.
- 7.342 Microsoft submitted that the narrowest unit of demand, and therefore competition, is a workload. It submitted that a Windows Server VM is not a workload and does not become a workload by adding a SQL Server component.²²³¹ Rather, it submitted that at least some storage and networking are required to make a Windows Server VM functional.²²³² On that basis, Microsoft suggested that the minimum denominator that we should use to calculate our estimates of the importance of the input should include spend on, at least, Windows Server VMs, Storage, and Virtual Network.
- 7.343 In support, Microsoft submitted evidence that [X] of UK Windows Server customers (with spend above [X]) in 2023-24 also use [X].²²³³
- 7.344 Microsoft also submitted that a minimum bundle consisting of Windows Server VM, Storage and Virtual Network is itself unrealistically narrow, and inconsistent

²²³⁰ CC3 (Revised), paragraph 150.

²²³¹ [Microsoft's response to the PDR](#), paragraphs 145.

²²³² [Microsoft's response to the PDR](#), paragraph 163.

²²³³ Microsoft's submission to the CMA [X].

with the business reality of how customers consume cloud services.²²³⁴ We understand Microsoft's submission to mean that customers do not make separate decisions about where to place this set of services, such that spend on these services constitutes an unreasonably narrow base against which to compare the Microsoft software input costs.

7.345 In support of its submission that the relevant bundle of services against which to compare Microsoft software input costs should be broader than these three services, Microsoft submitted the following observations:

- (a) Azure customers that used Windows Server in 2024 also used [X] other services on average,²²³⁵ and larger customers consumed [X] eg customers in the >\$20M bracket consumed [X] services on average,²²³⁶
- (b) Only [X] Azure customers that used Windows Server and spent [X] consumed only [X] in 2024;²²³⁷

7.346 [X] of Azure customers that used Windows Server and spent [X] in 2024 also consumed [X].²²³⁸

7.347 [X] submitted that Microsoft's analysis is limited to showing that customers that purchase Windows Server VMs often also purchase an amorphous number of other services, but does not show that:

- (a) There is a significant specific overlap in the additional services being purchased other than the necessary storage and networking used with the VM; or that
- (b) Microsoft actually competes [X] on such a specific set of cloud services.

7.348 [X] further submitted that there is no standard bundle or set of cloud services that customers purchase from a single cloud provider, but rather, there is significant variability in customer purchasing patterns in terms of the types, combinations and quantities of cloud services that customers purchase, and the timing of those purchases.

7.349 [X] submitted that this variability supports the CMA's use of a narrow set of cloud services (specifically, those that are necessary to run and used in conjunction with Windows Server VMs) as it undermines Microsoft's position that there is a broader

²²³⁴ Microsoft's submission to the CMA [X].

²²³⁵ Microsoft's submission to the CMA [X].

²²³⁶ Microsoft's submission to the CMA [X].

²²³⁷ Microsoft's submission to the CMA [X].

²²³⁸ Microsoft's submission to the CMA [X].

standard set of cloud services that should be used to compare the licensing input costs to.²²³⁹

7.350 To support this position, [X] submitted the following evidence:²²⁴⁰

(a) [X]²²⁴¹

(b) [X]

(c) [X]²²⁴²

(d) [X]

7.351 [X] said that the variability in purchasing patterns and the fact that Windows Server VMs account for a substantial portion of Windows Server-based workloads dismisses Microsoft's claim that [X] would be able to cross-subsidise the losses incurred from sales of Windows Server VMs with potential profits from sales on other products.

7.352 In response to [X] position on the variability of purchasing patterns, Microsoft submitted that it does not matter that a cloud provider does not know ex ante which precise array of services a customer will buy when running a workload that includes Microsoft software. Microsoft said that the only thing the cloud provider would know with any certainty was the expected profitability of different customer segments on average and it can decide how much to prioritise each of the segments.²²⁴³

Our assessment

7.353 We agree with Microsoft that it is reasonable to include at least some allocation of customer spend on storage and networking within the comparator. This means that the minimum appropriate level of spend against which to compare the Microsoft software input costs is broader than spend on Windows Server VMs only and, therefore, our estimates using the minimum bundle as the comparator overstate the significance of the input to some extent.

7.354 We also considered Microsoft's analysis of the number of other services that Windows Server customers purchase on Azure. In our view, evidence showing that customers use a number of other services does not necessarily imply that decisions about those other services were influenced by decisions concerning Windows Server VMs, Storage and Virtual Network, or vice versa. In particular,

²²³⁹ [X] submission to the CMA [X].

²²⁴⁰ [X] submission to the CMA [X].

²²⁴¹ For example, [X].

²²⁴² For example, [X].

²²⁴³ Microsoft's submission to the CMA [X].

Microsoft does not show that any of these other services are dependent on or otherwise causally connected to usage of Windows Server VMs, storage and networking.

- 7.355 In addition, even if we did consider this broader set of services as a plausible minimum bundle, customer spend on these services is significantly lower than spend on VMs and Storage, besides SQL Database which also uses Microsoft software as an input. For example, as set out in above section on Distribution of cloud use by spend, while VMs and Storage account for [20-30]% and [20-30]% of Azure public cloud revenues respectively, all the other services listed above account for [10-20]% of revenues combined.²²⁴⁴ Therefore, we consider that including these services would not significantly impact the results of our analysis or alter our conclusions.
- 7.356 We have also considered [REDACTED] evidence regarding the variability in customer purchasing decisions. We agree that this evidence indicates there is no standard bundle of services that Windows Server customers typically purchase together beyond those strictly necessary to operate a Windows Server VM. This variability makes it more difficult for providers to compete through bundled offerings, as providers will not know in any single case how profitable a customer will turn out to be over time.
- 7.357 We acknowledge that providers may compete based on the statistical expectation that a Windows Server customer will also procure additional services. As such, it may be appropriate to compare Microsoft's software inputs to spend on a broader set of services than just those necessary to run a Windows Server VM.
- 7.358 However, the unpredictability of future purchasing decisions means that the expected value of customers' future spend may be lower at the point of competition than the total annual spend that we observe in the data in many cases.
- 7.359 In addition, the incentive to compete based on expected value depends on risk appetite. If a firm is risk averse, then it will place less value on uncertain future returns. In such cases it may be not be appropriate to compare the software inputs to customers' total annual spend in a particular year.
- 7.360 Overall, we consider that Storage and Virtual Network should be included in the minimum set of services with which to compare the Microsoft software input costs. However, we consider that the evidence shows there is no standard bundle of services that the Microsoft software serves as an input for. As such, we consider that it is informative to use the minimum bundle as the lowest comparator.

²²⁴⁴ CMA analysis of Microsoft's response to the CMA's information request [REDACTED].

Indicators of significance in the cost base

Microsoft's analysis of AWS' and Google's licensing costs

- 7.361 Microsoft submitted that Windows Server and SQL Server are the only software services potentially affecting downstream competition.²²⁴⁵
- 7.362 Microsoft submitted an analysis comparing AWS' and Google's global spend on Windows Server and SQL Server licences (via their SPLA contracts) to AWS' and Google's global cloud revenues from customers that use either Windows Server or SQL Server. The analysis covered 2021 to 2023.²²⁴⁶
- 7.363 The results of Microsoft's analysis were as follows:
- (a) AWS' 2023 SPLA licensing input costs for Windows Server and SQL Server accounted for [5-10]% and [10-20]% of revenues generated by customers that used each of the products, respectively; and
 - (b) Google's 2023 SPLA licensing costs for Windows Server and SQL Server accounted for [0-5]% and [0-5]% of GCP revenues from customers that used each of the products, respectively.
- 7.364 Microsoft submitted that input costs as a proportion of revenue within this range are 'unlikely to generate foreclosure because the input cost share is too small to move the needle on downstream customer pricing and therefore competitive strength'.²²⁴⁷
- 7.365 We note that the proportion of downstream costs or revenues that the upstream input must account for to be deemed important is context specific and there is no single cost proportion threshold that an input must meet.
- 7.366 Microsoft's analysis is subject to several analytical issues which lead to a risk of the analysis understating the significance of the relevant Microsoft software as an input for AWS and Google in competing for cloud customers. We set these issues out below.

²²⁴⁵ Microsoft's submission to the CMA [§§].

²²⁴⁶ [Microsoft's response to the Licensing working paper](#) dated 06 June 2024, paragraph 1.5. Note, these results assume that 100% of SQL Server usage relates to Enterprise Edition. See appendix T for a discussion of the SQL Server Editions and their role in determining estimates of the significance of the input.

²²⁴⁷ Microsoft's response to the Licensing working paper dated 06 June 2024, paragraph 6.10. Microsoft referred to the CMA's Private Motor Insurance MIR (Private Motor Insurance Market Investigation ('PMI'), Provisional Findings Appendices, appendix 9.2, para 11, 13.1-15) and the CMA's BT/EE and Virgin/O2 merger inquiries in support of its submission. The analyses carried out in these cases were specific to the markets assessed. We also note that in relation to the PMI, the CMA could not rule out a concern related to an input that accounted for 20% of the downstream average repair bill, and that in relation to the two merger inquiries referenced, the input cost share was just one factor in the CMA's assessment of the ability to foreclose rivals.

- 7.367 Our own analysis (set out further below) addresses three of these issues and concludes that the proportion of revenues that the licensing input accounts for is higher than that submitted by Microsoft.
- 7.368 First, Microsoft's analysis includes spend on all cloud services in the denominator. This means that it effectively compares AWS' and Google's licensing costs with everything their customers spend on cloud services. As shown above, while a measure of cost that incorporates all cloud services might be appropriate if customers made a single aggregated decision about all of their cloud spend, our evidence shows that the extent to which customers make aggregated decisions varies significantly and so it is not appropriate to only consider a comparison against customers' total spend with a particular cloud provider.
- 7.369 The evidence presented above shows that some customers make more disaggregated decisions. By taking the widest possible denominator that includes all services, Microsoft's analysis risks understating the significance of licensing spend in the cost base.
- 7.370 Second, Microsoft's analysis expresses AWS' and Google's licensing spend as a proportion of their respective revenues rather than as a proportion of their costs. We consider the latter would be more appropriate: using revenues as the denominator incorporates a margin which inflates the denominator and therefore risks understating the significance of the licensing spend to the cost base. We consider Microsoft's submissions on the margins of AWS and Google separately in this chapter.
- 7.371 Third, we are cautious about evidence based on current levels of usage of Microsoft software on AWS and GCP, because those usage levels may be affected by any foreclosure effect. For illustrative purposes, it is useful to consider the example of an upstream supplier that refuses to supply an important input to its downstream rival and where the rival loses a significant proportion of its business as a result. In that scenario, the input would represent zero percent of the downstream rival's cost base, despite it being an important input. Analysis based on current usage levels on AWS and GCP should be interpreted in light of the risk that they could understate the importance of the licensing input.
- 7.372 Fourth, Microsoft's analysis does not acknowledge that many customers use both Windows Server and SQL Server:
- (a) Out of all Azure customers that use Windows Server, [10-20]% also use SQL Server on a PAYG basis.

- (b) Out of all Azure customers that use SQL Server on a PAYG basis, [90-100]% also use Windows Server.²²⁴⁸
- (c) Customers that use both Windows Server and SQL Server PAYG account for [70-80]% of total Azure revenues.

7.373 Therefore, for many Windows Server customers, and almost all SQL Server PAYG customers, it is appropriate to consider the significance of the combined cost of licensing both products.

7.374 Fifth, Microsoft's analysis does not consider that the importance of some of the Microsoft software may increase in the future as customers continue to migrate on-premises workloads to the cloud, and face barriers to switching away from Microsoft products when doing so. We have considered the factors that may affect the importance of the input in the future below.

Our analysis of SPLA costs

7.375 We undertook our own analysis of the significance of Windows Server and SQL Server licensing input costs by estimating the proportion of Azure's cost base that Microsoft's licensing spend would account for if it paid the same input prices that it charges AWS or Google for Windows Server and SQL Server.

7.376 We consider this analysis helps to account for some of the issues relating to Microsoft's analysis above.

7.377 First, our analysis is not restricted to looking at licensing costs as a proportion of spend on all cloud services. As set out above, customers vary in terms of how aggregated their workload allocation decisions are. To account for this variation, we have compared the Windows Server licensing input costs to:

- (a) total customer spend on Azure services, and
- (b) customer spend on Windows Server VMs on Azure.

7.378 We have compared the SQL Server and combined Windows Server and SQL Server licensing input costs to:

- (a) total customer spend on Azure services, and
- (b) customer spend on Windows Server VMs and SQL Server licensing IP on Azure.

²²⁴⁸ Our analysis focuses on SQL Server PAYG usage as customers can BYOL to AWS and Google (ie BYOL usage is not subject to SPLA terms).

- 7.379 In each case, (a) (a broad purchasing decision) and (b) (a narrow purchasing decision) provide a lower and upper bound estimate of the significance of Microsoft's software inputs in competing for cloud customers purchasing different bundles of services.
- 7.380 As set out above, our evidence shows that for most customers the relevant comparator is narrower than customers' total spend on cloud services but wider than spend on an individual product or service, for example, a Windows Server VM. Therefore, we place more weight on estimates within the range that we present for each revenue bracket.
- 7.381 In addition, as also noted above, some storage and networking are necessary to make a Windows Server VM functional. Therefore, the narrowest set of downstream products that Windows Server can serve as an input for is a Windows Server VM, storage and networking. The minimum set of downstream services that SQL Server can serve as an input for is a Windows Server VM with SQL Server running on top, storage and networking. As such, using the minimum set of services as the comparator overstates the maximum importance of the Microsoft software input to some extent.
- 7.382 Second, to mitigate the risk that indicators based on current levels of usage on AWS and GCP may understate the significance of the licensing input, as these usage levels may be lower as a result of the licensing practices, our analysis uses Microsoft data on usage on Azure. In our analysis, we calculated how much it would cost Microsoft to host each Azure customer's Windows Server and SQL Server usage if it had to pay the same input prices that it charges to AWS and Google. We then compared these hypothetical input costs to the spend denominators described above.
- 7.383 If Microsoft were subject to the same licensing terms as AWS and Google, then we would expect usage levels on Azure to be lower than they are currently. Therefore, while Microsoft's analysis likely understates the significance of the input, our analysis likely overstates the significance to some extent, and we take this into account in our assessment. However, our estimates using the minimum bundle as the comparator are not impacted by this limitation, as both the numerator and denominator are functions of customers' usage of each software product on Azure rather than their relative usage of Microsoft workloads compared with other workloads.
- 7.384 Third, to account for the fact that [90-100]% of SQL Server PAYG customers also use Windows Server, we present the combined licensing input costs as a proportion of spend for customers that use both SQL Server PAYG and Windows Server (in addition to presenting results for each product separately).

- 7.385 Our analysis does not overcome the issue of using revenues rather than costs as the denominator as we do not have access to data on Microsoft's cost base with respect to each customer. It also does not address the potential for the importance of Microsoft software workloads to increase in the future. As such, our analysis may still understate the significance of the licensing input in the cost base because of either of these factors.
- 7.386 Additionally, our analysis uses data on usage of Windows Server and SQL Server on VMs (ie IaaS usage) only. However, SQL PaaS products such as Azure SQL Database are among the most popular cloud services: as set out above, SQL Database accounted for [0-10]% of total Azure revenues in 2024. This suggests that our analysis significantly understates the significance of SQL Server relative to customer spend across all Azure services (including spend on PaaS services).
- 7.387 As outlined above, we estimated the licensing input costs that Microsoft would incur to host its customers' Windows Server and SQL Server usage if it paid the same input prices that it charges AWS and Google. To do this, we multiplied each Azure customer's total 2024 vcore hours of usage of each product by the respective per-vcore hour price for each product in each of AWS' and Google's SPLAs.²²⁴⁹ We then took these hypothetical licensing input costs as a proportion of each customer's:
- (a) total spend on Azure services, and
 - (b) spend on Windows Server VMs (and SQL Server IP where relevant).
- 7.388 We present the mean average licensing input costs as a proportion of each denominator for customers in each of five revenue brackets to give an indication of how the significance of the input may vary in relation to customers of different sizes.
- 7.389 For SQL Server, we also present the median licensing input costs as a proportion of each denominator to account for the fact that some revenue brackets encompass a broad range of heterogeneous customers. In particular, the lower brackets include a high concentration of customers at the bottom end of the spend distribution and a long tail of customers with significantly higher spend.
- 7.390 As explained in full in appendix T, our analysis covers the majority of Azure spend as we have excluded customers with <\$10k spend, which account for a minority of total Azure spend.

²²⁴⁹ Virtual core hours (vcore hours) are hours of usage normalised for the number of core processing units being used to run a particular instance or operating system environment (OSE). For example, using Windows Server OS on a VM that uses 4 CPUs for one hour constitutes 4 vcore hours of usage.

7.391 Google's SPLA with Microsoft [X]. We present the results using Google's prices below and present the full results in appendix T.

7.392 Table 7.4 below shows Microsoft's Windows Server licensing input costs if it paid the same input prices that it charges Google as estimated according to the methodology set out above. These costs are expressed as a share of customer spend across all Azure services and Windows Server VMs on Azure on average for Windows Server customers in each revenue bracket.

Table 7.4: Windows Server input costs as a proportion of customer spend on all Azure services and Windows Server VMs on Azure, Google SPLA prices

Revenue bracket (\$)	As a proportion of all Azure spend (%)	As a proportion of Windows Server VM spend (%)
10k-1M	[10-20]	[60-70]
1M-5M	[5-10]	[60-70]
5M-10M	[5-10]	[70-80]
10M-20M	[5-10]	[70-80]
>20M	[5-10]	[80-90]

Source: CMA analysis of Microsoft's submission to the CMA [X]

7.393 As discussed above, we consider that the significance of the input is best represented by a figure within the range of values presented for each revenue bracket. Therefore, it is likely that Windows Server accounts for at least [5-10]% and as much as [80-90]% of the relevant spend denominator for customers in each bracket.

7.394 Table 7.5 below shows Microsoft's SQL Server licensing input costs if Microsoft paid Google's SPLA input prices. These costs are expressed as a share of customer spend across all Azure services and separately as a proportion of spend on Windows Server VMs and SQL Server IP on Azure for SQL Server customers in each revenue bracket.

Table 7.5: SQL Server input costs as a proportion of spend on all Azure services and spend on Windows Server VMs and SQL Server IP on Azure, Google SPLA prices

Revenue bracket (\$)	As a proportion of all Azure spend (%)	As a proportion of Windows Server VM and SQL Server spend (%)
10k-1M	[10-20]	[30-40]
1M-5M	[0-5]	[10-20]
5M-10M	[0-5]	[10-20]
10M-20M	[0-5]	[10-20]
>20M	[0-5]	[20-30]

Source: CMA analysis of Microsoft's submission to the CMA [X]

7.395 Our analysis suggests that SQL Server accounts for at least [0-5]% and as much as [30-40]% of the relevant spend denominator for across the five highest revenue brackets.

7.396 Table 7.6 below show Microsoft's SQL Server input costs if it paid Google's SPLA input prices. These costs are expressed as the median share of customer spend

on all Azure services and spend on Windows Server VMs and SQL Server IP spend for customers in each revenue bracket.²²⁵⁰

Table 7.6: Median SQL Server input costs as a proportion of customer spend on all Azure services and spend on Windows Server VMs and SQL Server IP on Azure, Google SPLA prices

Revenue bracket (\$)	As a proportion of all Azure spend (%)	As a proportion of Windows Server VM and SQL Server spend (%)
10k-1M	[5-10]	[20-30]
1M-5M	[0-5]	[10-20]
5M-10M	[0-5]	[10-20]
10M-20M	[0-5]	[10-20]
>20M	[0-5]	[10-20]

Source: CMA analysis of Microsoft's submission to the CMA [X]

7.397 The median results are generally lower than the means across the revenue brackets. However, these results suggest that the significance of the SQL Server input for any single customer falls within a similar range of proportions as the mean results presented above.

7.398 Table 7.7 below shows Microsoft's combined Windows Server and SQL Server licensing input costs if it paid Google's SPLA input prices. These costs are expressed as a share of customer spend across all Azure services and spend

Table 7.7: Combined Windows Server and SQL Server input costs as a proportion of customer spend on all Azure services and spend on Windows Server VMs and SQL Server IP on Azure, Google SPLA prices

Revenue bracket (\$)	As a proportion of all Azure spend (%)	As a proportion of Windows Server VM and SQL Server spend (%)
10k-1M	[20-30]	[80-90]
1M-5M	[10-20]	[70-80]
5M-10M	[10-20]	[80-90]
10M-20M	[5-10]	[80-90]
>20M	[5-10]	[90-100]

Source: CMA analysis of Microsoft's submission to the CMA [X]

7.399 As discussed above, we consider that the significance of the input in competing for any single customer falls somewhere in the range that we present for each revenue bracket. Therefore, our analysis suggests that Windows Server and SQL Server combined account for at least [5-10]% and as much as [90-100]% of the relevant spend denominator across all revenue brackets.

7.400 Overall, the results presented above indicate that Windows Server constitutes a significant part of Google's cost base in competing for cloud customers, and particularly for customers that allocate their cloud spend in a less aggregated way.

²²⁵⁰ ACR stands for Azure Consumed Revenue and denotes total customer spend on Azure services

For customers that use both Windows Server and SQL Server PAYG (which account for [70-80]% of total Azure spend), the combined Windows Server and SQL Server input accounts for an even more significant proportion of the cost base.

- 7.401 As outlined above, as we are comparing the licensing input costs to spend rather than costs, the significance of the Microsoft software input in the cost base of rival cloud providers is likely to be higher than the proportions presented above.

Microsoft analysis of Windows Server VM revenues as a percentage of total cloud spend by customer

- 7.402 Microsoft also submitted an analysis of Azure customers' relative spend on Windows Server VMs to support its view that Windows Server is not a sufficiently important input to give rise to the ability to foreclose. In particular, Microsoft submitted that (a) AWS and Google are not foreclosed in relation to the average Windows Server-using Azure customer, or (b) customer sub-segments that show [REDACTED] Windows Server usage do not amount to material demand for the purposes of assessing whether AWS or Google are foreclosed.²²⁵¹
- 7.403 Regarding Microsoft's submission that AWS and Google are not foreclosed in relation to the average Azure customer, Microsoft calculated spend on Windows Server VMs as a proportion of total Azure cloud spend of a customer (as an unweighted average) in 2022. Microsoft found that Windows Server VM spend accounted for [10-20]% of total spend across all customers on average.²²⁵² Further, customers in the revenue brackets [REDACTED] spent [10-20]% and [5-10]% on Windows Server VMs, respectively.²²⁵³
- 7.404 Microsoft submitted that these results suggest that Windows Server is of low relevance to customers' choice of public cloud overall and particularly to large customers,²²⁵⁴ which are more valuable for the retention or building of scale.²²⁵⁵ Therefore, Microsoft submitted that Windows Server is not sufficiently important to give Microsoft the ability to partially foreclose AWS and Google to the extent that it prevents them from gaining or retaining scale.
- 7.405 As set out above, Microsoft also submitted that AWS and Google are not foreclosed in relation to customer sub-segments with [REDACTED] Windows Server usage. In this regard, Microsoft considered three customer groups: the [REDACTED] revenue bracket, new joiners to Azure, and users that ran Window Server VMs, but not Linux VMs. These groups allocated [20-30]% [20-30]% and [20-30]%, respectively,

²²⁵¹ Microsoft's submission to the CMA [REDACTED].

²²⁵² Microsoft's submission to the CMA [REDACTED].

²²⁵³ Microsoft's submission to the CMA [REDACTED].

²²⁵⁴ [Microsoft's response to the Licensing working paper](#), paragraph 7.21.

²²⁵⁵ [Microsoft's response to the Licensing working paper](#), paragraph 1.3, part d.

of their Azure spend on Windows Server VMs. Microsoft submitted that the Windows Server spend of these customers accounted for between [0-5]% of total Windows Server-user ACR and said that this is therefore not a material IP cost impact on rivals.^{2256,2257}

7.406 We noted certain issues with Microsoft's analysis and its interpretation of the results, some of which we have also identified in relation to previous analyses:

- (a) Microsoft's analysis used the overall average spend on Windows Server VM as the denominator to estimate the importance of Windows Server VMs to customer decision making. The relevant denominator depends on the composition of the bundle that each customer purchases. To the extent that customers disaggregate their decision making, Microsoft's estimate will understate the importance of Windows Server VMs.
- (b) Microsoft's analysis expresses its customers' licensing spend as a proportion of their respective total spend rather than as a proportion of their costs. We consider that using revenues as the denominator incorporates a margin which inflates the denominator and therefore risks understating the significance of the licensing spend to the cost base.
- (c) We do not consider that [10-20]% as a lower bound estimate, is an insignificant portion of total revenues;
- (d) In relation to Microsoft's submissions on the proportion that large customers spend on Windows Server VMs, we note that:
 - (i) While large customers are key to growing scale, we have found that they are also more likely to multi-cloud (see chapter 3). As such, the correct denominator for larger customers is likely to be narrower than for smaller customers. That is, Microsoft software workloads likely comprise a larger share of the total spend that large customers allocate to Azure within a single purchasing decision.
 - (ii) While large customers do spend more individually, in 2022, the top two revenue brackets together account for just under [%] of total Azure revenues, whereas the \$10k-1M bracket ([%]) accounts for [%] of total Azure revenues alone.²²⁵⁸

7.407 In relation to Microsoft's submission on customer groups with [%] Windows Server usage, we consider it is not appropriate to present their Windows Server spend as a proportion of total spend on Azure across all customers. We consider that

²²⁵⁶ Microsoft's submission to the CMA [%].

²²⁵⁷ Microsoft conducted this analysis on 2023 data.

²²⁵⁸ CMA analysis of Microsoft's response to the CMA's information request [%].

Microsoft's submission shows that Windows Server is likely to be a particularly important input to those customer groups because Windows Server accounts for a significant portion of their total spend.

- 7.408 We agree that these customer groups represent a small portion of the market in comparison to other groups of customers, but the groups chosen by Microsoft for their analysis are not exhaustive. Other groups may exist with similar usage, so the significance of a selection of three groups may understate the overall relevance of the input across the customer base.
- 7.409 We consider it relevant to assess average spend on Windows Server (i) across all customers, (ii) across customers that use Windows Server, and (iii) within customer revenue groups.
- 7.410 Using the data made available to us, we ran our own analysis considering Azure customers that use any Windows Server VMs on Azure. We calculated these customers' spend on Windows Server VMs as a proportion of total spend on all Azure services on average. Our analysis suggests that, for customers that use Windows Server, Windows Server VM revenues account for [20-30]% of total spend across all Azure customers on average.²²⁵⁹
- 7.411 Finally, we note that Microsoft's analysis does not account for the fact that many customers use both Windows Server and SQL Server. Combined spend on Windows Server and SQL Server is likely to account for a larger proportion of total spend for customers on average. Indeed, in our own analysis, we found that combined Windows Server VM and SQL Server spend accounted for [30-40]% [§] of total spend across all Azure customers that use Windows Server on average and at least [10-20]% of total spend in each revenue bracket.

Significance in shaping downstream competition

- 7.412 As outlined above, even if an input constitutes a small part of the providers' cost base in monetary terms, it may still be significant in terms of affecting the overall quality or attractiveness of providers' competitive offerings. We have therefore also assessed the significance of Microsoft software inputs by reference to their scope to determine the quality of particular downstream cloud services.
- 7.413 In this section, we set out evidence on:
- (a) the impact of certain non-price factors on the quality of relevant cloud services.

²²⁵⁹ CMA analysis of Microsoft's response to the CMA's information request [§].

- (b) AWS' and Google's internal documents on the importance of Microsoft workloads;
- (c) the importance of Microsoft software to the quality of cloud-based VDI offerings;
- (d) the current and future significance of VDI as a cloud workload; and

Cross-cutting evidence on the importance of quality factors

7.414 In assessing the importance of the Microsoft software input for AWS and Google, we consider not only the proportion of their costs that it accounts for, but also the role it plays as a determinant of product quality or the rate of innovation. We set out the evidence we have gathered in relation to this below.

7.415 We have considered the importance of quality differences using evidence directly from customers. This is because, while a firm theoretically may choose to pass on or absorb input costs, quality differences are by nature 'passed on' and customers therefore experience any differences directly.²²⁶⁰ As such, the extent to which customers consider quality differences in their choice of cloud can provide an indication, when considered alongside the input cost as a percentage of the cost base, on the importance of the input.

7.416 We asked customers about the differences, if any, between using Microsoft software products²²⁶¹ on Azure compared to using those products on other public clouds in terms of price and general quality factors (functionality, access to or timing of software updates and availability). We set out their answers in relation to quality factors below, though note that some of the factors identified may also be captured under price differences²²⁶²:

- (a) Most customers that responded did not identify differences in quality factors between public clouds;²²⁶³
- (b) One customer said that for IaaS there were 'some' Microsoft services absent on non-Azure clouds. As an example, the customer identified Office365, and said that there is therefore a need to continue to use Azure for this service. This customer also said that using Azure entitles customers to more upgrade rights for pre-October 2019 licences;²²⁶⁴ and

²²⁶⁰ We note that it may be possible for a firm to compensate for lower quality by providing price discounts. However, the trade-off between quality and price is not clear, and there are some elements of quality that customer will be unwilling to compromise on for the sake of lower prices.

²²⁶¹ The question referred to Windows Server, SQL Server, Windows 10/11, Microsoft Office, Microsoft 365 and Active Directory/Azure AD specifically.

²²⁶² Where this is the case, AWS and Google are theoretically more likely to be able to compensate customers through offering price discounts.

²²⁶³ Responses to the CMA's information requests [X].

²²⁶⁴ [X] response to the CMA's information request [X].

- (c) One customer said that it was previously unable to access Microsoft 365 on AWS. The customer explained that recent changes meant that it is now technically possible, but that customer still uses Azure for Microsoft 365 workloads because it is less expensive.²²⁶⁵

7.417 We asked customers about the specific quality differences that Google, CISPE and/or AWS said affect Microsoft software products on other public clouds (see above section Differences between using Microsoft software products on Azure compared to on non-Azure clouds via SPLA). As above, some of the factors identified may also relate to price factors.²²⁶⁶

- (a) Most customers we contacted felt that they were unable to answer the question because they had not compared these differences across clouds or that any differences were not relevant because they have a small Microsoft estate on public cloud.²²⁶⁷
- (b) A handful of customers identified some non-price differences. For example:
 - (i) Some customers identified the cost and length of security updates, including Extended Security Updates (ESUs).²²⁶⁸ One customer said that, while ESUs are free on Azure, ‘updates for [Windows] servers outside of Azure are extremely expensive’ and that ‘there is a benefit in duration of ESUs provided in Azure’.²²⁶⁹
 - (ii) Other factors that a few customers identified included the features and/or functionality available on software;²²⁷⁰ minimum purchase requirement, eg duration of licences or number of virtual cores;²²⁷¹ and the basis by which you are charged for use of Microsoft products on public cloud.²²⁷²

7.418 In addition, the Jigsaw report indicated that non-price factors may have contributed to some customers’ reasons for choosing Azure. For example, one customer said that they were able to get critical security patches for longer if they went to Azure.^{2273,2274}

7.419 Overall, the evidence suggests that most customers are currently not aware of quality differences in using Microsoft software products on Azure compared to rival

²²⁶⁵ [REDACTED] response to the CMA’s information request [REDACTED].

²²⁶⁶ To the extent that customers’ responses related to non-price differences in hosting VDI on Azure compared to other public clouds, we have considered their answers below.

²²⁶⁷ Responses to the CMA’s information requests [REDACTED].

²²⁶⁸ Responses to the CMA’s information requests [REDACTED].

²²⁶⁹ [REDACTED] response to the CMA’s information request [REDACTED].

²²⁷⁰ Responses to the CMA’s information requests [REDACTED].

²²⁷¹ Responses to the CMA’s information requests [REDACTED].

²²⁷² [REDACTED] response to the CMA’s information request [REDACTED].

²²⁷³ [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), paragraph 7.2.3.

²²⁷⁴ We understand critical security patches to include ESUs.

clouds. As such, on the whole we do not consider that current quality differences indicate that the input is more important than its proportion of the cost base would suggest (with the exception of VDIs, which are discussed below). However, there may be some quality factors, particularly security updates, that a subset of customers consider when making their choice of public cloud. As set out in chapter 3, security and data protection are key considerations for some customers in their negotiations with providers. For these customers, the qualitative differences in using the Microsoft software may be particularly important in terms of affecting the overall quality or attractiveness of AWS' and Google's competitive offerings.

AWS and Google internal documents

- 7.420 We asked AWS and Google to provide internal documents that discuss their ability to compete with Azure.²²⁷⁵
- 7.421 Almost all of Google's documents suggested that Microsoft workloads are important to the cloud market.²²⁷⁶ For example:
- (a) One document states that Windows workloads are [REDACTED]. The document also states [REDACTED].²²⁷⁷
 - (b) One document states that [REDACTED].²²⁷⁸
 - (c) Another document, which appears to be an internal slide deck on Microsoft licensing, states that '[REDACTED]' and identifies that the total addressable market for Microsoft workloads is [REDACTED].²²⁷⁹
- 7.422 AWS provided one responsive document, which suggested that using figures from a third-party industry report, Microsoft workloads are a significant part of the industry and important to win. This document states that Windows on cloud is [REDACTED]% of the overall Windows market segment revenue and continues to grow at [REDACTED]% year-on-year. The document also states that, with \$[REDACTED] Windows workloads still running on premises, AWS needs to capture this opportunity with continued investment. The document further states that Windows migrations are lagging behind Linux due to the complexity of the traditional enterprise workload migration projects that have commercial licences involved.²²⁸⁰

²²⁷⁵ We note that some of these documents have been produced during the course of the investigation, and have therefore been interpreted accordingly.

²²⁷⁶ Google's response to the CMA's information request [REDACTED].

²²⁷⁷ Google's response to the CMA's information request [REDACTED].

²²⁷⁸ Google's response to the CMA's information request [REDACTED].

²²⁷⁹ Google's response to the CMA's information request [REDACTED].

²²⁸⁰ AWS' response to the CMA's information request [REDACTED].

7.423 Overall, the AWS internal document and Google's internal documents suggest that they consider Microsoft workloads are an important segment, which they need to compete for, but that they have struggled to do so.

The importance of the Microsoft software products to VDI offerings

7.424 The primary use case for Microsoft's client-side products, ie Windows 10/11, Visual Studio, and Microsoft's productivity suites, on the public cloud is as part of Virtual Desktop Infrastructure (VDI) instances. This subsection sets out evidence on the importance of each of these products to the competitiveness of cloud-based VDI offerings.

7.425 We have found that Microsoft has significant market power with respect to all of these client-side products.

7.426 If the Microsoft products are important inputs to VDI offerings and, if VDI workloads are a sufficiently important input to cloud services (either alone or in combination with other Microsoft software products), Microsoft's licensing practices with relation to these products may have the scope to have an impact on rivals' competitive offerings.

7.427 Evidence shows that Microsoft software is important to the supply of cloud-based VDI:

- (a) One cloud provider submitted that [60-70]% of all [X] run Microsoft Office, and approximately [90-100]% of [X] instances run Windows instead of Linux.²²⁸¹
- (b) One cloud provider submitted that [90-100]% of VDI workloads run Windows as the operating system. It submitted that this is a conservative estimate based on qualitative feedback from VDI partners and its own analysis of [X] VMs which showed [90-100]% are based on Windows.²²⁸² An internal document submitted by this provider outlines that Microsoft Office is a strong requirement for customers with VDI infrastructure.²²⁸³

Current importance of cloud-based VDI workloads

Azure revenue analysis

7.428 As detailed above, our analysis shows that VDI currently accounts for [0-5]% of Azure revenues. This is a small proportion in absolute terms, but also [X] on Azure. However, given the importance of the Microsoft software input to VDI and

²²⁸¹ [X] submission to the CMA [X].

²²⁸² [X] response to the CMA's information request [X].

²²⁸³ Google's response to the CMA's information request [X].

the potential of the licensing practices to reduce competition (see below), we expect the proportion of market-wide revenues that VDI accounts for to be lower.

- 7.429 As part of our initial evidence gathering, we asked some customers what proportion of their total cloud spend was on VDIs and whether this was increasing or decreasing. Some customers responded with the proportion of their total cloud spend that they allocate to VDI. Around half of these spent between 10-20%²²⁸⁴ while the other half spent less than 10%.²²⁸⁵
- 7.430 We also contacted additional customers some of which were active in industries which are more likely to use VDI.²²⁸⁶ In response, most customers we contacted spent 5-10% of their total cloud spend on VDI.²²⁸⁷ However some spent up to 20-30%,²²⁸⁸ and some spent 0-5%.²²⁸⁹

Internal documents

- 7.431 We have also seen internal documents which detail the current size of the market for cloud based VDI. These documents suggest that VDI workloads are small relative to overall cloud workloads.
- (a) Google submitted that market research company iMARC estimated that the market for cloud-based VDI (across public, private & hybrid cloud) was worth \$6.8 billion USD in 2022 and \$7.8 billion USD in 2023 globally.²²⁹⁰
 - (b) Microsoft submitted an internal document which outlines that the global desktop as a service market was forecast to be worth [REDACTED] in 2022.²²⁹¹
 - (c) Microsoft also submitted an industry report which outlines that the public cloud desktop as a service end user spending in 2022 was \$2.7 billion in 2023.²²⁹²
 - (d) By way of comparison, the same industry report provides figures for worldwide public cloud services end user spending. For reference, in 2023, IaaS end user spending is reported as \$143 billion, and PaaS is reported as \$142 billion.²²⁹³

²²⁸⁴ Responses to the CMA's information requests [REDACTED].

²²⁸⁵ Responses to the CMA's information requests [REDACTED].

²²⁸⁶ In particular, healthcare and financial.

²²⁸⁷ Responses to the CMA's information requests [REDACTED].

²²⁸⁸ Responses to the CMA's information requests [REDACTED].

²²⁸⁹ Responses to the CMA's information requests [REDACTED].

²²⁹⁰ Google's response to the CMA's information request [REDACTED]. [Cloud-Based VDI Market Scope & Dynamics | Forecast 2032 \(imarcgroup.com\)](#).

²²⁹¹ Microsoft's response to the CMA's information request [REDACTED].

²²⁹² Microsoft's response to the CMA's information request [REDACTED].

²²⁹³ Microsoft's response to the CMA's information request [REDACTED].


- 7.432 Our analysis suggests that VDI workloads make up a small proportion of Azure revenues. Customer evidence is mixed, with some customers spending relatively little on VDI workloads; however, for some VDI workloads are a relatively large portion of their overall cloud spend. Evidence from internal documents suggests that cloud based VDI workloads currently make up a relatively small percentage of overall cloud workloads. Overall, the evidence suggests that currently VDI workloads make up a relatively small portion of public cloud workloads.

Future importance of cloud-based VDI workloads

Usage rates of virtual machines on Azure

- 7.433 We undertook an analysis of the relative usage from 2020-2023 of VMs that use Windows Server, Linux and Windows 10/11 as the OS on Azure.
- 7.434 As noted above, we understand that usage of Windows 10/11 VMs represents the underlying compute resource used to run VDI instances on Azure. Therefore, to the extent that Windows 10/11 VM usage as a share of total VM usage on Azure is growing over time, this would suggest that VDI is becoming an increasingly important workload, ie that demand for VDI is outgrowing demand for other IaaS workloads. This could also be consistent with VDIs becoming a larger share of Microsoft's total cloud revenues in the future.
- 7.435 We used Microsoft data on the total annual vcore hours of usage of Azure VMs by UK customers over the Years 2020-2023, segmented by OS. We used this data to calculate shares of usage by each OS, as well as annual growth rates and CAGRs for each OS.
- 7.436 Figure 7.2 below shows shares of Azure VM usage by operating system over the Years 2020-2023.

Figure 7.2: Shares of Azure VM usage by operating system 2020-2023.



Source: CMA analysis of Microsoft's response to the CMA's information request [2293].

7.437 Our analysis suggests that the share of Windows 10/11 VM usage has increased from [2294] in 2020 to [2295] in 2023. Since 2020, Windows 10/11 VM usage has increased at a CAGR of [2296]%, compared with CAGRs for Windows Server and Linux of [2297]% and [2298]% respectively.²²⁹⁴ By way of comparison, the market for public cloud (IaaS and PaaS) in the UK grew by [30-40]% per year on a compound annual basis.²²⁹⁵

7.438 These relative growth rates suggest that VDI could grow to account for a larger proportion of Azure revenues in the future. Further, we also understand that Windows Server VMs can support VDI workloads.²²⁹⁶ As such, our analysis may understate the share of VM resource that is used to support VDI workloads on Azure.

7.439 We have not analysed VDI usage on other public clouds. However, given our understanding that the majority of cloud-based VDI workloads are deployed on

²²⁹⁴ CMA analysis of Microsoft's response to the CMA's information request [2293]. We note that Microsoft began offering VDI services in 2019.

²²⁹⁵ CMA analysis of first-party revenue data and IDC and Synergy data. Refer to chapter on Market structure and concentration for further information.

²²⁹⁶ [2293] response to the CMA's information request [2293].

Azure, we consider that this analysis provides a good indication of the general growth in demand for cloud-based VDI.

VDI provider views

- 7.440 We heard from VDI providers that the market for VDIs on the cloud (also referred to as Desktop-as-a-Service or DaaS) is expected to grow rapidly, with no third party VDI provider suggesting that this market would not grow less than 10% per year on a compound annual basis.
- (a) A VDI provider said that VDIs are becoming increasingly important for cloud customers due to the additional security, flexibility and performance they can offer compared to a traditional PC device.²²⁹⁷
 - (b) The same VDI provider told us that according to a report by a third party market intelligence firm, VDIs currently account for around 7% of the Enterprise PC market, but that they expect this to grow to 20-30% in the coming years.²²⁹⁸ The market intelligence firm also forecast that cloud-based VDI (DaaS) would generate an additional \$14-38 billion dollars for Azure over the next 3-5 years and \$75-100 billion in the long run. By comparison, the VDI provider submitted that AWS' current global annual revenue is around \$75-80 billion.²²⁹⁹
 - (c) We heard from another VDI provider that it has seen healthy growth in VDI deployments on the public cloud. In particular, where almost all of its VDI instances were deployed on-premises or private cloud in 2018, approximately 25% are now deployed on public cloud in 2024. However, it noted that much of this early growth was prompted by the pandemic.²³⁰⁰

Cloud providers' views

- 7.441 We asked cloud providers to provide internal documents relating to the market landscape for the supply of VDI/DaaS and/or the forecast of growth opportunity for VDI/DaaS in the cloud. We consider that public cloud is expected to grow in the region of 15-20% on a compound annual basis over the next five to ten years.²³⁰¹

²²⁹⁷ Note of meeting with [REDACTED].

²²⁹⁸ Note of meeting with [REDACTED].

²²⁹⁹ Note of meeting with [REDACTED].

²³⁰⁰ Note of meeting with [REDACTED].

²³⁰¹ We have used publicly available long run CAGR estimates for worldwide public cloud growth estimates as a baseline for growth of the public cloud. Some of these estimates are 15.1% CAGR 2023-2028 (Source: Markets and Markets, [Cloud Computing market by service model](#), accessed on 7 October 2024), 16.5% CAGR 2022-2032 (Source: Business research insights, [Cloud Infrastructure Services Market Size, Share, Growth, and Industry Analysis](#), accessed on 7 October 2024) and 19.5% CAGR 2023-2028 (Source: IDC Research, [Worldwide Public Cloud Services Revenues Grew 19.9% Year Over Year in 2023](#), accessed on 7 October 2024).

We note that in particular some of these estimates may not align with our definition of public cloud (for example the inclusion of SaaS), but we consider that despite this caveat, these estimates provide an indication of the future trajectory against which to compare VDI workloads.

Some internal documents point towards a growth in importance of VDI workloads on the public cloud, for example:

- (a) Microsoft submitted a third party market intelligence firm's report which said that it expected the global market for VDI on the public cloud to grow at 30.59% per year from \$5.4 billion USD in 2021 to \$20.3 billion USD in 2026.²³⁰² The report also projected the EMEA market for VDI to grow at a compound annual rate of 31.99% between 2021 to 2026.
- (b) Microsoft has also submitted a third party market intelligence report which outlines that Worldwide VDI spend from 2021-2026 is forecast to be driven predominantly by gains in IaaS (rather than gains in software or on premise server storage).²³⁰³

7.442 However, we have also received evidence from cloud providers that suggests that demand for cloud-based VDI may not outgrow the broader cloud services market, or outgrow it only marginally, for example:

- (a) An IMARC Group report referenced suggests that the compound annual growth rate for the market for cloud based VDI to be 14.1% for years 2024-2032, growing to \$26.6 billion USD by 2032.²³⁰⁴
- (b) Similarly, a Microsoft internal document forecasted the market for DaaS in the public cloud to grow at [10-20]% between 2022 and 2028 on a compound annual basis.²³⁰⁵ This growth is accelerated from [10-20]% CAGR (across all cloud deployment types) from 2017-2022.²³⁰⁶
- (c) The same Microsoft internal document dated January 2023 reported that [REDACTED] of all enterprise PC users will rely on DaaS within three years.²³⁰⁷

Customer views

7.443 We asked customers, some of which were active in industries which are more likely to use VDI,²³⁰⁸ whether they expected their spend on VDI as a proportion of total cloud spend to increase, decrease or remain the same. Most customers expected VDI as a proportion of their cloud spend to remain the same,²³⁰⁹ and a few customers expected VDI as a proportion of total cloud spend to increase.²³¹⁰

²³⁰² Microsoft's response to the CMA's information request [REDACTED].

²³⁰³ Microsoft's response to the CMA's information request [REDACTED].

²³⁰⁴ Imrac Group, [Cloud-Based VDI Market Scope & Dynamics](#), accessed on 4 October 2024.

²³⁰⁵ Microsoft's response to the CMA's information request [REDACTED].

²³⁰⁶ Microsoft's response to the CMA's information request [REDACTED].

²³⁰⁷ Microsoft's response to the CMA's information request [REDACTED].

²³⁰⁸ In particular, healthcare and financial services.

²³⁰⁹ Responses to the CMA's information requests [REDACTED].

²³¹⁰ Responses to the CMA's information requests [REDACTED].

Our assessment

- 7.444 VDI currently accounts for a small proportion of Azure UK public cloud revenues and likely a smaller proportion of market-wide revenues.
- 7.445 The evidence suggests that VDI is growing and is likely to become a more important cloud workload over time. In particular, the relative growth in demand for VDI compared with other IaaS workloads on Azure suggests that VDI is growing. As noted above, our analysis suggests that VDI may grow to account for a larger share of cloud demand in the future. However, we have also seen some evidence that suggests that VDI will not outgrow demand for other cloud services.
- 7.446 Overall, there is mixed evidence as to whether VDI workloads will grow to be more important relative to other public cloud workloads over time, but nonetheless, VDI workloads are a sufficiently important input into rivals' competitive offerings such that, when taken together with the other Microsoft software products, they contribute to Microsoft's ability to worsen rivals' competitive offerings at present.

Factors which may affect the importance of the input in the future

- 7.447 As set out above, our analysis suggests that Microsoft's software inputs currently constitute a significant part of AWS' and Google's cost base in the supply of cloud services. However, demand for cloud services is likely to evolve over time and this could increase or reduce the importance of the Microsoft software inputs relative to customers total Azure spend in the future.
- 7.448 This section considers evidence on whether the importance of Microsoft's software inputs will change in the foreseeable future.
- 7.449 If there is a significant pool of potential demand for Microsoft software workloads that remains on-premises, then these customers may migrate to the cloud in the future. If this migration causes Microsoft software workloads to grow at a faster rate than the total cloud market, then the Microsoft software input would grow to account for a larger proportion of total Azure spend.
- 7.450 However, Microsoft software could shrink as a proportion of total spend if non-Microsoft software workloads grow at a faster rate, thereby diluting the significance of Microsoft software workloads relative to other workloads.
- 7.451 Therefore, the impact of migration from on-premises to the cloud on the future importance of the Microsoft inputs depends on the interplay between the following factors:
- (a) the relative size of the on-premises IT and public cloud markets;
 - (b) the share of on-premises workloads that rely on Microsoft software;

- (c) the rate of migration from on-premises to the cloud;
- (d) the extent to which customers switch away from Microsoft software as they migrate to the cloud; and,
- (e) the relative growth in non-Microsoft software workloads.

7.452 The net effect depends on the relative strength of each factor. For example, a relatively large on-premises IT market comprising a high share of workloads that rely on Microsoft software suggests that there is a larger pool of potential demand for such workloads on the cloud. However, if migration to the cloud slows, or customers switch away from Microsoft software as they migrate to the cloud, less of this potential demand will be realised. Additionally, if workloads that do not use the Microsoft software grow faster than workloads that do use the Microsoft Software, then the effect of any influx in Microsoft software workloads will be diluted by a relatively greater increase in the denominator compared to the increase in the numerator.

7.453 The relative strength of these factors could alter our estimates of the importance of the input that compare the input to customers' total Azure spend; our estimates using the minimum bundle as a comparator would not be impacted by any change in relative demand for Microsoft software workloads. In addition, as explained above, we consider that it is appropriate to compare the Microsoft input to spend on the set of services that customers include in the same purchasing decision as those that use the Microsoft software.

7.454 This section presents evidence on each factor listed above and assesses the likely net impact on the future importance of the input relative to customers' total Azure spend.

The size of the on-premises IT market

7.455 A cloud provider submitted three pieces of evidence on the relative size of the on-premises IT market based on third-party data:

- (a) The cloud provider's analysis [REDACTED] suggests that the total addressable market for public cloud in the UK in 2022 comprised \$10.1 billion from customers that are using public cloud already and \$11.3 billion from on-premises customers;²³¹¹
- (b) A Goldman Sachs Research survey from September 2024 estimated that about 30% of workloads have migrated to the cloud thus far;²³¹²

²³¹¹ [REDACTED] response to the CMA's information request [REDACTED].

²³¹² Google's submission to the CMA [REDACTED]. [Cloud revenues poised to reach \\$2 trillion by 2030 amid AI rollout | Goldman Sachs](#)

- (c) [redacted] survey from 2024 found that [70-80]% of respondents had migrated [50-60]% or fewer of their workloads to the cloud.²³¹³

7.456 The same provider submitted quotes from senior Microsoft executives on potential cloud growth and further migration from on-premises:

- (a) Microsoft's CEO stated in July 2023 that 'it remains early when it comes to the long-term cloud opportunity',²³¹⁴ and told investors in early 2024 that 'Azure again took share as organisations bring their workloads to the cloud'.²³¹⁵
- (b) Microsoft's CCO said in December 2024 that 'there's still an awful lot of long tail cloud growth out there and we're still really only in the middle innings of cloud... [T]he hype cyclists switched to AI, but let's not forget that there is an awful lot of money yet to be had in migrating SQL Server, Windows Server, [and] VMware', and added 'there is a lot of headroom on the migration front'.²³¹⁶

7.457 Microsoft responded to this latter quote that it 'fully agrees there remains a large portion of IT workloads that are yet to be migrated to the cloud. And a material proportion of these workloads will include some spend on Windows Server and SQL Server'.²³¹⁷ However, Microsoft noted that it expects these workloads to 'comprise a small portion of total cloud spend as customers modernise to deploy in the cloud.' We assess evidence on the extent to which modernisation may reduce customers' reliance on Microsoft software in [Market Power] and below.

7.458 Based on data from our shares of supply analysis, we conducted our own analysis of the share of Server OS revenues generated from on-premises compared with cloud-based deployments. This suggests that 46% of total revenue from sales of Server OS products came from cloud deployments in 2023, and 54% from on-premises and other deployments.²³¹⁸ We note that this covers only paid deployments of server OS, and therefore covers paid versions of Windows, Linux and Unix variants. This suggests that a large proportion of total Server OS demand remains on-premises.

7.459 Overall, we note that these are relatively high-level estimates but there is consistency between third-party estimates, the views of Microsoft, and our own analysis. Therefore, we consider it likely that a significant proportion, and perhaps even the majority, of total demand for IT services remains on-premises.

²³¹³ Google's response to the CMA's information request [redacted].

²³¹⁴ [Microsoft's Annual Report 2023](#).

²³¹⁵ [Microsoft's Q1 2024 earnings call](#).

²³¹⁶ Google's submission to the CMA, [redacted].

²³¹⁷ Microsoft's response to the CMA's information request [redacted].

²³¹⁸ CMA analysis of IDC data.

The share of on-premises IT workloads that rely on Microsoft software

- 7.460 As set out in appendix R, Windows Server had a [70-80]% share of supply in the Server OS market deployed in non-cloud environments in 2022.²³¹⁹
- 7.461 Additionally, one party submitted a bespoke analysis [REDACTED] that suggests that Windows Server workloads account for 60-70% of all on-premises workloads running in the UK.²³²⁰
- 7.462 Given that on-premises workloads generally rely on a Server OS, shares of supply in Server OS provide direct evidence of the share of on-premises workloads that rely on Windows Server.
- 7.463 Overall, therefore, we consider that there is a substantial pool of potential demand for Microsoft workloads that remains on-premises.
- 7.464 Microsoft has also submitted data on the number of Windows Server and SQL Server licences sold to on-premises customers that would be eligible for the AHB in each year. As outlined above, customers with on-premises licences and an active Software Assurance subscription can effectively migrate those licences to Azure for no extra cost. This evidence shows that the number of active Software Assurance subscriptions for Windows Server has remained stable over the period 2020-2023, while the number of subscriptions for SQL Server has increased each year. This suggests that demand for Microsoft software licences that can be migrated to the cloud remains stable.²³²¹
- 7.465 However, as noted above, the proportion of this potential demand that will be realised on the cloud depends on the rate of migration to the cloud and the extent to which customers switch away from Microsoft software as they migrate workloads to the cloud.

The rate of migration to the cloud

- 7.466 A cloud provider submitted a [REDACTED] survey of US cloud customers undertaken in 2023 which suggests that workloads that migrated to the cloud from on-premises, and/or private hosted cloud, are expected to drive an increasing proportion of public cloud usage.²³²² This survey found that workloads migrated to the cloud were predicted to account for [50-60]% of total public cloud usage (excluding SaaS) in 2024, up from [10-20]% in 2022.²³²³

²³¹⁹ Here we refer to the metric of share of supply based on installed base, as we consider this most appropriate when considering the prevalence of Windows Server usage in non-cloud environments.

²³²⁰ [REDACTED] response to the CMA's information request issued [REDACTED].

²³²¹ Microsoft's response to the CMA's information request [REDACTED].

²³²² [REDACTED] submission to the CMA [REDACTED].

²³²³ [REDACTED] response to the CMA's information request [REDACTED].

7.467 We have placed limited weight on this survey, but we note that there is no evidence in the survey that the migration of on-premise workloads to the public cloud is reducing and that Microsoft software as an input is declining in importance.²³²⁴

7.468 As outlined above, we conducted our own analysis of the share of Server OS revenues generated from on-premises compared with cloud-based deployments. Looking at these results over time gives an indication of the rate of migration. Table 7.8 below shows the results of this analysis.²³²⁵

Table 7.8: Percentage of revenue from sales of Server OS deployed on the public cloud

2020	2021	2022	2023
39.5%	41.5%	43.7%	46%

Source: CMA analysis of IDC data.

7.469 The table shows that over the period 2020 to 2023 the percentage of revenues from sales of server OS deployed on the public cloud has increased from 39.5% to 46%. This analysis is consistent with rate of migration to the cloud being steady or perhaps increasing.

7.470 In addition, AHB usage accounted for [40-50]% of total Windows Server usage on Azure in 2024, up from [30-40]% in 2022.²³²⁶ It would be irrational for customers to deploy more expensive Windows Server PAYG VMs before exhausting their available licences that qualify for the AHB. This suggests that a relative increase in AHB usage is likely driven by new customers migrating to the cloud, rather than existing customers migrating their licences that qualify for AHB. Therefore, this evidence is also consistent with the rate of migration to the cloud increasing.

7.471 Microsoft has submitted that AI is speeding up migrations to cloud. In particular, it submitted that its internal research indicates that AI is speeding up customer migrations, with [X] of organisations listing AI as the top factor in accelerating cloud migration and modernisation, ²³²⁷ although from the report it is unclear what proportion of customers are referring to AI increasing migration and modernisation as distinct from each other. Evidence from cloud providers, as set out in chapter 4 on Impact of AI on cloud services, suggests that, in part, migration of on-premises workloads is being driven by a desire to use AI PaaS and SaaS services. We consider that AI may be an emerging factor which may speed up migration of services (including potentially Microsoft software related workloads) to the cloud.

²³²⁴ See appendix C.

²³²⁵ CMA analysis of IDC data.

²³²⁶ 2022 figure comes from dividing the total AHB Windows Server hours of usage from Microsoft's response to the CMA's information request [X] by the total hours of Windows Server usage in Microsoft's response to the CMA's information request [X] for the year 2022. For the 2024 figure, we made the equivalent calculation using Microsoft's submission to the CMA [X].

²³²⁷ Microsoft's submission to the CMA [X].

7.472 While the rate of migration to the cloud over time is inherently uncertain, we have not seen evidence to suggest that the rate of migration will slow.

The extent of switching away from Microsoft software

7.473 We consider customers' ability to switch away from the Microsoft software products above in Market power. Based on our assessment of the evidence, it appears unlikely that in the foreseeable future a large number of customers will be able to switch a significant number of their workloads away from relying on Microsoft software, including through modernisation in the public cloud.

Relative growth in non-Microsoft software workloads

7.474 As noted above, even if a significant pool of potential demand for Microsoft software workloads is migrated to the cloud, the significance of Microsoft software as an input relative to the rest of the cost base could be diluted if non-Microsoft software workloads grow at a faster rate.

7.475 We conducted our own analysis of the proportion of Azure customers' total spend that is allocated to Windows Server VMs and SQL Server IP in each year.²³²⁸ If this proportion were falling over time, then this could suggest that demand for Microsoft software workloads is growing at a slower rate relative to growth in workloads that do not make use of Microsoft software.

7.476 Table 7.9 below shows Azure customers' relative spend on Windows Server VMs over the period 2022-2024.

Table 7.9: Azure customers' relative spend on Windows Server VMs

Revenue bracket (\$)	May – Dec 2022 (%)	2023 (%)	2024 (%)
<10k	[30-40]%	[20-30]%	[20-30]%
10k-1M	[10-20]%	[10-20]%	[10-20]%
1M-5M	[10-20]%	[10-20]%	[10-20]%
5M-10M	[10-20]%	[10-20]%	[10-20]%
10M-20M	[10-20]%	[10-20]%	[10-20]%
>20M	[10-20]%	[10-20]%	[5-10]%
All customers	[10-20]%	[10-20]%	[10-20]%

Source: CMA analysis of Microsoft's submission to the CMA [X].

7.477 This analysis shows that customer spend on Windows Server VMs as a proportion of total Azure spend has declined over the last three years.

7.478 As outlined above, Windows Server has retained a stable share of OS usage on Azure VMs of around [40-50]% over 2021-2023, which suggests that Windows Server workloads are not declining in importance relative to other IaaS workloads. However, the analysis presented in Table 7.9 above also considers spend on PaaS services.

²³²⁸ Customers that do not use Windows Server are excluded from this analysis.

- 7.479 Therefore, these results suggest that PaaS services which do not include Microsoft software may be growing to account for a greater proportion of total Azure spend over time. This is consistent with our analysis of Microsoft's revenue data which shows that Virtual Machines revenues have declined from [REDACTED] of total Azure revenue in 2023 to [REDACTED] in 2024.²³²⁹
- 7.480 If this trend continues, then the Microsoft software could become less important relative to *total cloud spend* over time.²³³⁰
- 7.481 However, we note that the decline in relative spend on Windows Server VMs and SQL Server IP for customers in the lower revenue brackets has been smaller than in the higher revenue brackets. Analysis submitted by [REDACTED].²³³¹ Additionally, as noted above, for customers that joined between May 2022 and December 2023, [70-80]% of new Azure customers (accounting for [80-90]% of spend) use Windows Server in their first year.²³³² It is therefore possible that the decline in relative spend on Windows Server VMs is being driven by mature customers that are increasing their use of other workloads.

Growth of AI workloads

- 7.482 We have examined whether Microsoft software workloads could diminish in importance as more customers adopt AI services and their total Azure spend increases.
- 7.483 Microsoft has submitted that the foreclosure analysis is directly impacted by the rise of AI, as i) each hyperscaler has made AI central to their cloud strategies and spending, and ii) each hyperscaler has their own proprietary AI value propositions which they can sell alongside hosting legacy workloads with SQL Server and Windows Server.²³³³ Microsoft has also submitted that AI is relevant to Windows Server and SQL Server workloads as the power of AI relies on enterprises' own data which on premise is in legacy corporate databases and that SQL Server lends itself to being used in conjunction with AI.²³³⁴ It submitted that Google is particularly strong in the AI/database intersection.²³³⁵

²³²⁹ Comparison of the figures presented above in Distribution of cloud use by spend and those presented in the same section in the PDR (2024 vs 2023 data). Source: CMA analysis of Microsoft's response to the CMA's information request [REDACTED]

²³³⁰ As noted above, the factors discussed in this section enable us to consider whether the importance of the Microsoft software inputs relative to customers' total Azure spend will increase or decrease in the future. The factors discussed in this section would not alter our importance of the input calculations when considering importance of the inputs relative to customers' spend on Window Server VMs on Azure (or Windows Server VMs and SQL Server licensing IP on Azure).

²³³¹ Microsoft's submission to the CMA [REDACTED].

²³³² Microsoft's submission to the CMA [REDACTED].

²³³³ Microsoft's submission to the CMA [REDACTED].

²³³⁴ Microsoft's submission to the CMA [REDACTED].

²³³⁵ Microsoft's submission to the CMA [REDACTED].

- 7.484 We have considered whether AI will expand the ‘all Azure spend’ denominator, with respect to AI PaaS, and IaaS based on accelerated compute, and whether AI is likely to be included as part of bundles with Microsoft software.
- 7.485 We note that AI PaaS is included as part of the denominator considering total Azure spend. Evidence suggests that this currently accounts for a small proportion of cloud services revenues:
- (a) Customers are more commonly adopting AI PaaS rather than IaaS for accelerated compute. Our AI customer evidence suggests when customers considered AI capabilities, they considered access to FMs, as well as AI enabled PaaS and SaaS solutions rather than access to AI infrastructure. AI PaaS services are already accounted for in the ‘all Azure spend’ denominator (see chapter 4: Impact of AI on cloud services).
 - (b) At present, AI PaaS accounts for a small portion of cloud provider revenues. Evidence gathered suggests that currently most AI revenues are coming from a small number of customers using IaaS for accelerated compute (see chapter 4).
 - (c) However, we note that over time AI PaaS may account for a larger portion of revenues. This would increase the size of the ‘all Azure spend’ denominator and decrease the importance of Microsoft revenues overall.
- 7.486 Considering IaaS based on accelerated compute, evidence gathered suggests that a small number of customers contribute a large portion of these revenues (see chapter 4). Therefore, this would potentially stand to increase the relevant denominator for only a small subsection of customers at present.
- 7.487 With respect to SQL Server in particular, we also asked providers to explain the concept of ‘data gravity’ (see appendix R: Market power - supplementary framework and evidence). This concept suggests database and interfacing services should be hosted near to the data, and may support the hosting of data, database and other services alongside AI applications. Therefore, it may be the case that AI workloads may be a plausible service to be aggregated with Microsoft workloads, in particular SQL Server workloads. However, as we outline in below, in [counterstrategies section], it is not clear that AI workloads are currently a driving factor in customer choice of cloud.
- 7.488 We recognise that AI could increase the relevant denominator in the future but have not seen evidence to suggest that this is currently having or will have a material impact in the foreseeable future.

Our assessment of the likely importance of Microsoft software as an input in the future

- 7.489 This section has considered evidence on the likelihood that the relevant Microsoft software inputs will become more or less important relative to total cloud demand in the foreseeable future.
- 7.490 If the evidence suggested that demand for cloud services that use Microsoft software is likely to be significantly outgrown by demand for other services, then this could suggest that the Microsoft software may cease to be an important input in the future.
- 7.491 As noted above, the evidence presented in this section is only relevant to the consideration of whether the importance of the Microsoft software inputs relative to customers' *total Azure spend* will increase or decrease in the future.
- 7.492 The evidence shows that Microsoft software relative to total Azure spend will remain important in the foreseeable future.
- 7.493 There appears to be a significant pool of potential demand for Microsoft workloads that remains on-premises, and it appears likely that the rate of migration will at least remain stable.
- 7.494 Additionally, it appears unlikely that in the foreseeable future a large number of customers will be able to significantly switch away from Microsoft software, including through modernisation in the public cloud.
- 7.495 However, our analysis of the proportion of Azure customers' total spend that is allocated to Windows Server VMs and SQL Server suggests that customer spend on non-Microsoft software workloads has been growing faster than Microsoft workloads in recent years.

Our assessment

- 7.496 We have considered the importance of Microsoft software inputs, both in terms of their significance to the cost base of AWS and Google and their potential to shape downstream competition, for example by degrading the quality of AWS and Google's offerings. The proportion of downstream costs or revenues that the upstream input must account for to be deemed important is context specific and there is no single cost proportion threshold that an input must meet.
- 7.497 Our view is that Microsoft's analysis of AWS' and Google's SPLA costs understates the importance of the Microsoft software input.
- 7.498 Our analysis of the significance of Windows Server and SQL Server relative to competition for cloud services suggests that the Windows Server input accounts

for at least [5-10]% and as much as [80-90]% of the relevant spend denominator across all revenue brackets. In addition, for customers that use both Windows Server and SQL Server the combined Windows Server and SQL Server input accounts for at least [5-10]% and as much as [90-100]% of the relevant spend denominator across all revenue brackets.

- 7.499 As our analysis compares the licensing input costs to revenues rather than costs, it may understate the significance of the software input in their cost bases.
- 7.500 The estimates using either the minimum bundle or total spend on Azure services in each revenue bracket are unlikely to represent the significance of each product in competing for any single customer. Rather, values within the range between them are likely to be more reliable as indicators and we have therefore placed more weight on such values.
- 7.501 We have found that across all revenue brackets, Windows Server and SQL Server combined account for a material proportion of overall costs. In particular, we have found that input cost proportions for Windows Server and SQL Server combined in the middle of each range, which we consider to be a closer reflection of their importance, are very high. We have also found that input cost proportions for Windows Server on its own are also very high, suggesting it is an important input on its own.
- 7.502 AWS and Google state that Microsoft workloads are important in the cloud services markets and this is supported by some internal documents.
- 7.503 Our evidence on the future importance of the Microsoft inputs shows that they will likely remain important inputs in the foreseeable future.
- 7.504 We have also assessed whether particular features that determine the quality of Windows Server and SQL Server could mean that the input cost share percentages understate their importance as inputs. The evidence suggests that most customers are not aware of any differences in quality or features. Therefore, we consider that, for most customers, these inputs do not shape downstream competition beyond their significance in the cost base. For the subset of customers that do consider quality factors of Microsoft software in their choice of cloud, Microsoft software may be a particularly important input in terms of affecting the overall quality or attractiveness of AWS' and Google's competitive offerings.
- 7.505 Our assessment of the importance of the client-side Microsoft software products suggests that Windows 10/11 and the Microsoft productivity suites are important to the provision of cloud-based VDI services, while Visual Studio is an important input for particular customers. However, evidence of the potential for VDI to increase in importance over time is mixed.

7.506 We consider that Microsoft's client-side products do not constitute an important input on their own but they contribute to the overall importance of the Microsoft software inputs.

Conclusions on the importance of Microsoft software inputs to AWS and Google

7.507 Our assessment shows that, taken together, the relevant Microsoft software products are sufficiently important as inputs to AWS and Google in terms of affecting the overall quality or attractiveness of their competitive offerings. In particular, Windows Server and Windows Server and SQL Server combined are sufficiently important inputs to shape downstream competition on their own and they give Microsoft the ability to reduce competition in cloud services by materially disadvantaging AWS and Google.

Microsoft's supply of software: price and non-price differences

7.508 In this section, we set out evidence on the price and non-price differences between the Microsoft software that Microsoft offers to Azure customers and the Microsoft software input that it supplies to AWS and Google. We have considered the significance of any differences in price and/or quality in relation to customer decisions because the greater the differences in the input, the more likely it is that these have an effect on competition.

7.509 We consider:

- (a) a comparison of the prices that Microsoft charges customers to use Windows Server and SQL Server on Azure to the prices it charges AWS and Google for the software via their respective SPLAs;
- (b) an analysis submitted by a cloud provider which compares the licensing costs it pays for Windows Server to Azure's infrastructure price;
- (c) differences in list prices between Microsoft and AWS or Google; and
- (d) the evidence related to the price and non-price differences for the client-side products Visual Studio, Microsoft 365, Office and Windows 10/11.

Our comparison of the price that Microsoft charges Azure customers for Windows Server and SQL Server and the price it charges AWS and Google

7.510 We have assessed whether, and if so the extent to which, Microsoft's licensing practices gives rise to a price disadvantage for AWS and Google in competing for public cloud customers.

7.511 As part of this assessment, we compared the prices that Microsoft charges its own customers to use Windows Server and SQL Server with the input prices that

Microsoft charges AWS and Google via their respective SPLA contracts to resell Windows Server and SQL Server as part of their own cloud solutions.

- 7.512 This analysis is relevant to our assessment because the bigger the difference between the prices that Microsoft charges its own customers and the input prices to AWS and Google, the greater the potential for there to be a price disadvantage on AWS and Google (ie the bigger the price difference, the more difficult AWS and Google will find it to cover their other cloud costs while remaining competitive with Microsoft on price). For example, if as the result of this analysis we were to find no or only small price differences we would likely consider that Microsoft's licensing practices do not create any significant price disadvantages.
- 7.513 In order to assess the potential for there to be a price disadvantage that AWS and Google face when competing with Microsoft, we would ideally compare Microsoft's own Windows Server and SQL Server unit input costs with the prices it charges to AWS and Google for each product. However, we do not have access to data on Microsoft's licensing input costs. Therefore, as a proxy, our analysis compares the prices that Microsoft charges its customers to use Windows Server and SQL Server on Azure (ie its customer-facing prices) to the input prices that it charges AWS and Google.
- 7.514 As such, this analysis understates the potential for there to be a price disadvantage that AWS and Google face as it does not account for the margin that Microsoft likely charges over its own licensing input costs. That is, even if Microsoft charges a higher price to its customers than it charges to AWS and Google, the latter might still face a price disadvantage in competing for cloud customers.
- 7.515 At the outset, we note that (as set out in the background section) the AHB allows customers with qualifying existing on-premises Windows Server licences and Software Assurance subscriptions to effectively migrate these licences onto Azure without paying any additional licensing fees, whereas they would incur additional licensing fees if they migrated Windows Server workloads to AWS or Google. Whilst in practice customers' workloads are more complex in that even customers that qualify for the AHB may end up paying Windows Server licensing fees (eg because they also have PAYG workloads), this analysis focuses on the potential for the AHB to create a price difference.²³³⁶ We note that in this simplified scenario, we are comparing minimal costs to positive costs.

²³³⁶ Microsoft has submitted that this analysis creates a fictional comparison as customers cannot purchase a Windows Server licence separate to the VM that it runs on and, therefore, that a Windows Server VM is the smallest unit of demand that could make for a price comparison between Azure, AWS, and GCP (even if this is not the smallest unit purchased in reality). (Microsoft submission to the CMA [§].) Our analysis does not seek to replicate the price comparison that customers make across providers, but rather to assess the potential for the AHB to place AWS and Google at a price disadvantage. As set out in appendix T, the effective price that customers pay to migrate their on-premises Windows Server licences to Azure is the cost of Software Assurance, whereas AWS and Google must incur the SPLA price to compete for these workloads. It is therefore relevant to compare the price of Software Assurance with the SPLA price to assess the potential for the AHB to competitively disadvantage AWS and Google.

- 7.516 Once we have established whether there is the potential for there to be a price disadvantage on AWS and Google, we need to assess the significance of the price difference when considering different bundles of services purchased by customers that include the Microsoft software. As set out below, we do this by expressing the price difference as a proportion of customer spend on Windows Server VMs (and SQL Server IP in the case of SQL Server) on Azure, and customer spend across all Azure services.
- 7.517 We estimated the licensing input costs that Microsoft would incur to host Azure customers' Windows Server and SQL Server usage if it paid the same input prices that it charges AWS and Google. To do this, we multiplied each Azure customer's total 2024 vcore hours of usage of each product by the respective per-vcore hour price in each of AWS' and Google's SPLAs. We then calculated the difference between these estimated input costs to serve each Azure customer and each Azure customer's actual spend on Windows Server and SQL Server licensing IP. The difference between these figures serves as an indicator of the scale of the difference between Microsoft's customer-facing prices for each product and the input prices that it charges AWS and Google.
- 7.518 As such, our analysis effectively compares Azure's customer-facing prices (net of discounts) and the input prices that Microsoft charges AWS and Google, as both the IP spend figures and estimated input costs are multiples of the same volumes of usage.²³³⁷
- 7.519 We express these differences in SPLA input costs and licensing IP spend as (a) percentage differences, (b) a proportion of customer spend on Windows Server VMs (and SQL Server IP in the case of SQL Server) on Azure, and (c) customer spend across all Azure services. The percentage differences demonstrate the magnitude of the difference between the upstream and downstream prices. The differences as a proportion of each denominator provide an indication of the scale of the additional input costs that AWS and Google face when competing with Microsoft for customers that want to use each software product on public cloud, all else equal.
- 7.520 As set out above in 'Aggregation of workloads', it is appropriate to consider values within the range (between total Azure spend and an individual product or service), which are likely to be more reliable as indicators of the importance of the Microsoft software as an input. Most customers are not at the two extremes and therefore we have placed greater weight on values towards the middle of the range for the cost base.

²³³⁷ In other words, you could divide both the input cost and the IP spend by the vcore hours of usage to get the unit cost and unit price for each customer. As such, the percentage difference between these two figures equals the percentage difference between the SPLA input price and the effective customer-facing price that each customer pays per vcore hour.

- 7.521 We present the mean and median of each of (a), (b) and (c) for customers that use the relevant product in each revenue bracket to understand how the potential for the licensing practices to harm AWS' and Google's competitiveness may vary with respect to customers of different sizes.
- 7.522 We also present the median results to account for the fact that some revenue brackets encompass a broad range of heterogeneous customers. In particular, the lower brackets include a high concentration of customers at the bottom end of the spend distribution and a long tail of customers with significantly higher spend (which, accordingly, are more likely to use more expensive editions of SQL Server). This means that the mean results may be less representative of the majority of the customers within these brackets.
- 7.523 For Windows Server, due to the availability of data our analysis focuses on the difference in customer-facing and input prices relating to AHB usage only. appendix T sets out our full methodology including a discussion of the AHB IP spend data point.
- 7.524 For SQL Server, our analysis focuses on PAYG usage only as customers can bring their own SQL Server licences to AWS and Google. This means that AWS and Google would incur no input costs if they were to host Azure customers' SQL Server workloads that are eligible for the AHB on their clouds. That said, we note that marketing materials on Microsoft's website indicate that SQL Server may perform better and cost less when customers bring their own licences to Azure rather than AWS or GCP; however, it is unclear whether this is related to a licensing practice. If related to a licensing practice, such price and quality differences could potentially contribute to partial foreclosure of AWS and Google without impacting their input costs, but our analysis has not assessed that.²³³⁸
- 7.525 As discussed above, there are three editions of SQL Server that are available to license via AWS' and Google's SPLAs. Microsoft provided an estimate of its relative usage data for these editions on SQL PAYG by conducting a normalisation exercise of VM hours by number of cores ([X]).²³³⁹ We use these estimates in our analysis.
- 7.526 As explained in full in appendix T, we have excluded customers with <\$10,000 spend from this analysis. These customers account for a minority of total Azure spend. Therefore, our analysis still covers the majority of Azure spend.

²³³⁸ For example, the Azure website states that migrating SQL Server to Azure VMs rather than AWS EC2 instances is up to 23% cheaper and allows for up to 22% faster performance. See: [SQL Server on Azure Virtual Machines](#).

²³³⁹ [Microsoft's response to the CMA's PDR](#) dated 28 January 2025, page 68, footnote 52. Microsoft submitted this is not recorded in the ordinary course of business.

7.527 Google's SPLA with Microsoft includes [REDACTED]. We present the results using Google's prices below and present the full results in appendix T.

7.528 Finally, this analysis only considers usage of the Microsoft products as installed on VMs (ie IaaS usage). However, SQL Server PaaS services such as Azure SQL Database are among the most popular cloud services (as shown above). Therefore, this analysis may underestimate the difference in the prices that Microsoft charges Azure customers for SQL Server and the prices it charges AWS and Google, and particularly when compared to customers' total spend across all cloud services.

7.529 Table 7.10 below shows the difference between the input prices that Microsoft charges Google for Windows Server and its customer-facing prices relating to Windows Server AHB usage. These price differentials are presented as average percentage differences for customers in each revenue bracket.

Table 7.10: Average percentage difference between the input prices that Google pays for Windows Server and Microsoft's customer-facing prices

[REDACTED]

Source: CMA analysis of Microsoft's submission to the CMA [REDACTED] and Microsoft's response to Ofcom's information request [REDACTED]

7.530 Our analysis suggests that the input prices that Microsoft charges Google for Windows Server are at least [REDACTED]% higher than its customer-facing prices relating to Windows Server AHB usage.

7.531 Table 7.11 shows the median percentage difference between Microsoft's customer-facing prices and the input prices that Google pays for Windows Server for customers in each revenue bracket.

Table 7.11: Median percentage difference between the input prices that Google pays for Windows Server and Microsoft's customer-facing prices

[REDACTED]

Source: CMA analysis of Microsoft's submission to the CMA [REDACTED] and to Microsoft's response to Ofcom's information request [REDACTED].

7.532 The median percentage differences are generally lower than the means presented above. However, they still suggest that Microsoft charges Google a significantly higher price for Windows Server than its customer-facing prices relating to Windows Server AHB usage.

7.533 Table 7.12 below shows the difference between the input prices that Microsoft charges Google and Google for Windows Server and its customer-facing prices expressed as a proportion of total spend across all Azure services and spend on Windows Server VMs on Azure on average for customers in each revenue bracket.

Table 7.12: Average difference between the input prices that Google pays for Windows Server and Microsoft's customer-facing prices, as a proportion of total spend on Azure services and Windows Server VM spend on Azure

[REDACTED]

Source: CMA analysis of Microsoft's submission to the CMA [REDACTED] and Microsoft's response to Ofcom's information request [REDACTED].

- 7.534 Our analysis suggests that the difference between Microsoft's customer facing prices and the input prices it charges Google for Windows Server account for at least [REDACTED]% of customer spend across all Azure services and as much as [REDACTED] of spend on Windows Server VMs across all revenue brackets.
- 7.535 As discussed above, the additional costs that Google must face to match Microsoft's competitive offering to any single customer that wants to use Windows Server on public cloud and qualifies for the AHB falls within the range that we present for each revenue bracket. These results suggest that Google faces a disadvantage in competing for such customers, and this is particularly the case for customers that allocate their cloud spend in a less aggregated way.
- 7.536 Table 7.13 below shows the median difference between the input prices that Google pays for Windows Server and Microsoft's customer-facing prices expressed as a proportion of total Azure spend and spend on Windows Server VMs on Azure.

Table 7.13: Median difference between the input prices that Google pays for Windows Server and Microsoft's customer-facing prices, as a proportion of total spend on Azure services and Windows Server VM spend on Azure

[REDACTED]

Source: CMA analysis of Microsoft's submission to the CMA [REDACTED] and Microsoft's response to Ofcom's information [REDACTED]:

- 7.537 The median results are generally lower than the means presented above. However, they still suggest that Google faces a disadvantage in competing for customers that want to use Windows Server on public cloud and qualify for the AHB, and this is particularly the case for customers that aggregate their cloud spend in a less aggregated way.
- 7.538 Table 7.14 below shows the percentage difference between the input prices that Microsoft charges Google for SQL Server and its customer-facing prices relating to SQL Server PAYG usage on average for customers in each revenue bracket.

Table 7.14: Average percentage difference between the input prices that Google pays for SQL Server and Microsoft's customer-facing prices relating to SQL Server PAYG usage

[REDACTED]

Source: CMA analysis of Microsoft's submission to the CMA [REDACTED] and Microsoft's response to Ofcom's information request [REDACTED]

7.539 Our analysis suggests that the input prices are at least [REDACTED]% higher than Microsoft's customer-facing prices for customers.

7.540 Table 7.15 below shows the median percentage difference between the input prices that Google pays for Windows Server and Microsoft's customer-facing prices relating to SQL Server PAYG usage.

Table 7.15: Median percentage difference between the input prices that Google pays for SQL Server and Microsoft's customer-facing prices relating to SQL Server PAYG usage

[REDACTED]

Source: CMA analysis of Microsoft's submission to the CMA [REDACTED] and Microsoft's response to Ofcom's information request [REDACTED].

7.541 Table 7.16 below shows the difference between the input prices that Microsoft charges Google for SQL Server and its customer-facing prices relating to SQL Server PAYG usage expressed as a proportion of customer spend across all Azure services and spend on Windows Server VMs and SQL Server IP on average for SQL Server customers in each revenue bracket.

Table 7.16: Average difference between the input prices that Google pays for SQL Server and Microsoft's customer-facing prices relating to SQL Server PAYG usage, as a proportion of total Azure spend and spend on Windows Server VMs and SQL Server IP on Azure

[REDACTED]

Source: CMA analysis of Microsoft's submission to the CMA [REDACTED] and Microsoft's response to Ofcom's information request [REDACTED].

7.542 Our analysis suggests that the difference between Microsoft's customer facing prices and the input prices it charges Google for SQL Server account for at least [REDACTED] of customer spend across all Azure services and as much as [REDACTED]% of spend on Windows Server VMs and SQL Server IP across the four highest revenue brackets.

7.543 As discussed above, the additional costs that Google must face to match Microsoft's competitive offering to any single customer that wants to use SQL Server on a PAYG basis on public cloud falls within the range that we present for each revenue bracket. These results suggest that Google faces a disadvantage in competing for such customers, and this is particularly the case for customers that allocate their cloud spend in a less aggregated way.

7.544 Table 7.17 below shows the median difference between the input prices that Microsoft charges Google for SQL Server and its customer-facing prices relating to SQL Server PAYG usage expressed as a proportion of customer spend across all Azure services and spend on Windows Server VMs and SQL Server IP for SQL Server customers in each revenue bracket.

Table 7.17: Median difference between the input prices that Google pays for SQL Server and Microsoft’s customer-facing prices as a proportion of total Azure spend and spend on Windows Server VMs and SQL Server IP on Azure

[REDACTED]

Source: CMA analysis of Microsoft’s submission to the CMA [REDACTED] and Microsoft’s response to Ofcom’s information request [REDACTED].

- 7.545 The median results are generally lower than the means presented above. However, they also suggest that Google faces a disadvantage in competing for customers that want to use SQL Server on a PAYG basis on public cloud, and this is particularly the case for customers that allocate cloud spend in a less aggregated way.
- 7.546 Table 7.18 below shows the combined difference in the input prices that Microsoft charges Google for Windows Server and SQL Server and its customer-facing prices relating to Windows Server AHB and SQL Server PAYG usage expressed as a proportion of customer spend across all Azure services and spend on Windows Server VMs and SQL Server IP on average for SQL Server customers in the four highest revenue brackets.

Table 7.18: Average combined difference between the input prices that Google pays for Windows Server and SQL Server and Microsoft’s customer-facing prices relating to Windows Server AHB and SQL Server PAYG usage, as a proportion of total Azure spend and spend on Windows Server VMs on Azure

[REDACTED]

Source: CMA analysis of Microsoft’s submission to the CMA [REDACTED] and Microsoft’s response to Ofcom’s information request [REDACTED].

- 7.547 Our analysis suggests that the average difference between the input prices that Microsoft charges Google for Windows Server and SQL Server and its customer-facing prices relating to Windows Server AHB and SQL Server PAYG usage account for at least [REDACTED]% of spend across all Azure services and as much as [REDACTED]% of spend on Windows Server VMs and SQL Server IP across the revenue brackets.
- 7.548 As discussed above, the additional costs that Google must face to match Microsoft’s competitive offering to any single customer that wants to use Windows Server and SQL Server on a PAYG basis on public cloud falls within the range that we present for each revenue bracket. These results suggest that Google faces a disadvantage in competing for such customers, and this is particularly the case for customers that allocate their cloud spend in a less aggregated way.
- 7.549 Table 7.19 below shows the median combined difference in the input prices that Microsoft charges Google for Windows Server and SQL Server and its customer-facing prices relating to Windows Server AHB and SQL Server PAYG prices, as a

proportion of total spend on all Azure services and spend on Windows Server VMs and SQL Server IP on Azure.

Table 7.19: Median combined difference between the input prices that Google pays for Windows Server and SQL Server and Microsoft's customer-facing prices relating to Windows Server AHB and SQL Server PAYG usage, as a proportion of total Azure spend and spend on

[X]

Source: CMA analysis of Microsoft's submission to the CMA [X] and Microsoft's response to Ofcom's information request [X].

- 7.550 The median results are generally lower than the means presented above across most revenue brackets. However, they also suggest that Google faces a disadvantage in competing for customers that want to use Windows Server and SQL Server on a PAYG basis on public cloud, and this is particularly the case for customers that aggregate their cloud spend in a less aggregated way.
- 7.551 The magnitude of the percentage differences between Microsoft's customer-facing prices and the input prices that it charges AWS²³⁴⁰ and Google suggests that Microsoft supplies Windows Server and SQL Server to AWS and Google at a significantly higher cost than it incurs itself. As such, AWS and Google face a significant disadvantage in competing for customers' Windows Server and SQL Server workloads.
- 7.552 The differences between Microsoft's customer-facing prices and the input prices it charges AWS and Google as a proportion of customer spend provide an indication of the extent of the effect of Microsoft's licensing practices on competition. These results show that AWS and Google would face at least several percentage points of higher costs to compete for such customers, with these cost proportions increasing for customers that allocate their cloud spend in a less aggregated way.
- 7.553 While the difference in Microsoft's customer-facing prices and the input prices it charges AWS and Google for SQL Server as a proportion of total Azure spend is lower, SQL Server is predominantly used in conjunction with Windows Server²³⁴¹ (and exclusively on GCP) such that it is appropriate to consider the combined impact of the licensing practices relating to Windows Server and SQL Server.
- 7.554 Additionally, as noted above, Microsoft's marketing material indicates that there may be price and non-price differences between BYOL SQL Server usage on Azure compared with AWS and GCP. To the extent that these are related to licensing practices or restrictions, this could contribute to partial foreclosure

²³⁴⁰ See appendix T for the results of our comparison of Microsoft's customer-facing prices and the input prices it charges AWS for Windows Server and SQL Server.

We conducted the analysis for AWS and Google separately as they have separate SPLAs with Microsoft. Nevertheless, we note that the results are similar for both.

²³⁴¹ In particular, [80-90]% of SQL Server PAYG IaaS usage on Azure occurs on Windows Server VMs. This figure was calculated using the usage data for SQL Server segmented by underlying OS in Microsoft's response to the CMA's information request [X].

without impacting AWS' or Google's input costs, although our analysis has not assessed that.

- 7.555 Finally, our analysis does not consider SQL PaaS usage. As set out above, Azure's primary SQL PaaS service, SQL Database, accounts for [0-10]% of Azure's total revenues. As such, our analysis likely understates the significance of the licensing practices relative to customer spend across all cloud services.

A cloud provider's analysis of Windows Server licensing cost versus infrastructure price

- 7.556 A cloud provider submitted data analysis that calculated the difference between the amount customers pay to use Windows Server licences on Azure and the input cost to that cloud provider to host those workloads.²³⁴² In essence, the analysis calculates the amount that the cloud provider would have to absorb to match the price Azure charges. The analysis considered different instances on Azure.
- 7.557 The cloud provider submitted that the analysis demonstrates that an as-efficient competitor would not be able to match the retail costs Microsoft charges customers with the AHB for the use of its software on Azure.²³⁴³ [REDACTED].²³⁴⁴
- 7.558 The analysis uses a similar numerator to the one we calculated in our analysis above. In particular, the numerator is the estimated difference in the licence price that the cloud provider pays to Microsoft to host Windows Server workloads and the price that customers would pay to use Windows Server licences on Azure (ie the cost of Software Assurance).²³⁴⁵ The cloud provider uses the price of infrastructure for that workload on Azure as the denominator.²³⁴⁶
- 7.559 The table below shows the average difference in licensing related costs as a proportion of Azure infrastructure prices.

Table 7.20: The cloud provider's analysis of differences in licensing costs paid by that cloud provider compared to customer facing price for Azure customers

[REDACTED]

Source: CMA analysis of [REDACTED] submission to the CMA [REDACTED].

- 7.560 The table shows, for example, that the average difference in price that the cloud provider pays to Microsoft to host Windows Server workloads and the price

²³⁴² [REDACTED] submission to the CMA [REDACTED].

²³⁴³ [REDACTED] submission to the CMA [REDACTED].

²³⁴⁴ [REDACTED] submission to the CMA [REDACTED].

²³⁴⁵ [REDACTED].

²³⁴⁶ [REDACTED] submission to the CMA [REDACTED].

customers would pay to use Windows Server on Azure accounts for [REDACTED] of Azure's infrastructure price under PAYG plans on general purpose instances.

7.561 The cloud provider submitted that the results show that the difference in licence-related cost is greater than [REDACTED]²³⁴⁷ [REDACTED].²³⁴⁸

7.562 The cloud provider said that it would therefore need to offer considerable discounts to match Azure's price. It added that, in fact, when the difference in the cost of Windows Server licence is greater than 100% of the Azure infrastructure price customers would have to pay more on the provider's cloud than Azure even if it priced the infrastructure at zero.²³⁴⁹

7.563 We note that:

- (a) The analysis compares the price that the cloud provider would pay to host Windows Server and the price customers pay to use Windows Server licences if they qualify for AHB. As such, the analysis is only relevant for customers with pre-existing licences for Windows Server with Software Assurance or with a qualifying subscription licence;
- (b) The denominator used (ie the Azure infrastructure price for that instance) is narrow and we consider that most customers likely make their purchasing decisions considering a wider range of workloads. See above for a discussion on the appropriate denominator in this context. To the extent that customers include other workloads in their purchasing decision, it may be worthwhile for the cloud provider to absorb losses (or low margins) on Windows Server workloads in order to win the customer's other workloads.
- (c) [REDACTED].²³⁵⁰ [REDACTED].
- (d) We do not know the relative usage of each instance on Azure. To the extent that some instances are used more frequently by customers, it may be appropriate to place greater weight on the difference in licensing cost calculated for those instances.²³⁵¹ Nevertheless, we note that even if the most used instances are on the lower end of the estimates, the difference in licensing costs still accounts for a material portion of infrastructure prices.

7.564 We consider that the cloud provider's analysis suggests that the difference in licensing cost that it pays to host Windows Server workloads and the price customers that qualify for AHB would pay to use Windows Server on Azure is a material proportion of Azure's infrastructure price for those instances. While the

²³⁴⁷ [REDACTED]. Submissions to the CMA [REDACTED].

²³⁴⁸ [REDACTED]. Submissions to the CMA [REDACTED].

²³⁴⁹ [REDACTED] submission to the CMA [REDACTED].

²³⁵⁰ [REDACTED] response to the CMA's information request [REDACTED].

²³⁵¹ The analysis currently weights each instance equally.

point estimates may be an overestimate, we consider that adjusting for [✂] and the relative usage of instances likely will not have a significant impact on our interpretation of the results.²³⁵² Further, while the denominator used is narrow, we consider that the results are consistent with the analysis we conducted, which uses a wider denominator.

List prices

- 7.565 Microsoft submitted a comparison of list prices for selected instances involving the use of Microsoft software across AWS, Azure and GCP. These are Windows VM PAYG prices which customers pay if they do not have a pre-existing Windows Server licence that qualifies for AHB. Therefore, this analysis is informative of the cloud providers' competitive offerings for customers without any qualifying licences.²³⁵³
- 7.566 For the instances provided by Microsoft, GCP has higher list prices than Azure by 5% to 93% for all but one of the relevant instances. We consider that general purpose instances (where GCP has a 5% higher list price) are likely the most common instance. For the same instances, AWS has equal or lower list prices than Azure's list prices.
- 7.567 We consider that the list prices analysis submitted by Microsoft indicates that for customers that do not have pre-existing licences that qualify for AHB on Azure, AWS' list prices are similar or better than Microsoft's, whereas Google's are at least marginally worse.

Price and non-price differences for the client-side products

- 7.568 In this section we consider whether Microsoft's licensing practices in relation to the client-side products impact the quality of AWS, Google's and other cloud providers' competitive offerings.
- 7.569 We consider the price and non-price differences for the client-side products and assess whether these differences may worsen AWS, Google's and other cloud providers' competitive offerings. The client-side products are those usually accessed through the desktop, for which the use case on the cloud is for use through a virtual desktop. These products are Visual Studio, Microsoft 365, Office and Windows 10/11.

²³⁵² In fact, adjusting for relative usage may either reduce or increase the point estimates.

²³⁵³ In relation to the difference in the customer facing price for customers with licences that qualify for AHB, Microsoft's website sets out that AHB results in cost savings including the licence for Windows Server, meaning the customer only needs to pay for the base compute rate of the VM, which Microsoft specifies is equal to the Linux rate for VMs. Source: [Azure Hybrid Benefit for Windows Server | Microsoft Learn](#), accessed on 06 June 2025

We note that a customer must have an active Software Assurance or qualifying subscription licence to benefit from AHB.

7.570 Table 7.21 outlines the licensing practices with respect to the client-side products.

Table 7.21 Availability of the Microsoft products for use in the public cloud

	Office desktop applications	Microsoft 365	Windows 10/11	Visual Studio
Microsoft	Can BYOL with software assurance	Microsoft 365 applications available to run in Azure. ²³⁵⁴	Can BYOL with software assurance	Can BYOL
AWS	Available under the AWS SPLA for use with EC2, Amazon Appstream and 2.0 and Amazon Workspaces. ²³⁵⁵ Cannot BYOL. ²³⁵⁶ Currently available under third party SPLA, [REDACTED]. ²³⁵⁷	Not available through the SPLA. ²³⁵⁸ BYOL to Amazon Workspaces for a subset of Microsoft 365 packages only. ²³⁵⁹ Not available on Amazon Appstream. ²³⁶⁰	Not available through the SPLA. ²³⁶¹ Cannot BYOL to the public cloud.* Currently made available for resale through third parties to host on Listed Providers infrastructure [REDACTED]. ²³⁶²	Available via the AWS SPLA for use with EC2. ²³⁶³ Limited ability to BYOL: Customers cannot BYOL the [REDACTED] of Visual Studio to AWS Workspaces. ²³⁶⁴
Google	Technically 'available' under the SPLA [REDACTED]. ²³⁶⁵ BYOL unavailable to the public cloud for licenses post-2019. ²³⁶⁶ Currently available under third party SPLA to use on GCP, [REDACTED].	Not available via BYOL or SPLA. ²³⁶⁷ Microsoft has commented that Google offers its own Workspaces solution, so it does not need to make this available to Google. ²³⁶⁸	Not available through SPLA. ²³⁶⁹ Cannot BYOL to the public cloud.* Currently made available for resale through third parties to host on Listed Providers infrastructure [REDACTED]. ²³⁷⁰	Available through the SPLA [REDACTED]. ²³⁷¹ Customers cannot BYOL. ²³⁷²
Non-Listed Providers	Can BYOL with Software Assurance.	Microsoft 365 Applications available to run in Non-Listed Provider clouds. ²³⁷³	Can BYOL with Software Assurance.	Can BYOL with Software Assurance.

*We understand that customers can BYOL for Windows 10/11 on dedicated hardware (private cloud) with the virtual desktop access (VDA) licences per user purchased under subscription from Microsoft.²³⁷⁴

²³⁵⁴ For access to the M365 apps on VDI on Azure you need to have a package which includes the [relevant licenses](#) (which are included in popular Microsoft 365 packages, including the business premium plan, and Microsoft E3/5 packages. Microsoft's submission to the CMA [REDACTED].

²³⁵⁵ [REDACTED] response to the CMA's information request [REDACTED].

²³⁵⁶ [REDACTED] response to the CMA's information request [REDACTED].

²³⁵⁷ [REDACTED] submission to the CMA, [REDACTED].

²³⁵⁸ [REDACTED] submission to the CMA [REDACTED].

²³⁵⁹ E3/E5, A3/A5, G3/G5, and Business Premium licenses only. Applications included may vary by package. [Migrating to Microsoft 365 Apps for enterprise at scale on Amazon WorkSpaces](#) accessed on 15 July 2024. AWS can host a customer's Microsoft 365 licence on shared hardware, but only if (a) AWS sells the customer a Windows Server licence under its SPLA, and (b) the customer brings its own existing Microsoft 365 licence. AWS submission to the CMA, [REDACTED].

²³⁶⁰ Note of meeting with [REDACTED].

²³⁶¹ Microsoft's response to the CMA's information request [REDACTED].

²³⁶² [REDACTED] submission to the CMA [REDACTED].

²³⁶³ [Microsoft licensing on AWS - AWS Prescriptive Guidance](#), accessed on 04 October 2024.

²³⁶⁴ [REDACTED] submission to the CMA [REDACTED].

²³⁶⁵ [REDACTED]. [REDACTED] submission to the CMA [REDACTED].

²³⁶⁶ [REDACTED] response to the CMA's information request [REDACTED].

²³⁶⁷ [REDACTED] submission to the CMA [REDACTED].

²³⁶⁸ Microsoft submission to the CMA [REDACTED].

²³⁶⁹ [REDACTED] submission to the CMA [REDACTED]; Microsoft's response to CMA information request [REDACTED].

²³⁷⁰ [REDACTED] submission to the CMA [REDACTED].

²³⁷¹ [REDACTED] submission to the CMA [REDACTED].

²³⁷² [REDACTED] submission to the CMA [REDACTED].

²³⁷³ Microsoft's submission to the CMA [REDACTED].

²³⁷⁴ [Bring Your Own Windows 10/11 licenses in WorkSpaces Personal - Amazon WorkSpaces](#), accessed on 15 July 2024; [Microsoft FAQ – AWS](#), accessed on 15 July 2024.

7.571 Table 7.21 shows that:

- (a) Google and AWS cannot resell Microsoft 365 or Windows 10/11 through their SPLAs.
 - (i) Google submitted that [REDACTED].²³⁷⁵
 - (ii) [REDACTED].²³⁷⁶ [REDACTED].
- (b) For Windows 10/11, customers cannot BYOL to the public cloud on AWS and GCP but can BYOL to the private cloud or 'dedicated infrastructure' with the VDA (Virtual Desktop Access) per user fee.
 - (i) One provider submitted that the VDA fee can increase customers costs by 30% to 110%.²³⁷⁷
 - (ii) AWS submitted that it is significantly more expensive and inefficient for the customer and cloud services provider to allocate dedicated hardware to an individual customer rather than using shared hardware, because shared hardware facilitates pooled resources that can be dynamically allocated among many customers as their respective individual needs fluctuate and scale over time.²³⁷⁸

7.572 Microsoft told us that it is a misconception that customers cannot run Microsoft 365 on Listed Provider clouds. It said that while it typically only makes its traditional Microsoft Office applications available for use on other clouds and not its cloud-backed Microsoft 365 Apps, this has changed for AWS which can now make Microsoft 365 Apps available on AWS Workspaces.²³⁷⁹ Microsoft also submitted that Google has its own popular productivity suite, Google Workspace, that are only available in Google's cloud, and so there is no rationale to force Microsoft to offer its competing product to them.²³⁸⁰ In any event, Microsoft said that Microsoft 365 is designed for use on personal computers, and while customers can run it in the cloud, it's relatively uncommon to do so compared to running them on a Windows PC or an Apple Mac.²³⁸¹

7.573 Microsoft said that its Microsoft 365 Apps are cloud-backed applications that were created to be part of Microsoft's Office 365 and Microsoft 365 SaaS solution, and

²³⁷⁵ [REDACTED]. Google's submission to the CMA, [REDACTED].

²³⁷⁶ [REDACTED]. [REDACTED] response to the CMA's information request [REDACTED].

²³⁷⁷ [REDACTED]. [REDACTED] submission to the CMA [REDACTED].

²³⁷⁸ The VDA fee does not apply to customers wanting to run Windows 10/11 on Azure or non-Listed Providers' private clouds. [REDACTED] submission to the CMA [REDACTED].

²³⁷⁹ Microsoft's submission to the CMA [REDACTED]. Since 1 August 2023, AWS customers have been able to BYOL their existing Microsoft 365 licences to AWS Workspaces. See: [Microsoft 365 Apps for enterprise now available on Amazon WorkSpaces services](#); Microsoft said that this was a concession made during its negotiation with AWS whilst the two were negotiating an agreement for Amazon to continue to be an end user of Microsoft 365. See, Microsoft's response to the CMA's information request [REDACTED].

²³⁸⁰ Microsoft's submission to the CMA [REDACTED].

²³⁸¹ Microsoft's submission to the CMA [REDACTED].

were designed to connect to Microsoft's cloud services, like OneDrive and SharePoint Online to enable new functionality like co-authoring that is only possible when connected to an integrated cloud service.²³⁸² Microsoft said that Microsoft Office and its associated server products are available in SPLA, but that Microsoft 365 Apps and its associated services – Exchange Online and SharePoint Online – are not, as it would not be possible to run these other cloud services in another cloud.²³⁸³

- 7.574 Microsoft submitted that customers can purchase subscriptions to Microsoft 365 and Windows 10/11s and run those applications in any non-Listed Provider cloud as well as Azure.²³⁸⁴ It said that customers that want to use a Windows based DaaS solution on Google's cloud can do so by bringing Windows Enterprise VDA licences to GCP and then buying Office from Google via Google's SPLA.²³⁸⁵ Customers that want to use a Windows-based DaaS solution from AWS can now bring the Microsoft 365 Apps to AWS WorkSpaces directly.²³⁸⁶
- 7.575 Google told us that Microsoft describes 365 as a SaaS product, but it also includes the client-side software which can be downloaded onto a machine. Google said that it has heard anecdotally that users prefer the local copies of the Microsoft Apps as the browser-based versions don't have feature parity, and don't provide as robust real time collaboration. It said that there should be no technical reason why a customer cannot bring the client-side versions of 365 to a VDI solution on third-party infrastructure.
- 7.576 [REDACTED]. But customers are still prohibited from bringing their own pre-purchased licences of that client-side software.²³⁸⁷
- 7.577 Microsoft has submitted that it was a misconception that Windows 10 and 11 are not available via SPLA and that this limits competition because it is important for DaaS.²³⁸⁸ We understand that Microsoft's submission is that the misconception relates to the fact that non-availability of these products via the SPLA limits competition because it is important for DaaS, rather than it being a misconception that these products are not available via the SPLA. We discuss this below. Microsoft noted that Windows 10/11 have never been available via the SPLA.²³⁸⁹
- 7.578 In response to the PDR, Microsoft submitted that Windows 10/11, Microsoft productivity suites and Visual Studio are SaaS offerings that do not impact the choice of IaaS provider. It gave the example of Gmail or Google Workspaces,

²³⁸² Microsoft's submission to the CMA [REDACTED].

²³⁸³ Microsoft's submission to the CMA [REDACTED].

²³⁸⁴ Microsoft's submission to the CMA [REDACTED].

²³⁸⁵ Microsoft's submission to the CMA [REDACTED].

²³⁸⁶ Microsoft's submission to the CMA [REDACTED].

²³⁸⁷ [REDACTED]. Note of meeting with [REDACTED].

²³⁸⁸ Microsoft's submission to the CMA [REDACTED].

²³⁸⁹ Microsoft's response to the CMA's information request [REDACTED].

which could be used to provide a user with a desktop. It submitted that the same is true of developer tools which do not in any way prohibit and distort customers' choice of IaaS provider.²³⁹⁰

7.579 We do not think it is correct to label these products only as SaaS offerings. We explain above that Windows 10/11 is a desktop operating system, Microsoft's productivity suites refers to both the desktop and web versions of Microsoft's popular productivity apps – some of which have a SaaS element – or are 'cloud backed', and that Visual Studio IDE is distinct from Visual Studio code – which may be delivered as SaaS. These products can either be installed on a physical desktop (on the individual computer) or a virtual desktop (installing the software on a virtual machine). Microsoft itself has two VDI offerings, one 'Windows 365' is delivered as SaaS, the other 'Azure Virtual Desktop', is delivered as PaaS.²³⁹¹ We have presented evidence above in this chapter on the importance of VDI workloads to Microsoft's cloud business. Our analysis does not include SaaS revenues.

7.580 As we discuss above, Microsoft products are important to offer as part of a comprehensive VDI solution. In view of the licensing practices outlined above, we consider AWS' and Google's competitive offering is adversely affected when offering VDI workloads. As outlined in this chapter, we have evidence which suggest that these restrictions impact choice of cloud provider for VDI workloads.

7.581 We have seen that [redacted] and the standard SPLAs [redacted].²³⁹²

Relevant price and non-price restrictions

7.582 Below, we consider the price and non-price restrictions for each of the client-side products.

Windows 10/11 multi-session

7.583 Customers are unable to use [redacted] to run Windows 10/11 in multisession mode. One cloud provider [redacted] submitted that this is only available for Azure customers and allows customers to simultaneously place multiple users on the same virtual instance, therefore lowering costs.²³⁹³

7.584 [redacted].²³⁹⁴

7.585 We have also received a submission from a small cloud provider which details its inability to offer this service. It said that giving every user its own Windows 11

²³⁹⁰ Microsoft's response to the PDR, paragraph 260.

²³⁹¹ Microsoft, [Azure Virtual Desktop | Microsoft Azure](#), accessed on 03 June 2025.

²³⁹² [redacted] response to the CMA's information request [redacted]. [redacted] response to Ofcom's information request [redacted].

²³⁹³ [redacted] submission to the CMA [redacted].

²³⁹⁴ [redacted] submission to the CMA [redacted].

virtual machine requires much more system resources such as RAM and CPU, as well as more disk space, and more cost to set up, and more cost to support, and that this makes its service more expensive compared to Microsoft's VDI service.²³⁹⁵

Microsoft 365 and Teams restrictions

- 7.586 AWS submitted that customers can BYOL Microsoft 365 on shared hardware on WorkSpaces, only if AWS sells the customer a Windows Server licence under its SPLA and the customer brings its own existing Microsoft 365 licence. This means that customers that already own a Windows Server licence must repurchase this [REDACTED], whereas this does not apply on Azure. Customers can also BYOL Windows Server to use with Microsoft 365 but are required to use 'dedicated hardware' (private cloud),²³⁹⁶ and therefore, cannot BYOL to the public cloud on this mode of deployment. [REDACTED].²³⁹⁷
- 7.587 [REDACTED] submitted that, based on its analysis, [REDACTED].²³⁹⁸
- 7.588 A small cloud provider also submitted that accessing the Microsoft 365 applications such as Word, Excel, Outlook, Access and PowerPoint through VDI facilitated by Windows Server on its cloud required some small businesses to upgrade their Microsoft 365 licence (to a more expensive licence which includes shared computer activation), whereas this was not the case if they want to use Azure Virtual Desktop, which can run Windows 10/11 OS.²³⁹⁹ It has submitted that this means the end customers total cost of ownership is higher when using Microsoft's competitors systems than its own.

7.589 [REDACTED].²⁴⁰⁰

Visual Studio subscriber software and restrictions

- 7.590 We understand that Microsoft sells some Visual Studio subscriber software products which provide developers with access to Microsoft software within Visual Studio to test and validate their applications as well as access to Visual Studio IDE.²⁴⁰¹ One cloud provider has submitted that Visual Studio subscriber software is only available on Azure and that it effectively offers unlimited non-production use of Windows Server on Azure for development and testing purposes.²⁴⁰²

²³⁹⁵ [REDACTED] submission to the CMA [REDACTED].

²³⁹⁶ [REDACTED] submission to the CMA [REDACTED].

²³⁹⁷ Note of meeting with AWS [REDACTED].

²³⁹⁸ [REDACTED] submission to the CMA [REDACTED]; [REDACTED] submission to the CMA [REDACTED].

²³⁹⁹ [REDACTED] submission to the CMA [REDACTED].

²⁴⁰⁰ [REDACTED] submission to the CMA [REDACTED].

²⁴⁰¹ [Visual Studio Subscriptions \(Formerly MSDN Licenses\) - Everything You Need To Know](#), accessed on 12 November 2024; [Available software downloads - Visual Studio Subscription](#), accessed on 12 November 2024; [REDACTED].

²⁴⁰² [REDACTED] response to Ofcom's information request [REDACTED].

- 7.591 We understand that the cloud provider is suggesting that customers who want to develop and test applications which run on Windows Server using Visual Studio on VDI on non-Azure clouds must also purchase additional Windows Server licences for development and testing through the SPLA in addition to the ('standalone') Visual Studio IDE software for use on Listed Provider clouds. In contrast, on Azure, customers can purchase these as a package.
- 7.592 While anecdotal, [redacted] licensing costs with respect to Visual Studio on [redacted] compared to on Azure. One example estimated that the cost of the development and testing licences on [redacted] this cloud provider was greater than the monthly licensing cost for Visual Studio IDE licences alone.²⁴⁰³ We understand this cloud provider's [redacted] argument to be that the need to purchase additional development and testing licenses on non-Azure clouds increases costs [redacted] above any differences in the licensing cost of Visual Studio IDE on Azure compared to non-Azure clouds.
- 7.593 AWS submitted that the [redacted] of Visual Studio continues to be restricted for customers accessing the software through BYOL on Amazon Workspaces.²⁴⁰⁴

Potential alternative solutions for Listed Providers

- 7.594 We understand that Windows Server can be used to create a Windows VDI environment, as well as Windows 10/11.²⁴⁰⁵ However, customers are subject to the licensing restrictions concerning Windows Server where they use Listed Providers. This means customers must pay SPLA input prices for Windows Server and customers cannot BYOL.²⁴⁰⁶ In addition, [redacted] submitted that Microsoft has announced end of life support for Microsoft Office on Windows Server, with all instances of Windows Server no longer supporting Microsoft Office as of 2025 anywhere but on Azure.²⁴⁰⁷ Therefore, we understand that post 2025, customers who base their VDI network on Windows Server not Windows 10/11 will be unable to use it to host Microsoft Office.
- 7.595 Further, while independent managed service providers can resell Windows 10/11 hosted on Listed Providers' cloud infrastructure, Microsoft has announced that this route will no longer be made available to customers on Listed Providers' cloud infrastructure from September 2025 onwards.²⁴⁰⁸

²⁴⁰³ [redacted] response to the CMA's information request [redacted].

²⁴⁰⁴ AWS' submission to the CMA [redacted].

²⁴⁰⁵ [redacted] response to the CMA's information request [redacted].

²⁴⁰⁶ Please refer above for a more detailed explanation of the Windows server licensing conditions.

²⁴⁰⁷ [redacted] submission to the CMA [redacted].

²⁴⁰⁸ See Description of the software licensing practices above outlining the 2025 changes [New licensing benefits make bringing workloads and licenses to partners' clouds easier](#), accessed on 4 October 2024.

Our assessment

- 7.596 We have compared the price that Microsoft charges Azure customers for Windows Server and SQL Server with the input prices it charges AWS and Google, and an analysis of Windows Server licensing cost compared to the infrastructure price. Both analyses address a similar question: how the input cost paid by AWS and Google compares to the customer facing price of Microsoft software on Azure.
- 7.597 While the analyses use slightly different methodologies, they both find that the input costs to AWS and Google are higher than Microsoft's customer facing price.
- 7.598 The fact that Microsoft has set AWS and Google's input costs higher than its own customer facing price in itself is evidence of the significance of the conduct and the potential for it to disadvantage them.
- 7.599 We also note that the difference in customer-facing prices and AWS and Google's SPLA input costs for Windows Server, and Windows Server and SQL Server combined, is a material portion of Windows Server VM spend and a smaller, but not insignificant, portion of those customers' total spend.
- 7.600 We consider that the analysis submitted by a cloud provider is consistent with our own and suggests that there is a significant disparity between the input cost paid by the cloud provider and the customer facing price of Microsoft software on Azure for customers with Windows Server licences that qualify for AHB.
- 7.601 We have found that list prices for select instances of Windows Server VMs are generally slightly higher on GCP than Azure and similarly priced across AWS and Azure. List prices are informative of the cloud providers' prices for customers without any pre-existing licences that qualify for AHB; it is not informative of prices for customers with qualifying licences.
- 7.602 We have also found that for Windows 10/11, Microsoft 365, Office and Visual Studio, the licensing practices mean that for customers with existing licences, they cannot bring these to AWS and Google's public cloud (except specific Microsoft 365 licences to Amazon Workspaces) and for customers without existing licences that they are [✂], nor can they purchase Windows 10/11 and Microsoft 365 on either AWS or Google.
- 7.603 Potential workarounds for AWS and Google wishing to offer Microsoft software for access via VDI are more costly compared to accessing them via VDI on Azure, and some routes are going to be closed in the future due to licensing changes.

Conclusions on price and non-price differences

- 7.604 There is a significant difference between Microsoft's customer facing prices on Azure and the prices it charges AWS and Google on their SPLAs. This, together

with the difference in other price and non-price factors is likely to harm competition.

The impact on AWS and Google's competitiveness from Microsoft's licensing practices

7.605 In the remainder of this chapter, we set out the evidence on the extent to which observed outcomes are consistent or inconsistent with Microsoft's licensing practices having had an impact on AWS and Google's competitiveness, resulting in an adverse effect on competition in the cloud services markets. We consider:

- (a) AWS' and Google's submissions on their pass-through of input costs to customers;
- (b) profit margins and mark-ups;
- (c) relative usage of Microsoft software products on Azure compared to AWS;
- (d) customer evidence on the reasons for their choice of cloud provider; and
- (e) examples of lost or missed customers from AWS and Google.

AWS' and Google's submissions on pass-through of input costs

7.606 In general, we would expect that a profit-maximising firm that experiences higher input costs will pass through at least some of these higher costs to customers.²⁴⁰⁹ In this section we set out AWS' and Google' submissions on how Microsoft's SPLA licensing terms have affected their competitive offerings. They have each taken a different approach and we set out both of these below.

7.607 One of them submitted that it faced significant price increases over the SPLA for Windows Server and that, while it initially tried to absorb the additional cost within its business, it ultimately had no choice other than to raise prices for its own customers.²⁴¹⁰ It submitted that it would incur substantial losses if it served certain customers on the same pricing terms as Microsoft, and it is therefore unable to offset Microsoft's significant input costs.²⁴¹¹

7.608 It also submitted that, in any case, if a competitor could profitably absorb those input costs, it would not preclude a finding that there had been a weakening of competition.²⁴¹² It submitted that it would be incoherent to use elements of an as-

²⁴⁰⁹ See our reasoning in Margins and mark ups section below. The theoretical models and empirical evidence considered in the cost pass-through report prepared for the OFT generally indicated that the exact degree of pass through depends on a range of factors. [RBB Economics -cost pass through: theory, measurement and potential policy implications](#).

²⁴¹⁰ [] response to the CMA's information request [].

²⁴¹¹ [] submission to the CMA [].

²⁴¹² [] submission to the CMA [].

efficient competitor test, which may consider whether an equally efficient competitor can profitably absorb Microsoft's SPLA input costs, for partial foreclosure. It stated that a raising rivals' cost assessment is concerned with partial foreclosure, and tests the materiality of those wholesale charges relative to the overall cost of supplying the downstream product to end customers.²⁴¹³

7.609 Evidence on the extent to which the other cloud provider has passed through high input costs on Microsoft software is mixed:

- (a) It submitted it is forced to pass on some of the higher licensing costs to customers.²⁴¹⁴ It [REDACTED]. It said that the difference in price relates to the cost of Windows Server licence acquired through the SPLA which it passes through to its customers.²⁴¹⁵ It also said absorbing the costs was not viable in a lot of cases, [REDACTED].²⁴¹⁶
- (b) It submitted that Microsoft had been raising Windows Server licensing fees over the past several SPLA renewals and that it had historically absorbed these cost increases and not raised prices for its customers.²⁴¹⁷
- (c) It submitted that it had passed on cost benefits for existing services to customers in the form of lower prices for these services, including for [REDACTED].²⁴¹⁸ It said that, in this sense, there may be indirect pass on (as it may have decreased prices if it had lower input costs).²⁴¹⁹ We consider that its evidence and analysis on passing through cost decreases may have limitations.²⁴²⁰

7.610 Microsoft submitted that, while AWS could have easily matched AHB discounts, it elected not to and continued to sell Windows Sever VMs at a significant premium as compared to its free Linux VM offering.^{2421,2422}

7.611 Microsoft submitted that AWS and Google can and do offer competing discounts on Windows Server or non-Windows Server workloads. In this regard, Microsoft said that AWS offers a discount of up to \$200 per Windows Server for migrations

²⁴¹³ [REDACTED] submission to the CMA [REDACTED].

²⁴¹⁴ [REDACTED] submission to the CMA, [REDACTED].

²⁴¹⁵ [REDACTED] submission to the CMA, [REDACTED].

²⁴¹⁶ Transcript of hearing with [REDACTED].

²⁴¹⁷ [REDACTED] response to Ofcom's information request [REDACTED].

²⁴¹⁸ [REDACTED] submission to the CMA [REDACTED].

²⁴¹⁹ Transcript of hearing with [REDACTED].

²⁴²⁰ [REDACTED].

²⁴²¹ Microsoft's submission to the CMA [REDACTED].

²⁴²² In comparison, AWS submitted that even if AWS were to fully absorb Microsoft licensing costs, such as buying an additional copy of a pre-existing licence for customers, which would be 'an economic absurdity', customers still must purchase VDA add-on licenses that add costs for them on top of the additional and unnecessary license which AWS would theoretically fully fund. AWS' submission to the CMA [REDACTED].

of at least 40 servers and \$250 for migrations of more than 80 Windows Servers.²⁴²³

- 7.612 The submissions by the first cloud provider, as set out above, are consistent with high costs leading to it charging a higher price and having a less competitive offering. This evidence has also shown a consistent picture over time, with changes in input prices being associated with changes in customer-facing prices.
- 7.613 In relation to the second cloud provider, there is some evidence that the input prices that Microsoft charges it appear to be associated with relatively high customer-facing prices, as set out above. However, the cloud provider has historically not passed through changes in input costs over time.
- 7.614 We consider this evidence on pass-through alongside the other evidence as part of our assessment.

Profit margins and mark-ups

- 7.615 In this section, we consider submissions by Microsoft, AWS and Google on the profit margins that are available (or would be available) to Google and AWS when supplying customers that use Windows Server and SQL Server on the cloud. Each of these providers:
- (a) made submissions relating to how we should consider margins in our assessment, and in particular what margins tell us about whether AWS and Google may be partially foreclosed when competing to supply some customers of cloud services that use Windows Server and/or SQL Server; and
 - (b) submitted their own estimates of what they considered to be the relevant margin or mark-up to consider.
- 7.616 Below, we first consider the submissions on margins as evidence of foreclosure in principle, before considering the specific estimates of margins submitted by Microsoft, AWS and Google.

Profit margins as evidence of partial foreclosure

Cloud provider submissions

- 7.617 Microsoft submitted that the critical issue is whether Microsoft's prices are materially too high or rivals' margins too low. Microsoft further submitted that its analysis of the margins it would earn if it were subject to the SPLA licensing costs

²⁴²³ Microsoft's submission to the CMA [X].

is critical to our consideration of an AEC because it provides clear evidence that competitors can effectively compete across the market. It submitted that this is because the analysis shows that the margins Google and AWS could earn are so high that, on any reasonable threshold, they provide evidence of there being no AEC.²⁴²⁴ Microsoft also submitted a separate piece of analysis which estimated the proportion of customers for which it would make positive gross margins.²⁴²⁵

- 7.618 Microsoft further submitted that the CMA has to find a benchmark for the foreclosure-free counterfactual, which it specified is the critical margin thresholds at which Google will compete effectively for Microsoft workloads.²⁴²⁶
- 7.619 AWS and Google submitted that margins are not determinative of whether there is an AEC.²⁴²⁷ AWS submitted that Microsoft cannot simply dismiss concerns that relate to partial foreclosure by arguing that AWS should absorb all the additional costs.²⁴²⁸

Our assessment

- 7.620 We consider that a material increase in input costs that is not common across all competitors is likely to lead to a material competitive disadvantage for those competitors that are affected, as long as the input is sufficiently important.
- 7.621 Rivals' margins can be relevant evidence of partial foreclosure. However, they do not provide determinative evidence and should be interpreted in the round with other evidence. We set out our reasoning below.
- 7.622 If a rival that incurs a high input cost cannot make a profit supplying the relevant product (ie their margin is negative), this would be evidence in support of foreclosure. However, it does not follow that any positive margin would be inconsistent with foreclosure or that there is a particular margin threshold associated with foreclosure:
- (a) As set out in our guidance, the relevant test for partial foreclosure in the context of an AEC assessment is whether a rival is materially disadvantaged and consequently competes less effectively.²⁴²⁹ In this context, we consider that a firm may still face a material disadvantage despite being able to earn positive margins.
 - (b) To illustrate this, we found it helpful to consider the example of a rival that is being foreclosed by high input costs. These costs can reduce the rival's

²⁴²⁴ [REDACTED] submission to the CMA [REDACTED].

²⁴²⁵ Microsoft's submission to the CMA [REDACTED].

²⁴²⁶ [Microsoft's response to the CMA's PDR dated 28 January 2025](#), paragraph 181.

²⁴²⁷ AWS' submission to the CMA [REDACTED]; Google's submission to the CMA [REDACTED].

²⁴²⁸ AWS' submission to the CMA [REDACTED].

²⁴²⁹ [CC3 \(Revised\)](#), paragraph 269.

overall profits in two ways. First, if the rival passes through some of the high cost, they may expect to lose some customers. Second, if the rival absorbs some of the high cost by accepting a lower margin, they will earn less profit on the customers they retain.

- (c) In most scenarios, we would expect that the firm will pass through some of the input cost increase. The extent of pass through will depend on which price level maximises its profits (ie minimises the impact of the higher input cost level), which in turn partly depends on the demand response.²⁴³⁰ Accordingly, while higher input costs are likely to result in the reduction of margins, the existence of positive margins is not evidence that foreclosure is not taking place.²⁴³¹

- 7.623 In other words, while we would generally expect foreclosure to result in the lowering of margins, a firm with positive margins may still be partially foreclosed.
- 7.624 Further, there is a range of positive margin levels (associated with different pass-through responses) at which foreclosure may be happening. We therefore do not consider it appropriate to define a bright line or a strict threshold for margins, below which we would automatically find that rivals' ability to compete is materially affected and they consequently compete less effectively. However, as set out below we find it informative to look at industry-wide margins as one factor to inform our assessment.
- 7.625 More generally, we do not consider it necessary or appropriate to define a threshold or a benchmark in relation to each feature of the market (as it is unreasonable to consider them in isolation), let alone define them for particular elements of each feature. We have set out the benchmark of a well-functioning market elsewhere in this report.
- 7.626 For the reasons above, we consider margins to be relevant but not determinative as an indicator of partial foreclosure. We have instead focused our assessment on the extent to which the evidence on margins is consistent or inconsistent with partial foreclosure.

²⁴³⁰ The less responsive customers are to price changes, the more likely the rival firm is to pass on the input cost increase and protect its margin. In a market with customer differentiation and ability to price discriminate, the extent of pass through may also vary across customer groups.

²⁴³¹ Simple numerical examples can demonstrate this principle. The following example illustrates how margins may be an unreliable indicator of partial foreclosure due to pass-through effects. Consider a stylised scenario where Firm A supplies an input to Firm B at a unit cost of 50. Firm B sells the final product to its customer for 100, earning a 50% margin. If Firm A raises its input price to 100 and Firm B absorbs the full increase, keeping its customer price at 100, its margin falls to 0%. This would suggest that Firm B makes zero margin and cannot profitably supply the product, resulting in a reduced incentive to compete. Conversely, if Firm B passes through the entire cost increase by raising its price to 150, it earns a 33% margin, but its product becomes less competitive. This illustrates that observing positive margins may not be inconsistent with partial foreclosure.

Estimates of margins

- 7.627 In this section we set out the following analyses that Microsoft, AWS and Google have submitted:
- (a) estimates of Google's and AWS' actual mark-ups and margins on Windows Server and SQL Server;
 - (b) estimates of the margin Microsoft would earn if it was subject to the SPLA licensing costs it charges AWS and Google for Windows Server and SQL Server.
- 7.628 Microsoft argued that AWS and Google's estimated margins are inconsistent with material foreclosure effects, as the margins that Google and AWS could earn are positive and ample.²⁴³² Google submitted that the estimated margins are too low to support effective competition.²⁴³³ As discussed in more detail in the counterstrategies section, Google also submitted that the gross margins on its Windows Migration incentive programme are well below margins typically observed for the supply of cloud infrastructure services. AWS submitted that once a broader and more meaningful set of costs are considered ([REDACTED]) the resulting margins affect rivals' ability to compete effectively.²⁴³⁴

Google and AWS' mark-ups and margins

- 7.629 In this section, we set out evidence on AWS and Google's margins and mark-ups. As mark-ups are calculated as a percentage of costs and margins as a percentage of revenues (which include profits), mark-ups are always higher. We have taken this into account when making comparisons.
- 7.630 In relation to actual mark-ups and margins we have two sets of estimates:
- (a) Microsoft submitted estimates of the mark-ups that AWS and Google charge their current customers for the use of Windows Server (Microsoft's analysis did not include SQL Server).
 - (b) Google submitted data on the total revenues it earned and costs it incurred on Microsoft workloads, including both Windows Server and SQL Server, which we have used to calculate margins.²⁴³⁵
- 7.631 Microsoft estimated the mark-up on the Windows Server software licence cost by first taking the customer-facing price for a Windows VM and subtracting both the

²⁴³² Microsoft's submission to the CMA [REDACTED]; Microsoft's submission to the CMA [REDACTED].

²⁴³³ Google's submission to the CMA [REDACTED].

²⁴³⁴ AWS' submission to the CMA [REDACTED].

²⁴³⁵ We also requested AWS' total revenues and costs incurred directly from licence sales of Windows Server and SQL Server, but it did not have such data available.

customer-facing price for a Linux VM of the same VM specification and AWS'/Google's SPLA costs per vcore hour. It then divided this by AWS'/Google's SPLA costs per vcore hour.²⁴³⁶ Based on this, Microsoft submitted that AWS and Google currently charge list prices for a Windows Server VM with a markup over their SPLA costs of approximately [10-20]%.²⁴³⁷ Microsoft submitted that these results show that there is no impediment to competition.²⁴³⁸

7.632 Google submitted data on its total revenues and costs incurred directly from licence sales of Windows Server and SQL Server.²⁴³⁹ We used this data to calculate margins and found that Google made an estimated margin of [less than 15%] on Windows Server workloads between 2021 and 2023. In relation to SQL Server, the margins we calculated varied according to which edition of SQL Server was used. Between 2021 and 2023, Google made an estimated margin of:

- (a) [20-30]% to [20-30]% on SQL Server Standard, which accounted for [40-50]% of usage in 2023;
- (b) A negative margin [X] on SQL Server Web, which accounted for [30-40]% of usage in 2023; and
- (c) [-10- -5]% to [0-5]% on SQL Server Enterprise, which accounted for [10-20]% of usage in 2023.

Margins Microsoft would earn if subject to SPLA licensing costs

7.633 Microsoft submitted several sets of analyses (both before and after the publication of the PDR) setting out the margins it would earn if it were subject to the SPLA licensing costs it charges AWS and Google for Windows Server and SQL Server.²⁴⁴⁰ In our assessment of the margins analyses provided by Microsoft before the PDR, we noted that these gross margins did not appear particularly high in the context of software markets.²⁴⁴¹

²⁴³⁶ The price of a Linux VM is used as a proxy estimate for the cost of and mark-up on compute.

²⁴³⁷ [X] submission to the CMA [X].

²⁴³⁸ [X] submission to the CMA [X].

²⁴³⁹ [X] response to the CMA's information request [X]. We also requested AWS' total revenues and costs incurred directly from licence sales of Windows Server and SQL Server, but it did not have such data available.

²⁴⁴⁰ Microsoft's estimates of the margins it would earn if it were charged SPLA prices were put forward by Microsoft as the outcome of an as-efficient competitor test, aimed at establishing whether an equally efficient rival could profitably supply if faced with the relevant input cost. [X] submission to the CMA [X]. [X] submission to the CMA [X]. We note that use of an as-efficient competitor test, which focuses on whether an as-efficient competitor earns a negative margin, is not provided for in our guidance and we do not consider that an application of an as-efficient competitor test is necessary or determinative for the statutory questions we must answer in the context of this market investigation. To the extent an as-efficient competitor test focuses on whether a rival is loss-making or not, it does not account for a rival being materially disadvantaged even while continuing to earn a positive margin. We note that this type of analysis can be used as a tool when assessing potential infringements of competition law but as noted above, we are not investigating whether the conduct in question infringes CA98.

²⁴⁴¹ PDR, paragraphs 6.461 and 6.457. Specifically, Microsoft submitted that it would make a margin of [10-20]% and [0-5]% in the UK in 2022 and 2023, respectively; and [0-5]% and [-5-0]% globally in 2022 and 2023, respectively on standalone Windows Server VMs and that considering broader packages would show even more favourable economics

- 7.634 In response to the PDR, Microsoft made further submissions containing additional estimates of the gross margins it would make if it were subject to the SPLA licensing costs.
- 7.635 The main difference between the different sets of estimates submitted by Microsoft is the extent to which different categories of cloud spend are included in the denominator over which gross margins are calculated. As set out in the ‘Aggregation of customers’ spend’ section above, this reflects the variability around how customers aggregate their spend, which in turn affects the calculation of margins.
- (a) The narrowest base of cloud spend used in Microsoft’s analyses is customers’ spend on components necessary to run Windows Server VMs, which Microsoft submitted was all of customers’ spend on non-Linux (ie WS) compute, storage and virtual network. Microsoft submitted that these are a worst-case or lower-bound extremes as they do not include other near-essential IaaS services.²⁴⁴²
 - (b) The widest base of cloud spend Microsoft used was customers’ total spend on Azure. We note that these are effectively the upper bound estimates.²⁴⁴³
- 7.636 Microsoft submitted estimates for various Windows Server customer groups: all Windows Server customers; those eligible for the AHB discount; customers using both Windows Server and SQL Server; and those using both services who are also AHB-eligible. Where SQL Server customers are considered, the estimates are provided as ranges reflecting different assumptions about SQL Server edition usage and [REDACTED].²⁴⁴⁴
- 7.637 The results of Microsoft’s analysis for the minimum set of services required to run Windows Server VMs are presented in Table 7.22 below. They show that Microsoft estimates that it would make the following margins if it were subject to SPLA licensing costs:
- (a) a gross margin of [20-30]% on Windows Server customers, which would reduce to [20-30%] when considering only AHB vCore hours; and

to AWS and Google. We consider that these results are superseded by Microsoft’s updated analyses. [REDACTED] submission to the CMA [REDACTED].

²⁴⁴² Microsoft also submitted results that included bandwidth spend. We have focused on the figures excluding bandwidth, though we have considered both. We note, however, that including bandwidth generally has a limited effect on the estimated margins, as bandwidth spend is, for example, generally small compared to storage. [Microsoft’s response to the CMA’s PDR dated 28 January 2025](#), paragraphs 158, 159, 240 and Table 1. [REDACTED] submission to the CMA in response to the PDR [REDACTED].

²⁴⁴³ This is because the margins on non-Windows Server VM related workloads are assumed to be higher in this analysis.

²⁴⁴⁴ Microsoft also submitted figures for Windows Server-intensive customers, defined as Windows Server users who do not use Linux (ie 100% Windows Server compute). These customer groups are relatively small, so we have not included these figures in our main analysis and have placed greater weight on the other customer groups.

- (b) a gross margin of [10-20]% to [20-30]% on customers who use Windows Server and SQL PAYG IaaS, which would reduce to [5-10]% to [10-20]%% when considering WS AHB vCore hours and SQL PAYG IaaS. These analyses conservatively assumed that customers only utilized WS VMs with the AHB discount, which rivals needed to match while still incurring SPLA costs.²⁴⁴⁵

- 7.638 The results of Microsoft's analysis for a customer's total Azure spend (ie the widest base) are also presented in Table 7.22 below. These account for SQL revenue and licensing costs. They show that Microsoft estimates that if it were subject to SPLA licensing costs it would make margins of [40-50]% and [50-60]% for Windows Server and Windows Server customers eligible for AHB, respectively.²⁴⁴⁶
- 7.639 Microsoft used the above analysis to set out the proportion of customers for which it would make positive gross margins if it were subject to SPLA licensing costs²⁴⁴⁷. Microsoft submitted that Amazon and Google could still profitably supply even the narrowest subset of cloud workload components necessary to use Windows Server.²⁴⁴⁸
- 7.640 Both Google and AWS made submissions in relation to Microsoft's methodology for the above analysis.²⁴⁴⁹
- 7.641 Google submitted that only block storage is necessary to run Windows Server VMs and that by including all storage Microsoft overstates the minimum margin that it would earn if it were subject to the SPLA licensing costs.²⁴⁵⁰ Google submitted gross margin estimates based on Microsoft's methodology adjusted to include only block storage and networking used in conjunction with Windows Server VMs (which Google estimated using information on its own customers).²⁴⁵¹ In response, Microsoft submitted that object storage should be included in the margin calculation as customers using Windows Server are likely to also consume object storage. Microsoft also submitted that Google's calculations for the block

²⁴⁴⁵ These figures do not include bandwidth spend. [X] submission to the CMA in response to the PDR [X].

²⁴⁴⁶ [X] submission to the CMA [X].

²⁴⁴⁷ For the minimum bundle, Microsoft submitted that it would make a positive margin on [90-100]% of Windows Server customers, and [80-90]-[90-100]% when accounting for SQL usage. Across total cloud spend, Microsoft submitted that it would make a positive margin on [90-100]% of Windows Sever customers and [90-100]% of Windows Server and SQL customers. [X] submission to the CMA[X].

²⁴⁴⁸ [Microsoft's response to the CMA's PDR dated 28 January 2025](#), section 14.2.

²⁴⁴⁹ [X] submission to the CMA [X]. [X] submission to the CMA [X].

²⁴⁵⁰ [X] submission to the CMA [X].

²⁴⁵¹ [X]. submission to the CMA [X]. Specifically, it did this by applying the ratio of customer spend on block storage to spend on Windows Server VMs, [X], to the Windows VM spend of Microsoft's customers. Google also made similar adjustments to estimate the proportion of networking spend that is associated with Windows Server VM usage. Using the information in Microsoft's post-PDR submissions based on Microsoft's own data, Google also provided additional calculations of the margin Microsoft would earn on Windows Server VMs (excluding storage and networking) if it paid the SPLA costs. Google submitted calculations estimating the gross margins Microsoft would earn on Windows Server VMs if it paid the SPLA cost, estimating margins of [-10 to -5]% for all Windows Server customers and [-40 to -30]% for Windows Server AHB customers. [X] submission to the CMA [X].

storage usage of customers appears to be based on a small customer sample, which means it may be unreliable.²⁴⁵²

- 7.642 The results of Google's sensitivity analysis for the minimum set of services required to run Windows Server VMs are presented in Table 7.22 below, and show that if Microsoft were subject to SPLA licensing costs it would make:
- (a) a gross margin of [10-20]% on Windows Server customers, which would reduce to [-10 to -5]% when considering only Windows Server AHB customers.
 - (b) a gross margin of [5-10]% to [10-20]% on customers who use Windows Server and SQL Server, which would reduce to [-20 to -10]% to [-10 to -5]% when considering only Windows Server AHB customers.
- 7.643 Google submitted that the gross margins estimated by Microsoft are too low to support effective competition. Specifically, focusing on Microsoft's estimates of the margins on what Microsoft submitted are the minimum bundles²⁴⁵³, Google argued that these were insufficient, noting that Microsoft's estimated gross margins are significantly lower than those typically observed in the cloud services sector and the broader economy.²⁴⁵⁴
- 7.644 AWS submitted that operating margins are not accounted for in Microsoft's cost estimates and also provided the results of a sensitivity test to illustrate the effect of this limitation. Specifically, AWS adjusted Microsoft's methodology by estimating costs based on Azure's operating margin, rather than using Microsoft's estimate of gross margins.²⁴⁵⁵ Google made a similar argument but did not provide a set of results that adjusted for this.²⁴⁵⁶
- 7.645 In response to AWS' submission, Microsoft submitted that gross margins are the appropriate measure, as firms typically assess these when deciding whether to bid for a customer and set prices based on variable rather than fixed costs.²⁴⁵⁷
- 7.646 The results of AWS' sensitivity test conducted using Azure's operating margin for the minimum set of services required to run Windows Server VMs are presented in Table 7.22, and show that if Microsoft were subject to SPLA licensing costs it would make:

²⁴⁵² Furthermore, Microsoft submitted that its analysis likely understates the margin for customers eligible for AHB, as it conservatively assumed that these customers used only Windows Server VMs with the AHB discount. [X] submission to the CMA [X].

²⁴⁵³ In particular, the minimum bundles include [X].

²⁴⁵⁴ [X] submission to the CMA [X].

²⁴⁵⁵ In particular, AWS applies a [X] AWS also submitted an analysis that used AWS' own lower operating margin as an input into the analysis which yielded lower estimated margins. [X] submission to the CMA [X].

²⁴⁵⁶ Specifically, Google noted that using a cost measure for the long-run average incremental cost is the more appropriate benchmark in an as-efficient competitor test. [X] submission to the CMA [X].

²⁴⁵⁷ [X] submission to the CMA [X].

- (a) a gross margin of [10-20]% on Windows Server customers, which would reduce to [5-10]% when considering only customers eligible for the AHB discount; and
- (b) a gross margin of [5-10]% on customers who use Windows Server and SQL, which would reduce to [-5-0]% when considering only customers eligible for the AHB discount.²⁴⁵⁸

7.647 AWS submitted that the costs estimated by Microsoft, once operating margins are accounted for, demonstrate that it would often be impossible for an as-efficient competitor of Azure to make a positive margin on a workload of a given customer subset, and that this provides sufficient evidence that Microsoft's competitors are partially foreclosed from competing for at least some customer workloads.²⁴⁵⁹

Table 7.22: Estimates of the margin Microsoft would earn if subject the SPLA licensing costs by customer group in the UK (2024)

UK Customer group	Microsoft submissions Margin on minimum set of services	Margin on all cloud spend	Google Sensitivity Margin on minimum set of services	% AWS sensitivity Margin on minimum set of services
Windows Server	[20-30]	[40-50]	[10-20]	[10-20]
Windows Server AHB	[20-30]	[50-60]	[-10 to -5]	[5-10]
Windows Server and SQL	[10-20] to [20-30]	-	[5-10] to [10-20]	[5-10]
Windows Server AHB and SQL	[5-10] to [10-20]	-	[-20 to -10] to [-10 to -5]	[-5-0]

Source: Microsoft's submission to the CMA [REDACTED]; [REDACTED] response to the CMA's information request [REDACTED]; AWS' submission to the CMA [REDACTED]; Microsoft's submission to the CMA [REDACTED].

Note: Calculations for Microsoft and Google for SQL are presented as a range, reflecting four different methodologies combining SQL edition splits [REDACTED]. The calculations for Microsoft and Google do not include bandwidth in the minimum bundle or account for associated revenues and costs, whereas AWS includes these. Estimates of the minimum margin for the Windows Server and Windows Server AHB customer groups exclude SQL spend and SPLA costs, whereas the all-cloud spend includes both.

Our assessment of margin estimates

7.648 Below, we set out our views on the gross margin analyses by Microsoft, AWS and Google presented above. Overall, we believe that that each analysis has important limitations (additional to the conceptual limitations already set out at the start of this section relating to the use of margins as an indicator of whether foreclosure may be expected to occur), which affect the way in which we interpret some of the individual estimates.

7.649 Before we set these out, we note that AWS' and Google's actual mark-ups and margins will always only capture those customers who have been won/retained and by definition would not include any customers lost due to foreclosure. This limitation does not affect the estimated margins Microsoft would earn if it were

²⁴⁵⁸ [REDACTED] submission to the CMA [REDACTED].

²⁴⁵⁹ [REDACTED] submission to the CMA [REDACTED].

charged the SPLA prices. However, the latter estimates are subject to other uncertainties and limitations, as set out below.

- 7.650 First, Microsoft's estimates of the mark-up charged by Google and AWS (see above 7.631) do not account for the fact that AWS and Google discount their list prices²⁴⁶⁰ and this analysis will therefore tend to overstate their mark-ups for some customers. Microsoft's estimates of mark-ups on the SPLA costs appear [REDACTED].²⁴⁶¹
- 7.651 Second, as set out above in 'Aggregation of customers' spend', we agree with Microsoft that it is reasonable to include at least some allocation of customer spend on storage and networking in the minimum bundle. However, we consider that the evidence submitted by Google indicates that only block storage is technically required to run Windows Server VMs,²⁴⁶² and for the purposes of establishing the lower bound for potential gross margins only the components that are necessary should be included. However, we agree with Microsoft that Google's block storage estimate appears to be based on a small customer sample.²⁴⁶³
- 7.652 Third, in line with AWS' and Google's submissions, we consider that when considering a rival's incentive to compete, to the extent that certain operating costs are incremental to supplying Windows Server VMs these should be included.²⁴⁶⁴ We also consider that the same applies to the cost of capital, which reflects the opportunity cost of invested capital and the need for a rival to make an economic profit. However, it is unclear to what extent the adjustments made by AWS effectively separate out the relevant incremental parts of these costs. We note that AWS does not use Windows Server VM-specific cost data and relies on a simplistic method for accounting for potentially relevant operating and fixed costs. We also consider that gross margin is still an appropriate measure of relative profitability in the context of a partial foreclosure assessment where other available benchmarks are also expressed as gross margins.
- 7.653 Finally, drawing inferences about whether AWS and Google could profitably supply customers based on the estimated margins Microsoft would earn if it were charged the SPLA prices assumes that Google and AWS have the same cost structure as Microsoft. Microsoft submitted that it would be untenable to suggest that AWS is less efficient than Microsoft, and that Google is also not an inefficient

²⁴⁶⁰ Cloud providers submitted that they discounts at the deal level. [REDACTED] submission to the CMA [REDACTED].

²⁴⁶¹ [REDACTED] response to the CMA's information request [REDACTED].

²⁴⁶² [REDACTED] submission to the CMA [REDACTED].

²⁴⁶³ [REDACTED] submission to the CMA [REDACTED].

²⁴⁶⁴ Microsoft presents its submitted margins analysis as an application of the as-efficient competitor test framework, which was developed in the context of margin squeeze exclusionary abuse cases. In margin squeeze cases, as-efficient competitor tests typically use the cost measure of long-run average incremental cost, rather than cost measures based on gross or operating margins. Measures of long-run average incremental cost include both variable and fixed costs directly attributable or incremental to supplying the product, often covering operational costs and the relevant industry cost of capital. The use of long-run average incremental cost, as a cost floor for downstream activity is consistent with the principle that, in the long run, a rational profit-maximising firm would have no economic reason to provide a service below this level.

competitor.²⁴⁶⁵ We have not undertaken a detailed analysis of cloud providers' costs but note that this means that Microsoft's margin estimates may not be accurate approximations of AWS' and Google's margins.

- 7.654 Because of the above limitations, we do not consider that any one of the specific sets of estimates is entirely accurate. However, taken together, they offer a degree of insight into the overall level of margins and we have considered this (within the context of the wider question about the relevance of margins to partial foreclosure) as part of our assessment as set out below.

Conclusions on margins and mark-ups

- 7.655 AWS and Google's current profit margins can be useful indicators but they are not determinative of partial foreclosure. As set out above, profit maximising behaviour may lead to increased costs being at least partly passed on to customers in the form of higher prices. In fact, the foreclosure effect may be stronger if the firm passes on the cost increase to customers. As a result, partial foreclosure could still be taking place, despite positive margins being generated. Similarly, estimates of the margin the foreclosing firm would earn if subject to the input costs it charges its rivals implicitly assume it does not raise its prices. Therefore, for both sets of margin estimates, the analysis does not directly address the relevant question of whether competition from its rivals is being weakened.
- 7.656 In addition to these conceptual difficulties in interpreting margins, in this case there are also practical difficulties:
- (a) Margins are particularly difficult to estimate in the cloud services markets because, as set out above in 'Aggregation of customers' spend', the extent to which customers aggregate their demand for cloud services varies. Even for the minimum set of services calculations, there is an uncertainty around the estimates of the proportion of storage and networking strictly required to run Windows Server VMs. Because of this uncertainty there is a wide range of plausible margin estimates.
 - (b) As we explained above, we do not consider that there can be a bright line threshold for margins which we would use to automatically assess whether rivals are being foreclosed. This would be the case in all markets. However, in some markets we may have access to reliable data on what margins rival firms make when they are not subject to foreclosure, which we could use to create a benchmark to compare against. Because of the complexity of the cloud services markets, we consider this to be more difficult in these markets compared to some other markets.

²⁴⁶⁵ Microsoft's response to the CMA's PDR dated 28 January 2025, footnote 236

7.657 Notwithstanding these limitations, our view is that:

- (a) Despite the uncertainty around demand aggregation, as set out above in 'Aggregation of customers' spend', it is clear that the appropriate set of services to use in the estimations is narrower than customers' total spend on cloud services and wider than an individual product or service (the minimum bundle) for most customers. Therefore, we can be reasonably confident that, for an average customer, the margins fall somewhere between those two extremes.
- (b) Although we lack a clear benchmark to compare against, we consider that margin levels in this industry (ie not only for Microsoft workloads) are generally high, reflecting that cloud services are characterised by low variable costs and high fixed costs. [X].²⁴⁶⁶
- (c) These are significantly above Microsoft's calculations of mark-ups and our own calculations of Google's actual margins for Windows Server and SQL workloads (including those for their Windows Migration incentive programme).
- (d) These are all above Microsoft's estimates of margins it would make if it were subject to SPLA costs based on total cloud spend ([50-60]%), which would be the upper bound in this set of calculations.
- (e) They are significantly above the margins based on the minimum set of services, regardless of whether we consider Microsoft's estimates or Google's estimates.
 - (i) Our estimates of Google's actual margins are [less than 15%] for Windows Server. For certain editions of SQL Server these margins are negative.
 - (ii) The estimated margins Microsoft would earn if it paid SPLA licensing costs range from [20-30]% according to Microsoft's analysis to [10-20]% in Google's estimates.
 - (iii) The above estimates are [less than 30%] and [less than 20%] for the subgroup of customers who also use SQL Server.
- (f) For customers eligible for AHB they range between [20-30]% and [-10 to -5]%, depending on whether we consider Microsoft's or Google's estimates and whether they use SQL Server.

²⁴⁶⁶ See the analysis of gross margins in appendix E for more detail.

7.658 We therefore consider that, despite the uncertainties and limitations of this analysis, the level of these estimated mark-ups and margins is consistent with partial foreclosure. We consider this finding below in the round with other evidence.

Relative usage

Microsoft, AWS and Google submissions

- 7.659 Microsoft submitted that relative Windows Server VM usage is not a lens which measures competition.²⁴⁶⁷ Microsoft submitted that it is not surprising that it is winning workloads that involve its software from customers who are on average bigger users of Microsoft software. It submitted that the justification for this is that these are entities with whom Microsoft will typically have close and longstanding relationships. Microsoft also submitted that it should not be surprising that Microsoft wins workloads for which it offers a discount, and that this (the AHB, as a discount) is strong evidence of price competition. It submitted that AWS and Google are capable (and do) compete with their own discounts.²⁴⁶⁸
- 7.660 Microsoft submitted that if its licensing practices allow it to win more customer workloads, one may expect Azure to have a disproportionate share of larger customers because large enterprise customers tend to have on-premise licences that would benefit from AHB. [REDACTED].²⁴⁶⁹
- 7.661 AWS submitted that the fact that Microsoft product usage is higher on Azure than on AWS or Google by hundreds of percentage points, and that the difference has increased significantly since the introduction of the licensing restrictions beginning in October 2019,²⁴⁷⁰ clearly indicates the effect of the licensing restrictions on Microsoft's customers and competitors, with no other event in that timeframe which could have affected usage in the same way.²⁴⁷¹
- 7.662 Google submitted that the especially high usage of Windows Server and SQL Server on Azure compared to on Google, as noted in the PDR, were consistent with Microsoft's licensing practices harming the UK public sector and UK enterprises alike.²⁴⁷²

²⁴⁶⁷ [Microsoft's response to the PDR](#), paragraph 227.

²⁴⁶⁸ [Microsoft's response to the PDR](#), paragraph 229.

²⁴⁶⁹ Microsoft's submission to the CMA [REDACTED].

²⁴⁷⁰ We note that the 2019 changes meant that customers of Listed Providers could no longer use their pre-existing licences for Microsoft software on a BYOL basis on dedicated hardware, but did not affect the licensing terms for shared hardware. See 'Timeline of licensing practices' for further detail.

²⁴⁷¹ [AWS' response to the CMA's PDR dated 28 January 2025](#), paragraph 74.

²⁴⁷² [Google's response to the CMA's PDR dated 28 January 2025](#), paragraph 12.

Our analysis

- 7.663 We compared Windows Server and SQL Server usage levels across Azure, AWS and Google. To the extent that usage of Windows Server and SQL Server is significantly high on Azure than on AWS and Google, this would be consistent with Microsoft's licensing practices having had an impact on AWS and Google's competitive offerings. Similar levels of usage across each cloud would be more consistent with a lack of impact on AWS and Google's competitive offerings.
- 7.664 AWS and Microsoft provided us with data on the average usage (in vcore hours) of Windows Server and SQL Server by customers on their respective platforms in each year from 2020-2023; Google provided us with equivalent data for the years 2021-2023.²⁴⁷³ We used this data to compare the relative average usage rates of Windows Server and SQL Server on each platform over time.
- 7.665 Microsoft's data includes usage of SQL Server licences sold on a BYOL and PAYG basis, as well as licences acquired from third parties, across all editions of SQL Server.²⁴⁷⁴ AWS' and Google's data only includes usage of SQL Server licences purchased through their respective SPLAs and therefore includes PAYG usage of Enterprise Edition, Standard Edition and Web Edition only.²⁴⁷⁵ Therefore, our main analysis (which compares usage rates over time) overestimates the difference in usage of SQL Server on Azure compared with AWS and GCP to some extent. To counter this, we present a sensitivity analysis that compares PAYG usage of Enterprise, Standard and Web editions in 2022 only.
- 7.666 This analysis cannot determine with certainty whether there is causal link between Microsoft's licensing practices and relative usage levels. Variation in usage levels across providers could in principle be explained by other factors, including for example a more general preference to use Microsoft software on Azure, the length of customer relationships with Microsoft, or the use of different software products that are substitutes for the Microsoft software on AWS and GCP.²⁴⁷⁶ Therefore, the results of this analysis must be considered in the context of other evidence, including on the availability of substitutes for each product and the relative position of the providers in cloud infrastructure services.

²⁴⁷³ Virtual core hours (vcore hours) are hours of usage normalised for the number of core processing units being used to run a particular instance or operating system environment (OSE). For example, using Windows Server OS on a VM that uses 4 CPUs for one hour constitutes 4 vcore hours of usage.

²⁴⁷⁴ Microsoft's time-series data includes usage of the Enterprise, Standard, Web, Express, Developer, and Power BI editions of SQL Server.

²⁴⁷⁵ [X] response to the CMA's information request [X].

²⁴⁷⁶ Average customer usage rates were calculated as the total number of vcore hours used for Windows Server and SQL Server divided by the total number of customers. Whilst in principle differences in average Windows Server and SQL usage could be partly due to the differences in total average vcore hours (ie which include non Windows Server and non-SQL hours) we have not received any evidence that Azure customers use a significantly greater number of total vcore hours and therefore consider it appropriate to compare average vcore hours per customer used on Windows Server or SQL server between Azure, AWS, and GCP.

7.667 The figure below shows the average usage of Windows Server over the years 2020-2023 on Azure and AWS and the years 2021-2023 on GCP.

Figure 7.3: [REDACTED].

Source: CMA analysis of responses to CMA's information requests [REDACTED].

7.668 Our analysis suggests that Windows Server usage on Azure is significantly higher than on AWS and GCP. In particular, average usage on Azure was [over 200%] and [over 300%] higher than on AWS and GCP in 2023, respectively. The difference between usage on Azure and AWS and GCP has increased significantly since 2021.²⁴⁷⁷

7.669 The figure below shows the relative average usage of SQL Server over the years 2020-2023 on Azure and AWS and the years 2021-2023 on GCP.

Figure 7.4: [REDACTED].

Source: CMA analysis of responses to CMA's information requests [REDACTED].

7.670 Our analysis suggests that SQL Server usage on Azure is significantly higher than on AWS and GCP. In particular, average usage on Azure was [over 1000%] and [over 1500%] higher than on AWS and GCP, respectively in 2023. The difference in relative usage on Azure and AWS has increased since 2020, whereas the difference on Azure and GCP has decreased since 2021 although from a significantly higher starting point.

7.671 However, as explained above, to some extent this graph overstates the difference in usage of SQL Server on Azure compared with AWS and GCP. We present our sensitivity analysis below.

7.672 The figure below shows the relative average usage of SQL Server Enterprise, Standard, and Web editions licensed on a PAYG basis in 2022. We also include AHB usage of these three editions and total average usage across all editions and deployment types for comparison.

Figure 7.5: [REDACTED].

Source: CMA analysis of responses to CMA's information requests [REDACTED].

7.673 Our sensitivity analysis suggests that when comparing PAYG usage of Enterprise, Standard, and Web editions, usage of SQL Server on Azure was significantly higher than on AWS and GCP in 2022. PAYG usage on Azure was [over 500%] and [over 800%] higher than on AWS and GCP, respectively.

²⁴⁷⁷ One of the internal documents submitted [REDACTED]. As such, it is difficult to interpret the graphs from this internal document. Instead, we place more weight on the growth rates we have calculated ourselves. [REDACTED] response to the CMA's information request [REDACTED].

7.674 AHB (BYOL) SQL Server usage makes up less than half of total SQL Server usage on Azure. As noted above, customers may receive a lower price and higher quality service when they bring their own SQL Server licences to Azure rather than AWS or GCP. As such, we expect BYOL usage to comprise a larger proportion of total SQL Server usage on Azure than AWS or GCP. We do not, therefore, expect the difference in total SQL Server usage to differ significantly from those presented in figure above.

Our assessment

- 7.675 We considered Microsoft's submission that, if its licensing practices allow it to win more customer workloads, it may be expected that it will win a disproportionate share of large customers because large enterprise customers tend to have on-premises licences that would benefit from AHB. We make the following observations:
- (a) First, it is not necessarily the case that large customers are more likely to have pre-existing licences that qualify for AHB for the relevant Microsoft software. Digital native customers (which are unlikely to have pre-existing licences that qualify for AHB for the relevant Microsoft software) may also have high cloud spend.
 - (b) Second, our analysis on relative usage provides a more direct assessment of the extent to which Microsoft is winning cloud workloads that use its software.
- 7.676 The results of our relative usage analysis suggest that customers disproportionately choose to deploy their Windows Server and SQL Server workloads on Microsoft. As outlined above, there are multiple factors that could drive such differences in usage. However, given the magnitude of the differences, even if just a portion of these is driven by the licensing practices this would still indicate that the latter are having an effect on customer choice.
- 7.677 As outlined above, we have considered these results alongside evidence on Microsoft's market power in the relevant software markets. In particular, customers finding it difficult to switch away from the respective software products suggests that Microsoft's licensing practices may have led to a worsening of AWS' and Google's competitive offerings for Microsoft software workloads, which in turn drives customers to disproportionately choose to deploy these workloads on Azure.
- 7.678 We agree with Microsoft's submission that Windows Server VM usage is not a lens which measures competition; however, we do consider it alongside other evidence as an indicator which can be consistent or inconsistent with the licensing practices having an impact on AWS and Google's competitive offerings. For example, if levels of Windows Server usage were broadly similar across the

different cloud providers, this could suggest licensing restrictions are not affecting rivals' competitive offerings. However, if there is a large disparity of Windows Server usage between Microsoft and its rivals, then there are multiple factors that could be responsible.

- 7.679 As we outline above, licensing restrictions, customer preferences, customers using substitutes for Microsoft software on other clouds, as well as the length of customer relationship, could be responsible for the disparity in relative usage. However, as we describe above, the magnitude of the difference indicates that licensing practices are likely having an effect. In addition, given the evidence on the difficulty of modernisation and the evidence that many customers are unwilling or unable to switch away [REDACTED], we consider that Microsoft's licensing practices are likely a contributing factor to the disparity in usage.

Possible AWS and Google counterstrategies

- 7.680 We have assessed whether AWS and Google may have counterstrategies at their disposal that allow them to protect their competitiveness against the effects of Microsoft's licensing practices and to prevent any harm to competition.²⁴⁷⁸ We have considered a range of counterstrategies including migration incentives, modernisation and strength in non-cloud areas. We also asked AWS and Google to provide internal documents that discuss their ability to compete with Azure and/or any counterstrategies that they could adopt (or have already adopted) in relation to their ability to compete. These documents discussed modernisation efforts (presented above), as well as other counterstrategies, which we present below.²⁴⁷⁹
- 7.681 Google submitted that the counterstrategies were [REDACTED] ineffective in helping it to win Microsoft workloads, which it said is evidenced in the fact that customer usage of Windows Server on GCP remains [REDACTED] low.²⁴⁸⁰
- 7.682 Below we summarise the evidence from Google's submissions and its internal documents. The evidence suggests that Google has various strategies to try to win Microsoft workloads. However, Google's internal documents generally do not appear to discuss the success of these counterstrategies.
- (a) Targeting customers based on the proportion of their estate that uses Windows vs Linux. One of Google's internal documents from 2022 divides customers based on the proportion of the customer's usage of Windows Server and suggests it is less effective in competing for customers with relatively higher Windows Server use. In particular, the document [REDACTED]. The

²⁴⁷⁸ CC3 (Revised), paragraph 283.

²⁴⁷⁹ We note that some of these documents have been produced during the course of the investigation, and have therefore been interpreted accordingly.

²⁴⁸⁰ Google's response to the CMA's information request [REDACTED].

document suggests [REDACTED]. The document states Google's assessment at the time that [REDACTED].²⁴⁸¹

- (b) Offering credits [REDACTED].²⁴⁸² [REDACTED].²⁴⁸³
- (c) Google submitted that [REDACTED].²⁴⁸⁴ Google also submitted a supporting internal document which shows the gross margins on [REDACTED] are between [REDACTED] (which Google notes is well below margins typically observed for the supply of cloud infrastructure services).²⁴⁸⁵ The document stated that over a ten month period since launch, [REDACTED]²⁴⁸⁶ [REDACTED].²⁴⁸⁷
- (d) [REDACTED], this internal document suggests [REDACTED]. The same document notes that this incentive can act as a trigger for modernisation.²⁴⁸⁸ This may suggest Google expects to be able to sell other services in the future in addition to those in the original deal.
- (e) Optimising licence usage to reduce cost. For example, [REDACTED].^{2489, 2490}
- (f) One of Google's documents states that licence optimisations are [REDACTED] in reducing licensing costs.^{2491, 2492} Another internal document suggests [REDACTED].²⁴⁹³ However, Google submitted that it is currently [REDACTED].²⁴⁹⁴

7.683 In one of these documents, Google states that Windows Server is generally challenging as new licences are ineligible for BYOL on GCP, but that BYOL of SQL Server is generally unproblematic because licence mobility through Software Assurance is available.²⁴⁹⁵

7.684 Google submitted that given the critical importance of traditional enterprises' workloads (which it says are dependent on Microsoft software), it has [REDACTED] aimed at winning Windows workloads notwithstanding the significant hurdles. Google added that [REDACTED] do not come close to offsetting the adverse impact of Microsoft's practices.²⁴⁹⁶

²⁴⁸¹ Google's response to the CMA's information request [REDACTED].

²⁴⁸² Google's response to the CMA's information request [REDACTED].

²⁴⁸³ This document sets out [REDACTED].

²⁴⁸⁴ Google's response to the CMA's information request [REDACTED].

²⁴⁸⁵ Google's response to the CMA's information request [REDACTED].

²⁴⁸⁶ Google's response to the CMA's information request [REDACTED].

²⁴⁸⁷ Google, [AnnualReport2024_Cover_0414.indd](#), page 40, accessed on 19 June 2025.

²⁴⁸⁸ Google's response to the CMA's information request [REDACTED].

²⁴⁸⁹ Google's response to the CMA's information request [REDACTED].

²⁴⁹⁰ This is supported by an internal document, which states that certain technical adjustments offered by Google (namely disabling simultaneous multithreading) can save [REDACTED] of licensing cost, [REDACTED].

²⁴⁹¹ [REDACTED].

²⁴⁹² Google's response to the CMA's information request [REDACTED].

²⁴⁹³ [REDACTED]. Google's response to the CMA's information request [REDACTED].

²⁴⁹⁴ Google's response to the CMA's information request [REDACTED].

²⁴⁹⁵ Google's response to the CMA's information request [REDACTED].

²⁴⁹⁶ Google's response to the CMA's information request [REDACTED].

- 7.685 AWS submitted that customers of all providers benefit from improving the efficiency of their architecture and reducing costs.²⁴⁹⁷ AWS has similar initiatives to Google's as presented above to provide discounts to customers with Microsoft licences or work to reduce licensing costs in the cloud.²⁴⁹⁸
- 7.686 AWS provided one document discussing its counterstrategies, which it says is an analysis of the state of the 2025 EC2 Windows business that identifies the challenges and opportunities for this business area, with a particular focus on Microsoft's licensing restrictions.²⁴⁹⁹
- 7.687 This document suggests that AWS has pursued a number of counterstrategies. These include:
- (a) Optimising compute for Windows licence included EC2s. The document states that the Optimize CPUs feature could allow the customer to save about 50% of the vCPU-based licensing cost while still using the same memory and IOPS of a full-sized instance. [REDACTED].²⁵⁰⁰
 - (b) Offering customers credits of \$250 (or \$200) for every Windows Server rehosted to EC2 once they migrate at least 80 (or 40) servers in a calendar month. [REDACTED].²⁵⁰¹
 - (c) Accelerating Windows migrations with MAP Partner cash fund. This appears to be a cash pool, in the form of AWS service credits or partner investments, to assist with Windows migration planning, readiness, and implementation. [REDACTED].²⁵⁰²
- 7.688 Similar to Google, AWS divides EC2 Windows customers based on the proportion of the customer's usage of Windows Server [REDACTED]. In the document, AWS says it is [REDACTED]. The composition of customer types that AWS currently has is affected by both the win rate of those customers, and whether those customers [REDACTED]. Indeed, AWS' document also suggests that [REDACTED]. This document gave the example of [REDACTED].²⁵⁰³
- 7.689 The document suggests that AWS expects its counterstrategies will deliver a [REDACTED]% year-on-year growth rate for EC2 Windows in 2025, compared to [REDACTED]% in 2022, [REDACTED]% in 2023, and [REDACTED]% in 2024. In comparison, the document says that Linux grew [REDACTED]% in 2024.²⁵⁰⁴

²⁴⁹⁷ AWS' submission to the CMA [REDACTED].

²⁴⁹⁸ For example AWS, [ESG-Showcase-Evolve-OLA-Data-27-Apr-2022 final.pdf](https://aws.amazon.com/microsoft/), <https://aws.amazon.com/microsoft/>, accessed on 09 July 2025, also see below.

²⁴⁹⁹ AWS' response to the CMA's information request [REDACTED].

²⁵⁰⁰ AWS' response to the CMA's information request [REDACTED].

²⁵⁰¹ AWS' response to the CMA's information request [REDACTED].

²⁵⁰² AWS' response to the CMA's information request [REDACTED].

²⁵⁰³ AWS' response to the CMA's information request [REDACTED].

²⁵⁰⁴ AWS' response to the CMA's information request [REDACTED].

- 7.690 Microsoft has submitted that one reason for the difference in relative usage across clouds is that AWS and Google have the incentive to reduce customers' use of Microsoft software through modernisation.²⁵⁰⁵ We discuss modernisation earlier in this chapter, and find that the evidence suggests that, while some customers are modernising workloads, modernisation attempts by AWS and Google have had a limited effect on moving customers away from the Microsoft software products.
- 7.691 Microsoft has also submitted that AWS and Google have various other strengths in the provision of cloud services,²⁵⁰⁶ and have invested significantly to dilute the cost of Microsoft IP.²⁵⁰⁷ In this section we present evidence on other potential counterstrategies of AWS and Google, including their strengths in other areas and initiatives aimed at reducing the cost of Microsoft software.
- 7.692 Microsoft has submitted that Google is an AI leader and that Google's leadership in AI is helping it win customer migrations, including those involving Microsoft IP (such as Windows Server and SQL Server).²⁵⁰⁸ It has also submitted, referring to market research from Gartner, that Google is strong in AI/database integration (which is relevant to consider with respect to SQL Server workloads).²⁵⁰⁹
- 7.693 We have gathered evidence on providers' strength in AI (set out in 'Impact of AI' chapter). We have not seen evidence which suggests Google is currently benefitting from strength in AI to win Microsoft software workloads. We note that AI may be a factor in promoting migration to the cloud (addressed above in factors which may affect the importance of the input in the future). However, elsewhere in this report, we explain that AWS, Microsoft and Google have very similar competitive strengths in AI and show customer evidence which suggests that any strength a cloud provider might have in AI will not necessarily translate into winning 'traditional' or 'large' enterprise customers at scale (ie to become their main cloud provider). In addition, while AI [REDACTED].²⁵¹⁰

Our assessment

- 7.694 We considered possible AWS and Google counterstrategies to assess whether they can protect their competitiveness against the impact of Microsoft's licensing practices. A firm can have counterstrategies in place and still be materially disadvantaged by the practices it is seeking to counteract. For example, reducing costs from a high base may still result in high costs. Therefore, we have considered the scope and success of any counterstrategies in place.

²⁵⁰⁵ Microsoft's response to the PDR [REDACTED]. Microsoft's submission to the CMA [REDACTED].

²⁵⁰⁶ Microsoft's submissions to the CMA [REDACTED]. We address AWS' and Google's strength in providing other software products as part of the conceptual framework above.

²⁵⁰⁷ Microsoft's submission to the CMA [REDACTED].

²⁵⁰⁸ Microsoft's submission to the CMA [REDACTED].

²⁵⁰⁹ Microsoft's submission to the CMA [REDACTED].

²⁵¹⁰ Google's response to the CMA's information request [REDACTED].

- 7.695 Overall, the evidence shows that AWS and Google have counterstrategies in place which are largely similar: offering credits, optimising compute (to reduce licensing costs) and targeting customers based on the proportion of the customer's usage of Windows Server.
- 7.696 Customer targeting by AWS and Google may also be consistent with partial foreclosure as it may suggest they are less effective in competing for customers with large Microsoft estates. [REDACTED], both AWS' and Google's internal documents suggest [REDACTED].
- 7.697 However, AWS' internal document, set out above, also suggests [REDACTED]. Our findings on modernisation above are that the evidence suggests that while some customers are modernising workloads, modernisation attempts by AWS and Google have had a limited effect on moving customers away from Microsoft software products. Considering licence optimisation, to the extent that optimisation of costs through reducing compute can also be achieved on Azure, this may not improve AWS' and Google's offering relative to Azure.
- 7.698 We also note Google's internal document, set out above, that states SQL Server is generally unproblematic because licence mobility is available. We set out above that customers are able to use the BYOL route to use their pre-existing SQL Server licences on AWS and Google and as such, have considered SQL Server PAYG in the relevant analysis detailed above.
- 7.699 While some evidence suggests that AWS and Google are able to improve their win rate of customers using these strategies, overall we consider that they have not been able to significantly mitigate against the effect of Microsoft's licensing practices.

Customers' choice of cloud providers

- 7.700 In this section, we set out evidence on how licensing costs may affect customers' choice of cloud providers. We have considered:
- (a) Customer evidence on their perception of the difference between providers' competitiveness in terms of hosting workloads that use Microsoft software, the factors that they consider when choosing their cloud provider and the role of Microsoft's licensing practices in their choices;
 - (b) similar evidence, specifically in relation to VDI, this is the primary use case for the 'client side' relevant Microsoft software products in the public cloud;
 - (c) VDI providers' views on customers choice of public cloud for VDI services and selection factors when choosing a public cloud for VDI; and

- (d) AWS' and Google's examples and internal documents in relation to customers' choice.

Perceived differences in competitive offerings

- 7.701 We gathered evidence from randomly selected Microsoft, AWS and Google customers from (i) different revenue brackets²⁵¹¹ for licensing specific questions and (ii) large customers for general public cloud questions, including questions concerning the importance of different factors in their choice of cloud. We have also considered customer evidence from the Jigsaw report.
- 7.702 We gathered evidence from customers on their perceptions of any price differences between using Microsoft software products²⁵¹² on Azure compared to other public clouds:
- (a) most of these customers said that there were price advantages from using Microsoft software products on Azure;²⁵¹³
 - (b) some of these customers indicated that the price advantage arose from receiving AHB on Azure compared to needing to repurchase Microsoft licences for use on other clouds.^{2514,2515} For example, one customer said that moving to a non-Microsoft cloud writes off the investment in perpetual licences and the Software Assurance paid on them.²⁵¹⁶
- 7.703 We also asked customers about quality differences in using Microsoft software products on Azure compared to rival clouds. We set out their responses in the Importance of the input section above. In summary, most customers that responded to this question did not identify quality differences in using Microsoft software products across different clouds. However, there may be some quality factors, particularly security updates, that a subset of customers consider when making their choice of public cloud.

²⁵¹¹ In particular, the revenue brackets were < \$10,000, \$10,000 – \$1Million, \$1Million – \$5Million, \$5Million – \$10Million, \$10Million – \$20Million, > \$20Million.

²⁵¹² The question referred to Windows Server, SQL Server, Windows 10/11, Microsoft Office, Microsoft 365 and Active Directory/Azure AD specifically.

²⁵¹³ Responses to the CMA's information requests [X].

²⁵¹⁴ We acknowledge that it is not technically possible to BYOL for Windows Server to Azure because Microsoft is included as a Listed Provider. However, Microsoft's website says that the AHB means that the licence for Windows Server is covered, so the customer only needs to pay for the base compute rate of the VM. ([Azure Hybrid Benefit for Windows Server](#), accessed on 18 November 2024). Further, evidence from customers suggests that customers perceive it is possible to BYOL to Azure. As such, we consider that the financial effect is that of BYOL to Azure.

²⁵¹⁵ Responses to the CMA's information requests [X].

²⁵¹⁶ [X] response to the CMA's information request [X].

Selection factors

7.704 Customer evidence provided us with important insights into the effect of Microsoft's licensing practices on choice of public cloud provider. In particular:

- (a) We sought to understand how customers select their public cloud service provider. We asked Azure customers whether Microsoft licensing practices had an effect on their choice of public cloud provider. When provided with a list of potentially relevant factors, including the ability and ease of obtaining licences,²⁵¹⁷ some of these customers described licensing as being relevant to their choice of Azure.²⁵¹⁸
- (b) We also asked large customers to rate the importance of cost and ability to use software licences.²⁵¹⁹ Many of these customers said it was important or very important in their choice of public cloud;²⁵²⁰ many of these customers were neutral;²⁵²¹ a few customers rated it as unimportant or not important at all.²⁵²² In their explanations of those ratings, some customers explicitly mentioned the ability to bring on-premises licences to the public cloud was important.²⁵²³ One of these customers submitted that Microsoft makes the cost and ability to use software licences with competitors such as AWS difficult, particularly with respect to Windows Server and SQL Server. It said that it is uncompetitive to use a competitor cloud for Microsoft software products use cases because it is expensive, complex and difficult.²⁵²⁴
- (c) The Jigsaw report stated that participants often struggled to precisely detail why their organisation uses Azure, beyond describing it as a natural choice for both technical and financial reasons.²⁵²⁵ Licensing terms appear to contribute to these reasons. For example, one customer identified the ability to port licences to use on Azure compared to having to re-license on AWS, and said that this made some workloads more competitive on Azure. Another customer said that they were able to get critical security patches for longer if they went to Azure.^{2526,2527}

²⁵¹⁷ We asked customers to consider when they were selecting Azure as their public cloud and explain the relevant factors and their importance in this decision. In particular, we asked customers to consider the relative importance of (i) familiarity with the Microsoft ecosystem; (ii) the ability and/or ease of obtaining licences for Microsoft software products in their choice of Azure; and (iii) already using Azure for back-end management of other apps.

²⁵¹⁸ Responses to the CMA's information requests [§].

²⁵¹⁹ We have specified in the question that 5 means very important and 1 means not important at all. We have interpreted 4 as important, 3 as neutral, and 2 as unimportant.

²⁵²⁰ Responses to the CMA's information requests [§].

²⁵²¹ Responses to the CMA's information requests [§].

²⁵²² Responses to the CMA's information requests [§].

²⁵²³ Responses to the CMA's information requests [§].

²⁵²⁴ [§] response to the CMA's information request [§].

²⁵²⁵ We note that the Jigsaw report stated that the decision to go with Azure was often taken several years ago (over five, sometimes more than ten years ago). [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), paragraph 7.3.2.

²⁵²⁶ [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), paragraph 7.2.3.

²⁵²⁷ We understand critical security patches to include ESUs.

- 7.705 However, other factors also had a role in some customers' choice of cloud. The Jigsaw report noted that the role of software licensing in the decision to go with Azure is difficult to unpick, and added that participants did not single out licensing as a key factor on its own influencing their decision.²⁵²⁸
- 7.706 There was a broad consensus among the customers we heard from that existing skills and familiarity with the Microsoft ecosystem were also factors in choosing Azure as their cloud provider.²⁵²⁹ A few of these customers identified themselves as a 'Microsoft shop' or a 'Microsoft first' organisation.²⁵³⁰ The Jigsaw report similarly found that pre-existing use of Microsoft was often closely related to participants' original take up of Azure, with some participants identifying themselves as 'Microsoft shops'.²⁵³¹
- 7.707 Licensing factors did not play a role in some customers' choice of cloud:
- (a) In relation to the selection factors customers use when choosing a public cloud, a few customers said that they did not consider licensing in their choice of Azure.²⁵³² One of these customers said that licensing was not a persuasive factor in its decision making.²⁵³³ Similarly, as discussed above, a few large customers rated the cost and ability to use software licences as unimportant or not important at all in their choice of public cloud.²⁵³⁴ Of these, two customers indicated that licences were not relevant for their applications²⁵³⁵ (and therefore they are not customers of Microsoft licences for their cloud applications) and another said that licence costs are bundled in for their use cases.²⁵³⁶
 - (b) As set out above, most customers that responded to our question did not identify quality differences between using Microsoft software products across public clouds.²⁵³⁷ They therefore said that their choice of cloud was not

²⁵²⁸ We note that the Jigsaw report stated that the decision to go with Azure was often taken several years ago (over five, sometimes more than ten years ago). [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), paragraph 1.4.30-31.

²⁵²⁹ We asked customers to consider the relative importance of (i) familiarity with the Microsoft ecosystem; (ii) the ability and/or ease of obtaining licences for Microsoft software products in their choice of Azure; and (iii) already using Azure for back-end management of other apps. The other factors listed by the CMA were price, including discounts or cloud credits; service quality; AI capabilities; number and location of data centres; existing relationship with the cloud provider; range of cloud infrastructure services offered by the cloud provider; range of services offered by ISVs; cloud-specific skills of your employees; ease of integration with your existing IT (such as private cloud or traditional IT); ease of integration with other public clouds; and data sovereignty requirements. Customers also had the option to identify and rate other factors. Responses to the CMA's information requests [REDACTED].

²⁵³⁰ Responses to the CMA's information requests [REDACTED].

²⁵³¹ [Cloud Services Market Investigation Qualitative Customer Research](#) conducted by Jigsaw (2024), paragraph 7.3.1.

²⁵³² We asked customers to consider when they were selecting Azure as their public cloud and explain the relevant factors and their importance in this decision. In particular, we asked customers to consider the relative importance of (i) familiarity with the Microsoft ecosystem; (ii) the ability and/or ease of obtaining licences for Microsoft software products in their choice of Azure; and (iii) already using Azure for back-end management of other apps. Responses to the CMA's information requests [REDACTED].

²⁵³³ [REDACTED] response to the CMA's information request [REDACTED].

²⁵³⁴ Responses to the CMA's information requests [REDACTED].

²⁵³⁵ Responses to the CMA's information requests [REDACTED].

²⁵³⁶ [REDACTED] response to the CMA's information request [REDACTED].

²⁵³⁷ Responses to the CMA's information requests [REDACTED].

affected. For example, one customer said that it considered other factors, such as scalability and flexibility.²⁵³⁸

7.708 We also asked customers whether any differences in using Microsoft software products²⁵³⁹ on Azure compared to other public clouds affected their choice of supplier; some of these customers said that it did,²⁵⁴⁰ and some customers said that it did not.²⁵⁴¹

Selection factors for VDI customers

7.709 Because the primary use case for some of the Microsoft software products (namely Windows 10/11, Microsoft 365 and Visual Studio) on the public cloud is as part of VDI instances, we also identified additional customers active in industries that are more likely to use VDI, and some of which were randomly selected. The questions we asked them were similar to those we asked other customers but were specifically framed in the context of VDI. We supplemented this customer evidence by evidence gathered from VDI providers, which we have set also out in this section.

7.710 We asked customers specifically about their choice of cloud provider for their VDI workloads. Almost all customers we heard from responded that they use VDI.²⁵⁴² Some of these customers host VDI on-premises.²⁵⁴³ Considering the customers we heard from that host VDI on public cloud:

- (a) many customers host their VDI network on Azure;²⁵⁴⁴
- (b) some customers host their VDI network on both Azure and AWS;²⁵⁴⁵
- (c) a few customers host their VDI network both on Azure and on-premises;²⁵⁴⁶ and
- (d) one customer did not host any of its VDI network on Azure, hosting across both AWS and IBM.²⁵⁴⁷

7.711 We asked some customers why they chose the public cloud to host their VDI workloads.

²⁵³⁸ [REDACTED] response to the CMA's information request [REDACTED].

²⁵³⁹ The question referred to Windows Server, SQL Server, Windows 10/11, Microsoft Office, Microsoft 365 and Active Directory/Azure AD specifically.

²⁵⁴⁰ Responses to the CMA's information requests [REDACTED].

²⁵⁴¹ Responses to the CMA's information requests [REDACTED].

²⁵⁴² Responses to the CMA's information requests [REDACTED].

²⁵⁴³ Responses to the CMA's information requests [REDACTED].

²⁵⁴⁴ Responses to the CMA's information requests [REDACTED].

²⁵⁴⁵ Responses to the CMA's information requests [REDACTED].

²⁵⁴⁶ Responses to the CMA's information requests [REDACTED].

²⁵⁴⁷ [REDACTED] response to the CMA's information request [REDACTED].

- (a) Many customers that responded mentioned costs²⁵⁴⁸ and whether the VDI provider is an existing supplier²⁵⁴⁹ as important factors.
- (b) Some customers that responded mentioned price and non-price differences relating to Microsoft products as reasons for choosing their public cloud provider:
 - (i) Some of these customers mentioned choosing Azure as they already use Microsoft software;²⁵⁵⁰
 - (ii) Some mentioned choosing Azure to enable Windows 10/11 multisession use (a capability not available on other public clouds);²⁵⁵¹
 - (iii) Some mentioned licensing costs or lower costs for Microsoft products.²⁵⁵²

7.712 We asked customers who reported use of Azure for their VDI workloads about any role that Microsoft's licensing practices played in their choice of provider for VDI:

- (a) A few customers said that they chose Azure to make effective use of their existing Microsoft software licences and Azure commitments;²⁵⁵³
- (b) However, one customer said that the Azure VDI offering was superior to competitors, and another said that there is no difference in the costs of licensing Microsoft software on Azure and AWS.²⁵⁵⁴

7.713 We also asked some further customers (who were not selected based on where they deployed their VDI network) whether any factors associated with the relevant 'client' Microsoft software products influenced their decision around which cloud provider to host their VDI workloads.

- (a) Some customers mentioned at least one factor associated with relevant Microsoft products including:
 - (i) Some customers reported preferring to host VDI workloads in the same public cloud as other Microsoft products;²⁵⁵⁵
 - (ii) One customer referred to the importance of being able to run the Windows 10/11 multisession feature;²⁵⁵⁶

²⁵⁴⁸ Responses to the CMA's information requests [X].

²⁵⁴⁹ Responses to the CMA's information requests [X].

²⁵⁵⁰ Responses to the CMA's information requests [X].

²⁵⁵¹ Responses to the CMA's information requests [X].

²⁵⁵² Responses to the CMA's information requests [X].

²⁵⁵³ Responses to the CMA's information requests [X].

²⁵⁵⁴ Responses to the CMA's information requests [X].

²⁵⁵⁵ Responses to the CMA's information requests [X].

²⁵⁵⁶ [X] response to the CMA's information request [X].

- (iii) One customer referred to the fact that it did not have to pay the VDA²⁵⁵⁷ per-user licence on Azure.²⁵⁵⁸

7.714 We asked customers whether there was any benefit to hosting other workloads on the same cloud as their VDI. Customers' responses depended on their use cases for VDI.

- (a) Many of the customers that responded reported benefits (or hypothetical benefits) to having some workloads located in the same public cloud as their VDI network.²⁵⁵⁹

- (i) Workloads which some customers we contacted mentioned as preferable to host in the same public cloud as their VDI network included data,²⁵⁶⁰ core line of business applications,²⁵⁶¹ key cloud-based workloads²⁵⁶² and some customers reported unspecified applications or workloads.²⁵⁶³

- (ii) The primary reason for customers we heard from hosting workloads in the same public cloud as their VDI network was latency.²⁵⁶⁴ For example, one customer that used a VDI to run intensive workloads said that the underlying storage and compute must sit as close as possible to mitigate latency and performance issues.²⁵⁶⁵

- (iii) One customer explained that, while there are benefits to hosting other workloads in the same public cloud as their VDI network, that it would be the customer's overall public cloud strategy which would determine its location of VDI workloads, not the other way around.²⁵⁶⁶

- (b) However, some customers did not consider this in deciding where to deploy other cloud workloads or did not consider it beneficial to deploy additional workloads in the same public cloud as their VDI network.²⁵⁶⁷

²⁵⁵⁷ The VDA license refers to the Virtual Desktop Access license referred to above.

²⁵⁵⁸ [REDACTED] response to the CMA's information request [REDACTED].

²⁵⁵⁹ Responses to the CMA's information requests [REDACTED]. We did not ask customers about the size and/or value of these workloads, and it may be for some customers that these are a small portion of their public cloud spending.

²⁵⁶⁰ Responses to the CMA's information requests [REDACTED].

²⁵⁶¹ [REDACTED] response to the CMA's information request [REDACTED].

²⁵⁶² [REDACTED] response to the CMA's information request [REDACTED].

²⁵⁶³ Responses to the CMA's information requests [REDACTED].

²⁵⁶⁴ Responses to the CMA's information requests [REDACTED].

²⁵⁶⁵ [REDACTED] response to the CMA's information request [REDACTED].

²⁵⁶⁶ [REDACTED] response to the CMA's information request [REDACTED].

²⁵⁶⁷ Responses to the CMA's information requests [REDACTED].

VDI and cloud providers' views on VDI

- 7.715 We also asked VDI providers²⁵⁶⁸ about where their customers choose to deploy their VDI networks in the public cloud and which factors influence customers' choice of cloud for VDI services.
- 7.716 We asked a VDI provider about where customers chose to host their VDI workloads. It said that the vast majority of its customers who host their VDI workloads on the public cloud did so on Azure.²⁵⁶⁹
- 7.717 We asked the same VDI provider why customers choose to deploy their VDI network on Azure, and whether customer choice is affected by licensing restrictions in relation to the Microsoft software products. It said that in its experience, customer choice of public cloud for VDI workloads is influenced by three factors:
- (a) where the customer's public cloud workloads are already based;
 - (b) attractive commercial licensing conditions on Azure, including the availability of multi-session desktops; and
 - (c) licensing conditions in particular with respect to Microsoft 365.²⁵⁷⁰
- 7.718 Another VDI provider said that licensing restrictions make Windows more expensive to run on AWS and Google Cloud. It said that customers are unable to use Google Cloud on 80-90% of opportunities currently in the market because of the restrictions on licensing for Office/Microsoft 365.²⁵⁷¹
- 7.719 The same VDI provider said that it is preferable for customers to host other workloads in the same cloud as their VDI workloads, describing this as the 'desktop gravity', meaning that wherever a customer's desktop goes it drags everything else with them. As such, in this VDI provider's view, not allowing Microsoft 365 to run on Google and forcing it to run on Azure, drags every workload on to Azure.²⁵⁷² The VDI provider also said that it has seen around 30-50 customers that moved their entire data set into Azure so that their data sets would be closer to their virtual desktops.²⁵⁷³
- 7.720 AWS and Google have submitted that they have lost customers due to Microsoft's restrictions with respect to the 'client-side' products.

²⁵⁶⁸ VDI providers are providers of software which runs a customer's VDI network on the public cloud, on premises or in the private cloud

²⁵⁶⁹ Note of meeting with [REDACTED]

²⁵⁷⁰ Note of meeting with [REDACTED].

²⁵⁷¹ Note of meeting with [REDACTED].

²⁵⁷² Note of meeting with [REDACTED].

²⁵⁷³ Note of meeting with [REDACTED].

- (a) AWS submitted that it had lost [REDACTED] USD due to customers that had either left the provider's VDI solutions, or expressed interest in using the services but ultimately chose not to due to licence restrictions in 2022 and part of 2023.²⁵⁷⁴ This provider's total revenue for its VDI solutions was around [REDACTED] USD in 2022.²⁵⁷⁵
- (b) Google submitted that as a result of Microsoft's restrictions and the resulting financial penalty imposed on GCP customers (in particular from the inability to BYOL post-2019 Windows licenses, additional costs for ESUs and VDA add on licenses and the inability to purchase Microsoft 365 through the SPLA or BYOL), GCP's ability to attract Microsoft workloads is significantly impeded.²⁵⁷⁶

7.721 We asked cloud providers for internal documents relating to the market landscape for the supply of VDI/DaaS and/or the forecast of the growth opportunity for VDI/DaaS in the cloud.²⁵⁷⁷ We also asked Google to provide any internal documents relating to its decision not to offer its own first party VDI service and the motivations behind that decision. In response to these requests Google submitted internal documents which detail its business plans and potential challenges with respect to the provision of VDI services:

- (a) An internal Google document suggested that one reason it has not been able to offer a first party VDI service [REDACTED]. The same Google internal document also noted that Windows Server [REDACTED] is significantly more expensive than Azure (via Azure Hybrid Benefit) and some VDI workloads such as Microsoft 365 apps are not eligible for resale [REDACTED] and are not BYOL eligible, limiting Google's ability to enter this market as a first party VDI provider.²⁵⁷⁸
- (b) The same Google internal document outlined that while it is working with partners [REDACTED] that this workaround:
 - (i) adds friction for customers;
 - (ii) requires them to repurchase licenses for software which they already own; and
 - (iii) will only be available until 2025.²⁵⁷⁹

²⁵⁷⁴ AWS' response to the CMA's information request [REDACTED].

²⁵⁷⁵ AWS' response to the CMA's information request [REDACTED].

²⁵⁷⁶ Google's response to the CMA's information request [REDACTED]. In support of its response Google provided the internal documents summarised below.

²⁵⁷⁷ We note that the majority of documents responsive to this question are summarised as part of the Importance of the input section above.

²⁵⁷⁸ Google's response to the CMA's information request [REDACTED].

²⁵⁷⁹ The availability of third party SPLA for use on listed provider's public clouds is detailed above. Google's response to the CMA's information request [REDACTED].

- (c) We have also seen Google internal documents which illustrate that Google has lost out on several IaaS deals in relation to the provision of cloud infrastructure services due to Microsoft's licensing practices, even going so far as to describe it as a 'blocker' to its business.²⁵⁸⁰ One internal document provided a sample list of lost VDI deals, where out of [REDACTED] deals presented, [REDACTED] were reported as being lost directly due to Microsoft's licensing restrictions:
- (i) [REDACTED].
 - (ii) [REDACTED].²⁵⁸¹
- (d) Google has also submitted [REDACTED] as a non-exhaustive set of examples where Google has failed to win a customer contract for their VDI workloads on Google Cloud due to licensing restrictions related to Microsoft Office.²⁵⁸²
- (i) One deal valued at [REDACTED] USD was said to have been lost due to client deciding to move their VDI to Azure due to Microsoft licensing.²⁵⁸³
 - (ii) A migration deal to GCP was blocked due to the customer requiring Microsoft 365 on Virtual Desktop Access (VDA). It described that Microsoft 365 is not allowed to be installed on VDA and therefore 'it is a full blocker to migration.'²⁵⁸⁴

7.722 As outlined above, similarly in this instance we consider it appropriate to place limited weight on these examples. While they illustrate the scope for the licensing practices to give rise to an effect on Google's competitive offering, we consider it appropriate to attach greater weight to the customer evidence we have gathered directly and in a consistent way.

7.723 We asked AWS about the effect of customers being able to BYOL some Microsoft 365 licences to Amazon Workspaces. AWS said that the ability to host Microsoft 365 on Workspaces has [REDACTED]. It said that it was not possible to host Microsoft 365 for many years and that the ability to use the Microsoft product was key while customers made their cloud journey, [REDACTED].²⁵⁸⁵

7.724 Microsoft has submitted that it was a misconception that Windows 10 and 11 [not being available through the SPLA] limits competition because it is important for DaaS.²⁵⁸⁶ Microsoft said that DaaS is a very small portion of the market for cloud services and that, in addition, the SPLA had always been a programme to offer

²⁵⁸⁰ Google's response to the CMA's information request [REDACTED].

²⁵⁸¹ Google's response to the CMA's information request [REDACTED].

²⁵⁸² Google's response to the CMA's information request [REDACTED].

²⁵⁸³ Google's response to the CMA's information request [REDACTED].

²⁵⁸⁴ Google's response to the CMA's information request [REDACTED].

²⁵⁸⁵ Note of meeting with AWS [REDACTED].

²⁵⁸⁶ Microsoft's submission to the CMA [REDACTED].

server products to be used to build cloud services, whereas client products are designed to be installed on individual PCs for individual use.²⁵⁸⁷

Customer examples from AWS and Google

- 7.725 Since the PDR, we have requested further examples of internal documents from AWS and Google that contain discussions about prospective or current customers that they were not able to compete for due to Microsoft's licensing practices, as well as (as presented above) requesting evidence on customers they have been able to win due to adopting counterstrategies to Microsoft's licensing practices. This evidence has been added to the examples previously provided by AWS and Google and both have been set out below.²⁵⁸⁸
- 7.726 AWS and Google provided customer examples that they said show the effect of Microsoft's licensing practices.²⁵⁸⁹ Both provided anonymised customer examples and internal documents, and Google additionally provided hypothetical customer examples. We set out an overview of these customer examples and internal documents below.
- 7.727 Before setting out this evidence, we make the following observations which are relevant to its interpretation:
- (a) The examples have been selected by Google and AWS, so we are unable to assess how representative they are of typical customer experiences. Nevertheless, we consider the examples provide evidence on the impact of Microsoft's licensing practices for some AWS and Google customers.
 - (b) The examples and documents submitted by AWS and Google include some companies that operate outside the UK + EEA. As there are similarities between different geographic markets, such as the same products, similar market participants, and similar practices we consider it illustrative of how Microsoft's practices can affect customers.
- 7.728 In some of the examples, we are unable to assess the extent to which licensing costs were passed on to customers. AWS submitted [a number of] [REDACTED] anonymised customers that it said illustrated the harmful impact of Microsoft's licensing practices. AWS said these customers abandoned plans to migrate workloads to AWS, chose to migrate workloads away from AWS or incurred significant additional expenses as a result of Microsoft's licensing practices.²⁵⁹⁰ For example, AWS said that one of these customers, a [REDACTED], was previously a

²⁵⁸⁷ Microsoft's submission to the CMA [REDACTED].

²⁵⁸⁸ We note that some of these documents have been produced during the course of the investigation, and have therefore been interpreted accordingly.

²⁵⁸⁹ Google's response to the CMA's information request [REDACTED]; AWS' submission to the CMA [REDACTED].

²⁵⁹⁰ AWS' submission to the CMA, [REDACTED].

customer of AWS but migrated its workloads to Azure because it needed to use Windows Server but did not want the additional expense of having to purchase new licences for use on AWS when it could use its existing licences on Azure.²⁵⁹¹

7.729 AWS also provided documents [REDACTED] which it said clearly indicate the challenges AWS faces when competing for workloads that are impacted by Microsoft's licensing restrictions.²⁵⁹²

- (a) Some of these documents [REDACTED] appear to consider AWS' inability to offer multi-session [REDACTED].²⁵⁹³ [REDACTED].²⁵⁹⁴ AWS provided another document which shows that, as AWS is unable to offer multi-session, AWS would have needed to incur a negative margin of [REDACTED]% to meet a customer's requirements.²⁵⁹⁵
- (b) One document discusses the implications of a change to Microsoft's SPLA program on independent software vendors (ISVs) and states that changes to Microsoft licensing practices will result in higher costs and reduced flexibility for ISVs [REDACTED] and that AWS will need to [REDACTED].²⁵⁹⁶
- (c) Some of these documents [REDACTED] appear to discuss discounts AWS is [REDACTED].²⁵⁹⁷ For example, one document suggests that the inability for customers to BYOL meant that [REDACTED], and AWS therefore had less flexibility on the discount level they could offer while [REDACTED].²⁵⁹⁸ Another one of these documents suggests that AWS has to increase its discount from [REDACTED] to [REDACTED] to be competitive with Microsoft, and that not taking any mitigating measures would risk AWS losing [REDACTED] of its revenue from this customer, and potentially future SaaS revenue.²⁵⁹⁹
- (d) One document presents customer cases studies of where customers incur additional costs from running the same Windows Server and SQL Server workloads on AWS compared to Azure as a direct result of Microsoft's licensing restrictions, by showing a direct comparison of AWS and Microsoft Azure's costs without any optimisation. The case studies suggest that AWS' compute optimisation [REDACTED].²⁶⁰⁰ However, AWS clarified that since Azure offers cost optimisation to all of its customers, for the purposes of this issue it is incorrect to compare Azure's non-optimised costs in the calculations with

²⁵⁹¹ AWS' submission to the CMA [REDACTED].

²⁵⁹² AWS' response to the CMA's information request [REDACTED].

²⁵⁹³ AWS' response to the CMA's information request [REDACTED].

²⁵⁹⁴ AWS' response to the CMA's information request [REDACTED].

²⁵⁹⁵ AWS' response to the CMA's information request [REDACTED].

²⁵⁹⁶ AWS' response to the CMA's information request [REDACTED].

²⁵⁹⁷ AWS' response to the CMA's information request [REDACTED].

²⁵⁹⁸ AWS' response to the CMA's information request [REDACTED].

²⁵⁹⁹ AWS' response to the CMA's information request [REDACTED].

²⁶⁰⁰ AWS' response to the CMA's information request [REDACTED].

AWS' optimised costs, because the same reduction could apply to the customer's costs on Azure.²⁶⁰¹

- 7.730 Google provided a number of [REDACTED] anonymised customer examples that it said show where Microsoft's licensing practices caused Google to lose the relevant customer opportunity. For example, Google said that one of these customers, [REDACTED], was forced to move to Azure because the inability to BYOL Windows Server licences made the cost of moving to a non-Azure cloud [REDACTED]. Google said that while it was willing to offer aggressive discounts, it could not permanently take on the additional costs of licences that the customer would incur on an ongoing basis. Google also said that this customer [REDACTED].²⁶⁰²
- 7.731 Google provided documents [REDACTED] which it said are examples of where its ability to compete has been impaired by Microsoft's licensing practices. Google submitted that while in many of these cases it offered additional discounts, credits [REDACTED], it nevertheless lost the deal.²⁶⁰³
- (a) Many [REDACTED] of these documents, which are generally internal email chains between Google employees, appear to show that [REDACTED] compared to [REDACTED] as a result [REDACTED].²⁶⁰⁴ In some of these cases, [REDACTED].²⁶⁰⁵
- (i) For example, one of these documents sets out [REDACTED], and some [REDACTED] of these explicitly mention the cost of Microsoft licensing as a challenge to winning the opportunity. For instance, one lost opportunity specifies that Google's cloud SPLA was [REDACTED] times as expensive as Azure's.²⁶⁰⁶
- (ii) Another one of these [REDACTED] indicated that the customer would be willing to pay a [REDACTED]% premium to use GCP since it was their primary and preferred cloud. However, the customer found that AHB meant that Google was several [REDACTED] times more expensive than Azure and it ultimately decided to migrate their application to Azure.²⁶⁰⁷
- (b) Several of these documents set out that customers are going through third party ISVs for Microsoft licences [REDACTED]²⁶⁰⁸ [REDACTED].²⁶⁰⁹ Nevertheless, some of

²⁶⁰¹ AWS' response to the CMA's information request [REDACTED].

²⁶⁰² Google's response to the CMA's information request [REDACTED].

²⁶⁰³ Google's response to the CMA's information request [REDACTED].

²⁶⁰⁴ Google's response to the CMA's information request [REDACTED].

²⁶⁰⁵ Google's response to the CMA's information request [REDACTED]. In other cases, it was unclear whether Google ultimately won or lost the customer.

²⁶⁰⁶ Google's response to the CMA's information request [REDACTED].

²⁶⁰⁷ Google's response to the CMA's information request [REDACTED]; Google's response to the CMA's [REDACTED].

²⁶⁰⁸ Google's response to the CMA's information request [REDACTED].

²⁶⁰⁹ Google's response to the CMA's information request [REDACTED].

these documents suggest that the changes to Microsoft's licensing terms mean [REDACTED].²⁶¹⁰ [REDACTED].²⁶¹¹

- (c) Some [REDACTED] of these documents appear to show that Google lost customers due to Microsoft's licensing restrictions on Microsoft Office (which is used on VDI).²⁶¹² For example, one document identifies Google's (in)ability to offer Microsoft Office SPLA licences as [REDACTED].²⁶¹³ We have not received any internal documents which report on the success of Google's license manager solution; therefore we are unsure whether this has been introduced and may be having an impact on Google's ability to compete for customers' VDI workloads.
- (d) Some [REDACTED] of these documents appear to relate to other Microsoft licensing terms that contributed to Google not winning the customer. This [REDACTED]²⁶¹⁴ [REDACTED].²⁶¹⁵

7.732 Separately, Google also provided an internal document discussing the licensing costs of a customer [REDACTED].²⁶¹⁶ [REDACTED].²⁶¹⁷

7.733 Google also produced five hypothetical customers and calculated the amounts they would pay for hosting workloads on Azure compared to its own cloud. Google submitted that these examples were produced by employees that regularly assess real customers' cloud services needs and Microsoft's licensing position. Google said that the examples therefore reflect real-world scenarios of typical traditional enterprise customers running Microsoft workloads on-premises that its sales team encounters.²⁶¹⁸ According to Google, total cloud costs were [several] [REDACTED] times more expensive on its cloud compared to on Azure, despite the IaaS generally being cheaper on the non-Azure cloud.²⁶¹⁹

7.734 In its response to the PDR, Microsoft submitted that there was a lack of concrete examples that demonstrate harm and that there was no material expression of customer concern according to UK enterprises themselves.²⁶²⁰ Microsoft also said that the few examples given are almost entirely ([REDACTED]) [REDACTED] companies and therefore not probative of partial foreclosure in the UK that benefits Azure.²⁶²¹

²⁶¹⁰ Google's response to the CMA's information request [REDACTED].

²⁶¹¹ Google's response to the CMA's information request [REDACTED].

²⁶¹² Google's response to the CMA's information request [REDACTED].

²⁶¹³ Google's response to the CMA's information request [REDACTED].

²⁶¹⁴ Google's response to the CMA's information request [REDACTED].

²⁶¹⁵ Google's response to the CMA's information request [REDACTED].

²⁶¹⁶ Google's submission to the CMA [REDACTED].

²⁶¹⁷ Google's response to the CMA's information request [REDACTED].

²⁶¹⁸ Google's response to the CMA's information request [REDACTED].

²⁶¹⁹ Google's submission to the CMA [REDACTED].

²⁶²⁰ [Microsoft's response to the CMA's PDR dated 28 January 2025](#), paragraph 192-194.

²⁶²¹ Microsoft's submission to the CMA [REDACTED].

- 7.735 We consider that the examples are illustrative of how Microsoft's practices can affect customers. The documents in particular provide some concrete examples of Microsoft's licensing practices having an impact on AWS and Google's competitiveness. While AWS has only been able to provide a relatively limited number of documents, Google provided a number that show that higher costs as a result of the SPLA or Microsoft's licensing restrictions were a factor in Google losing a number of opportunities and it being a concern in others. In one example provided, Google was able to win a customer by offering significant free credits to offset costs from Microsoft's licensing practices.
- 7.736 Overall, we consider that we can place some weight on the examples, and particularly those examples for which we have seen documentary evidence, as they demonstrate how some customers' choices are affected by Microsoft's licensing practices. In particular, these examples show that the concerns raised are not hypothetical. However, given the caveats set out above, we consider that the examples cannot be viewed as being representative of overall customer behaviour, and we have not treated them as such. Therefore, we have not used these examples in assessing the extent of the effect of Microsoft's licensing practices.

Our assessment of customer evidence and customer examples

- 7.737 Customer evidence shows that software licensing terms likely have an impact on customers' choice of cloud provider. It indicates that the cost, ease and/or ability to use Microsoft software licences are an important selection factor for many customers, and some particularly consider the ability to make use of their existing investment in licences in their decision. These customers are therefore more likely to choose Azure for running at least their Microsoft workloads, and possibly more widely.
- 7.738 The evidence also shows that pre-existing use of Microsoft software, and the associated skills developed, were very important selection factors for many Azure customers. Nevertheless, many of these customers indicated that licensing terms were also a consideration in their decision-making process. Further, we note that even for customers that would have chosen Azure due to their pre-existing use of Microsoft regardless of licensing terms, the licensing terms likely still influence future decision making. In particular, for Azure customers considering switching, licensing terms may result in an additional friction to doing so.
- 7.739 The evidence suggests that most customers are not aware of quality differences in their use of Microsoft software when choosing a public cloud service provider. But for some customers, non-price factors in relation to use of Microsoft software may influence choice of cloud; for example, the availability of security updates may impact a subset of customers, and for some customers considering VDI workloads the availability of Windows multisession capability is important. However, we

consider that for most customers quality factors alone do not appear to influence their decision making.

7.740 VDI providers' views are consistent with Microsoft's licensing practices having an effect on customer choice of cloud with respect to VDI workloads.

7.741 The examples and internal documents provided by AWS and Google show some instances of where Microsoft licensing practices resulted in AWS and Google being competitively disadvantaged. Google was able to provide a greater number of examples than AWS. As we are unable to assess how representative the documents are of customers' choices being impacted by the licensing practices, we have placed limited weight on them – other than in demonstrating that the concern is not purely hypothetical - and place greater weight on the customer evidence we have gathered directly.

Our assessment

7.742 We have assessed whether Microsoft's licensing practices may have an impact on the competitiveness of AWS and Google, and whether they have the ability to absorb or counteract the effect.

7.743 In considering whether they can absorb or counteract the effect of Microsoft's licensing practices, we have considered how AWS' and Google's competitive offerings compare to Microsoft's. Evidence from both AWS and Google is consistent with high input costs resulting in them charging a higher price and having less competitive offerings. One provider, [X], has passed through changes in the input cost, while the other, [X], has historically absorbed increases.

7.744 Due to a number of limitations and methodological difficulties, we have not used a single set of estimates for profit margins (and more generally consider margins to be less determinative of partial foreclosure than some of the other evidence). However, all estimates of mark-ups and margins for workloads associated with Windows Server appear to be below the margins larger cloud providers generate on their overall cloud business. This is true even for the upper bound estimates and the lower bound estimates are significantly below this level. The available evidence on mark-ups and margins is therefore consistent with partial foreclosure.

7.745 Average usage of Windows Server and SQL Server is much higher on Azure than on AWS and GCP, suggesting that customers disproportionately choose to deploy workloads using Windows Server and SQL Server on Azure. Given the evidence that customers are unlikely to switch from Windows Server and SQL Server (as set out in our assessments of market power) and given the level of the disparity, there is a high likelihood that the differences in relative usage are at least partially driven by Microsoft's licensing practices.

- 7.746 We considered Microsoft's submission that, if its licensing practices allow it to win more customer workloads, it may be expected that Azure will win a disproportionate share of large customers because large enterprise customers tend to have on-premises licences that would benefit from AHB. However, we note that it is not necessarily the case that large customers are more likely to have pre-existing licences that qualify for AHB for the relevant Microsoft software. Digital natives (which are unlikely to have pre-existing licences that qualify for AHB for the relevant Microsoft software) may also have high cloud spend. Further, our analysis on relative usage, as discussed in the paragraph above, provides a more direct assessment of the extent to which Azure is winning workloads that use Microsoft software.
- 7.747 Customers' views on using Microsoft software on public cloud show that they perceive AWS and Google to be more expensive than Azure, and that some (although not all) customers consider licensing factors in making their choice of cloud.
- 7.748 Finally, Microsoft's licensing practices with respect to Office, Microsoft 365, Windows Desktop and Visual Studio are likely to affect competition for VDI workloads on the public cloud. This view is supported by evidence from VDI providers and customer evidence which outlines that licensing and non-price factors are important in customer choice of cloud. Competition for VDI workloads may be the lower bound for any effect on competition with respect to VDI because some customers highlighted benefits to hosting other workloads in the same public cloud as their VDI network; therefore their choice of cloud for VDI may influence choice of cloud for other workloads.

Conclusions on the impact on AWS and Google's competitiveness from Microsoft's licensing practices

- 7.749 Microsoft's licensing practices in relation to Windows Server, and Windows Server and SQL Server together are adversely impacting AWS' and Google's competitive offerings. The impact of Microsoft's licensing practices on VDI workloads contributes to the adverse impact on AWS' and Google's competitive offerings.

Summary of our assessment and conclusions

Ability

- 7.750 Microsoft has significant market power in relation to each of Windows Server, SQL Server, Windows 10/11, Visual Studio and its productivity suites. Microsoft has moderate to high market shares in all of these products and customers are generally unable or unwilling to switch away from using this software.

- 7.751 AWS and Google have had limited success in moving customers away from the Microsoft software products through modernisation programmes. These programmes have not significantly affected Microsoft's market power in relation to these software products.
- 7.752 Cloud providers supply customers with various bundles of services that differ in value. The broader the set of services a customer buys, the less significant Microsoft software will be in relative terms and its price may therefore play a less significant part in the customer's cloud provider choice.
- 7.753 The extent to which customers make aggregated purchasing decisions varies significantly. As such, it is appropriate to consider a range of possible values for how big a share of the cost of cloud services they purchase Microsoft software accounts for. We have placed more weight on values towards the middle of the range for the cost base.
- 7.754 Windows Server alone, and when combined with SQL Server, accounts for a material proportion of costs for all sizes of customer and, for values in the middle of the range (ie between total spend on Azure and spend on Windows Server VMs (or Windows Server and SQL Server)) the proportions are very high.
- 7.755 VDI allows customers to use Windows 10/11, Visual Studio and the Microsoft productivity suites on cloud. VDI represents a small proportion of customers' spend on public cloud services. However, Windows 10/11 and the Microsoft productivity suites are particularly important to VDI and Visual Studio is important to VDI for some customers. As such, in combination with Windows Server and SQL Server workloads, VDI workloads add to the importance of Microsoft software as an input to cloud services.
- 7.756 In terms of quality differences between using the relevant Microsoft software products on Azure compared to AWS and Google, these are less important than cost. Most customers are not aware of quality differences between using Microsoft software on Azure compared to AWS or Google, because they have not made this comparison. However, some customers identified quality differences. For these customers, Windows Server, SQL Server and/or Windows 10/11 may be particularly important in terms of affecting the overall quality or attractiveness of AWS' and Google's as alternatives to Azure.
- 7.757 We considered whether the importance of Microsoft software as an input is likely to change in the future.²⁶²² The size of the on-premises IT market is large, and Microsoft has a large share of the server OS market, which indicates that there is a substantial pool of potential demand for Windows Server that could be migrated

²⁶²² As noted above, this assessment is only relevant to whether the importance of the Microsoft software inputs will increase or reduce in the future relative to customers total Azure spend.

to the public cloud. We have not found evidence that the rate of migration will slow. There is limited switching away from Microsoft software which means if Microsoft's licensing practices stay as they are these would continue to influence customers' choices of cloud provider.

- 7.758 Other (non-Microsoft software) workloads are growing, potentially faster than Microsoft workloads. However, even if, in the future, Microsoft workloads make up a lower proportion of customers' total spend they are likely to remain an important input in the foreseeable future.
- 7.759 In summary, Microsoft has the ability to partially foreclose AWS and Google through its software licensing practices.

Incentive

- 7.760 Microsoft sets a high input price for AWS and Google for hosting Windows Server and SQL Server. AWS and Google, Microsoft's closest competitors, are targeted specifically as Listed Providers, which do not include other smaller cloud providers.²⁶²³
- 7.761 The price that AWS and Google face for hosting Windows and SQL Server has increased substantially since 2018. The input cost for AWS and Google is higher than the customer-facing price that Microsoft charges its own customers that qualify for AHB to use Windows Server on Azure. For SQL Server, the input cost for AWS and Google is higher than Microsoft's PAYG customer-facing price.
- 7.762 Many customers are unlikely to switch away from the relevant Microsoft software products to use other software on AWS and Google. This means that the risk to Microsoft of losing software customers (and associated revenues) by raising AWS' and Google's costs is lower, so Microsoft has a greater incentive to partially foreclose them.
- 7.763 Microsoft has a high share of supply in both IaaS and PaaS, and IaaS in particular is a concentrated market, and so Microsoft is an important competitive constraint on AWS and Google. This means that customers that are deterred from choosing AWS or Google are likely to be captured by Microsoft. This may be more likely as customers switching away from AWS because of the cost of Microsoft software would likely face the same issue on Google, and vice versa.
- 7.764 IaaS and PaaS are high growth markets, and IaaS has material economies of scale. As such, the potential for Microsoft to win customers in IaaS and PaaS and position itself to capture future growth and realise further economies of scale may act as an additional incentive for it to partially foreclose rivals.

²⁶²³ As set out in our description of the timeline of licensing practices, Microsoft submitted that [X].

7.765 In summary, Microsoft has the incentive to partially foreclose AWS and Google through its software licensing practices and these licensing practices reflect that incentive.

Effect

7.766 We have considered whether Microsoft's licensing practices have the effect of reducing competition in cloud services markets by materially disadvantaging AWS and Google, who consequently compete less effectively. This is not a binary consideration of whether AWS and Google can compete but rather a consideration of whether they compete materially less effectively than they otherwise would and whether this harms competition in the relevant markets.

7.767 AWS and Google are being materially disadvantaged by Microsoft's licensing practices in relation to Windows Server and SQL Server. They pass through at least some of the input costs of Microsoft software, and customers generally perceive them to be more expensive than Microsoft. This evidence primarily speaks to differences in offerings for customers with pre-existing Microsoft software licences for Windows Server or SQL Server that qualify for AHB. For customers that do not have licences that qualify for AHB, AWS and Google have mostly similar list prices for certain instances of Windows Server VMs as Azure.

7.768 A comparison of AWS and Google's profit margins is not determinative of partial foreclosure, due to both conceptual difficulties in interpreting margin levels and practical difficulties in estimating them in this particular case, notably the uncertainty identified above around how customers purchase sets of cloud services. We have therefore considered ranges of estimates of AWS' and Google's margins. We consider these ranges to be consistent with partial foreclosure.

7.769 Windows Server and SQL Server are used disproportionately on Azure compared to on AWS and Google. We consider that the very large difference is at least partly because some customers' choice of cloud is influenced by Microsoft's licensing practices.

7.770 AWS and Google have counterstrategies in place including modernisation of workloads away from Microsoft software, offering credits, optimisation and targeting customers based on their software mix. However, these have had a limited effect and have not sufficiently mitigated the effect of Microsoft's licensing practices.

7.771 Most customers do not perceive quality differences with respect to running Microsoft software on AWS' or Google's clouds, compared to running them on Azure. However, some customers identify access to security updates and the

availability of Windows multisession capability for VDI workloads as important, which are advantages to using Azure.

- 7.772 Microsoft does not make Windows 10/11 and Microsoft 365 available to AWS and Google through their respective SPLAs, and customers cannot use the BYOL route to deploy their existing software licences on AWS' or Google's clouds (with the exception of certain Microsoft 365 licences on Amazon Workspaces). [X]. These restrictions mean that AWS' and Google's competitive offerings are directly affected by Microsoft's practices in relation to these software products. These products are important inputs to VDI offerings and, while VDI represents a relatively small proportion of usage of all cloud services, Microsoft's practices relating to these software products has an adverse effect on rivals' competitive offerings.
- 7.773 In terms of the overall impact of any partial foreclosure of AWS and Google: they are the most significant competitors to Microsoft by share of supply, both in IaaS and PaaS markets. In addition, the group of customers with pre-existing licences for Windows Server that qualify for AHB is significant in size and they are affected by the reduction of competition between Microsoft and AWS and Google. As such, the impact of Microsoft's licensing practices on AWS' and Google's competitive offerings is likely to affect a significant portion of the cloud services markets.
- 7.774 The size of the on-premises market is large, and there is a substantial pool of workloads that we expect to be migrated to the public cloud over the next few years. There is therefore the potential for this impact to increase as the transition from on-premises to the cloud provides opportunity for customers to assess their options.
- 7.775 In summary, Microsoft's licensing practices have the effect of reducing competition in cloud services markets by materially disadvantaging AWS and Google, who consequently compete less effectively.

Conclusion

- 7.776 Considering all the evidence in the round, Microsoft's licensing practices are adversely impacting the competitiveness of AWS and Google in the supply of cloud services, particularly in competing for customers that purchase cloud services which use the relevant Microsoft software as an input. As a result, Microsoft faces weaker competitive constraints from AWS and Google, its most significant competitors, which is reducing competition in cloud services markets.

8. Committed spend agreements

- We have assessed whether CSAs may lead to harm to competition by acting as a potential commercial barrier to switching and multi-cloud. We focussed on AWS and Microsoft based on their position as the two largest providers of cloud services.
- We have assessed the extent to which CSAs could lead to harming competition by either reducing the ability of smaller rival suppliers to compete for AWS' and Microsoft's existing customers, leading to weakening or marginalisation of those smaller suppliers; and/or reducing the incentive of rival suppliers to compete for each other's existing customers, leading to softening of competition between those suppliers.
- AWS' and Microsoft's CSAs cover a significant portion of the cloud infrastructure services markets and influence customers' choices in relation to workload allocation. However, in relation to the impact of CSAs on competition, our quantitative analysis indicates that the existence of AWS' and Microsoft's CSAs, as currently applied, do not adversely affect rivals' ability or incentive to compete against AWS and Microsoft.
- In view of the above, the way in which AWS and Microsoft currently apply CSAs does not constitute a feature leading to harm to competition in the provision of public cloud infrastructure services in the UK.
- Although we have not identified competition concerns arising from CSAs as currently applied, this does not mean that harm to competition could not materialise in the future. In particular, the link between sticky and contestable demand in AWS' and Microsoft's CSAs could harm competition if the share of sticky demand increases significantly and/or if AWS and/or Microsoft change the way their CSA discounts are applied by increasing the incentive of customers to concentrate their spend with them.

Introduction

- 8.1 Committed spend agreements or discounts (CSAs or CSDs) are agreements between a cloud provider and a customer in which the customer commits to spend a minimum amount across the cloud provider's cloud services over a period of years and in return, receives a percentage discount on its spend with that provider during those same years.²⁶²⁴
- 8.2 This chapter presents our analysis of the potential impact of these agreements on competition in connection with the supply of public cloud infrastructure services (cloud services) in the UK. We have examined these agreements as a potential commercial barrier to switching and multi-cloud.

²⁶²⁴ In this report, we use the terms CSA and CSD interchangeably to mean the same type of agreement.

- 8.3 Whilst Ofcom focused on discussing potential concerns arising from committed spend *discounts* in its Market Study, we considered the potential for competition concerns arising from committed spend *agreements*. This is because we consider committed spend *agreements* more accurately reflect the nature of the potential concern, which relates to the pricing structure of these agreements rather than the discount per se – in particular, the link that these agreements create between sticky and contestable demand, an element which is explained in our conceptual framework below. However, in order to be consistent with Ofcom’s terminology, we used the term ‘CSDs’ in our evidence collection. Therefore, with a few exceptions, the analysis of the evidence in the other sections of this chapter is framed by reference to ‘CSDs’.
- 8.4 We first set out our conceptual framework for conducting our assessment and then move to our analysis.
- 8.5 We focus mainly on AWS and Microsoft. We do so based on their position as the two largest providers of cloud services. By virtue of their positions, we consider that any impact on competition arising from their CSAs/CSDs is likely to be greater than any impact from CSAs/CSDs offered by Google or by smaller cloud providers.

Conceptual framework and analysis structure

- 8.6 The provision of discounts can clearly be beneficial to customers. However, when discounts are provided under certain conditions and/or are structured in certain ways, they can give rise to concerns about harm to competition.
- 8.7 Pricing structures that incorporate conditional discounts, such as CSAs, can be thought of as a form of price discrimination: while some customers pay lower prices if they meet conditions set by the supplier, others will pay higher prices if they do not meet those conditions.
- 8.8 Price discrimination more generally, and conditional rebates more specifically, is not always a cause for concern. However, one example of how a concern may materialise in practice is as follows:
- (a) a customer has some of its demand met by a supplier and the extent to which the customer can exercise effective choice over that demand is limited by factors such as lack of suitable alternatives or barriers to switching. We call this ‘sticky demand’ and we use this term to encompass different categories of demand, eg services that are ‘must-have’, services that are infeasible to switch away from, or services where high costs of switching make it difficult to switch away;

- (b) the customer also has a portion of demand that is more contestable: the customer would be willing and able to place that demand with an alternative supplier (we call this the ‘contestable demand’); but
- (c) the supplier of the ‘sticky demand’ imposes a condition such that the customer must place some or all of the contestable demand with them or otherwise pay higher prices (lose a discount) on the sticky demand.

8.9 The concern under these circumstances is that the prospect of paying a higher price for the sticky demand, brought about by the CSA, deters customers from considering alternative suppliers for their contestable demand. The incumbent supplier leverages its strong position over one portion of demand into a new segment where it would not otherwise have enjoyed the same strong position. Competition may be harmed to the extent that the conduct:

- (a) reduces the ability of smaller rival suppliers to compete for AWS’ and Microsoft’s existing customers, leading to weakening or marginalisation of those smaller suppliers, for example because they lose, or fail to achieve, economies of scale;²⁶²⁵ or
- (b) reduces the incentive of rival suppliers to compete for each other’s existing customers, leading to softening of competition between those suppliers. That is because these rival providers would make higher profits by competing for smaller portions of contestable demand than for the whole.²⁶²⁶

8.10 As a stylised example of how this concern could materialise in practice, assume a customer has 100 units of demand overall – 78 units represent sticky demand already placed with an incumbent provider and 22 units represent contestable demand. We also assume that the price before discounts is £1 for each unit. In this example, the customer has an existing CSA with the incumbent provider under which the customer is already meeting the spend target and is in receipt of a 10% discount. For example, if the spend target of the existing CSA was 75 units, the customer gets a 10% discount on all units purchased if it buys at least 75 units – the customer would achieve that in our example given it is already buying 78 units from the incumbent.


8.11 Now consider a scenario where the customer is offered a new CSA which gives an 11% discount on its entire spend (ie both sticky and contestable) on condition that it commits to buy at least 95 units from the incumbent provider.

²⁶²⁵ The scenario we considered is AWS’ and Microsoft’s CSAs reducing the ability of smaller rival suppliers to profitably compete for the contestable demand.

²⁶²⁶ The scenarios we considered are (i) AWS’ CSAs reducing the incentive of rival suppliers, including Microsoft, to compete for the whole contestable demand; (ii) Microsoft’s CSAs reducing the incentive of rival suppliers, including AWS, to compete for the whole contestable demand; or (iii) both.

- 8.12 Thinking about a rival that is competing for a share of the customer's 100 units of demand:
- (a) If the rival wins 5 units or fewer, the customer can still meet its 95-unit commitment to the incumbent provider.
 - (b) If the rival wins more than 5 units of the contestable demand, the customer would no longer be able to meet the 95-unit commitment to the incumbent and the incumbent's price would rise substantially – the customer would lose a 1 percentage point discount on all 95 units, including the sticky demand that the customer finds it difficult or infeasible to switch away. As a result, the discount the rival would have to offer rises sharply on that incremental unit to compensate the customer with an equivalent discount spread across only a very small number of units.
- 8.13 As shown in Figure 8.1 below, if the rival wins more than just six units, the unit discount it would have to apply to win them decreases in percentage terms as the number of units won increases. That is because the rival can 'spread' the compensation it has to give the customer (for losing its incremental 1% discount from the incumbent) over a larger base of units. For example, if the rival were to win 10 units, it would have to apply a unit discount of 20% to match the incumbent's price. If the rival wins 22 units, it would have to apply a unit discount of around 15%.


Figure 8.1: Discount that the rival must apply to win units of contestable demand



Source: CMA analysis.

- 8.14 As Figure 8.1 shows, in this example the incumbent supplier leverages its strong position over the 78 units of sticky demand into the 22 units of contestable demand – that is because the discount that the rival must apply is magnified by the compensation necessary to win the customer, bearing in mind that it will lose the incremental 1 percentage point discount from the incumbent.
- 8.15 Whether the CSA's impact on the ability and/or incentive of the rival to compete for the contestable demand will depend on the rival's profit margin:
- (a) The rival may not have the ability to profitably compete for more than five units if the discount structure is such that the rival would have to price below its own costs. For example, if offering a 15% discount on the 22 units would take the rival below its unit costs, then the rival will not be able to profitably compete for any of the contestable demand.
 - (b) The rival may not have the incentive to compete for more than five units, even if it would be profitable in absolute terms. This would be the case if, for example, the profit made from selling five units at a 11% discount is greater than the profit from selling 22 units at a 15% discount.
- 8.16 Figure 8.2 below shows a stylised example of what happens to the rival's total profits when assuming the rival's margin on the pre-discount unit price of £1 is 15%. Looking at the total profits the rival would make for different number of units of demand it wins, the rival is better off, ie would make higher profits, by winning five units (point B in Figure 7.2) rather than the whole 22 units (point C), or any number of units between five and 22. That is because there are not enough units for the rival to spread the compensation it has to give the customer for not getting the incremental discount from the incumbent. Therefore, while the rival has the ability to profitably win all the 22 units in this example, it does not have the incentive to compete for more than 5 units.

Figure 8.2: Total profit to rival for different units of contestable demand it wins



Source: CMA analysis.

- 8.17 We note the above is a stylised example and is not necessarily representative of the CSAs offered by AWS or Microsoft, nor is it intended to be reflective of the actual margins of rivals.
- 8.18 What this example shows is that, depending on the circumstances, an agreement involving a conditional discount creates a link between sticky and contestable demand that could have a negative impact on competition.
- 8.19 In order to examine whether the particular characteristics of CSAs used in the supply of cloud services are likely to harm competition, we have analysed the following issues:
- (a) The prevalence of CSAs. This provides an indication of the size of the segment that could be affected by these agreements.
 - (b) The extent to which CSAs affect customers' choices in relation to the allocation of workloads on public cloud. If these agreements in practice play a limited role in customers' choices on where to allocate workloads, then they would be less likely to harm competition.
 - (c) The extent to which the pricing structure of CSAs as currently applied has characteristics that may cause them to harm competition. In particular, we assess, in light of the particular structure of CSAs in cloud services, whether rivals would have the ability and incentive to compete against the incumbent's CSAs for the customers' contestable demand. We do so by drawing together the following factors which drive that assessment:
 - (i) The share of demand for the incumbent's services that is sticky. All else being equal, if a large share of the demand for services supplied by a

provider is sticky, then CSAs offered by that provider are more likely to harm competition. That is, such a large share would make it more likely that it will be uneconomical for the rival to overcome the incumbent's price increase on the sticky demand with a discount on the contestable demand.²⁶²⁷

- (ii) The relationship between discount rate and commitment in the incumbent's CSAs. Where customers receive relatively large discounts for relatively small commitments, the effect of CSAs is exaggerated.
 - (iii) The rival's profit margin. The lower a rival's margin, the lower the rival's ability and incentive to overcome any effect introduced by a CSA. Where rivals have low margins, they are more likely either to make a loss in absolute terms if they try to undercut the CSA or, where profitable, they are more likely to be deterred from undercutting the CSA because they would make greater profits by competing only for the customer's 'uncommitted' demand not covered by the CSA.
 - (iv) Exacerbating factors. There are other relevant factors which might amplify the impact of CSAs on competition, though they do not feature directly in our quantitative analysis. The first is the proportion of customer demand covered by the commitment. If the commitment of a newly agreed CSA covers a large proportion of the expected customer demand, customers might be less inclined to consider rivals for allocating new demand to and/or might feel compelled to spend on irrelevant services while the CSA is in place. Customers might behave that way to be sure of meeting the commitment on their CSAs. The second factor is the length of the CSA. The longer the contract length, the longer the period customers might be less inclined to consider rivals for allocating new demand to.
- (d) whether there are any potential benefits that may arise from CSAs even if they may harm competition. For example, some cloud providers have submitted that commitments from CSAs increase the predictability of demand for cloud services and that this helps providers plan their investments. We have assessed the extent to which these submissions are supported by evidence.

8.20 We note that cloud providers offer several other discount programmes besides CSDs. These range from discounts on specific services based on booking cloud

²⁶²⁷ For example, a customer has 1,000 units of demand with a cloud provider which are sticky and 30 which are contestable and that customer has a CSD with that provider giving a 2% discount with a commitment of 1,030 units. In this scenario, all else equal, a rival would have to apply a discount of around 69% to win the 30 contestable units ($(0.02 \times 1,030) / 30 \sim 0.69$). If instead the customer had 100 units of demand with that cloud provider which are sticky, all else equal, the rival would have to apply a discount of around 9% to win the 30 contestable units ($(0.02 \times 130) / 30 \sim 0.09$).

services capacity (eg servers) in advance, standard discount tiers which give increasing discounts to customers who spend/use more on specific services and various types of credits (eg credits to startups and migration credits, ie credits given to customers for migrating workloads from on-premises/private cloud to public cloud).²⁶²⁸

- 8.21 In this chapter, we have not focused on these other discounts in our assessment of CSDs, with the exception of individual service CSDs or cloud credits insofar as they are negotiated alongside CSDs and/or are part of the same CSD contract.²⁶²⁹ Our analysis focused on agreements for which customers commit to a certain level of spend, as this was the focus of Ofcom's reference and was an area of concern raised by stakeholders.

Stakeholder submissions

- 8.22 In this section, we present a summary of the submissions we have received which are relevant to CSDs. We have set out the main high-level arguments in this section and have addressed the detailed points in the specific sections below, which also include a wider range of views and submissions.
- 8.23 We have received a wide range of submissions, with some parties considering that CSDs raise competition concerns and others considering CSDs to be part of normal business practice that should therefore not be restrained.

Cloud providers and ISVs

- 8.24 Some cloud providers and ISVs considered that CSDs raise concerns:
- (a) Former UK Cloud employees submitted that, according to the Jigsaw primary research, even if customers are happy to receive a discount, this does create more barriers to entry for smaller cloud providers.²⁶³⁰
 - (b) IBM submitted that CSAs can reinforce the technical barriers to multi-cloud and present a challenge from customers' standpoint. Nevertheless, IBM also submitted that CSAs warrant a nuanced approach as they also allow customers to benefit from lower prices and stability/predictability of spend.²⁶³¹
 - (c) OVHcloud submitted that CSAs incentivise cloud users to put all their workloads within the services of one cloud provider and steer the ecosystem

²⁶²⁸ We consider cloud credits in chapter 5.

²⁶²⁹ We discuss cloud credits targeted to startups in chapter 5.

²⁶³⁰ [Former UK Cloud employees' response to working papers and updated issues statement](#), page 3.

²⁶³¹ [IBM's response to the Committed spend agreements working paper](#), page 1.

towards cloud providers offering CSAs, as these agreements also apply to the services available on its marketplace.²⁶³²

- (d) An ISV said that CSDs are reasonable business practices, but they raise barriers to entry and expansion for smaller cloud providers by incentivising customers to concentrate their business with one cloud provider.²⁶³³
- (e) An ISV said that CSDs incentivise customers to buy third-party services in the larger cloud providers' marketplaces, [§].²⁶³⁴

8.25 We have also received submissions that considered CSDs to be part of normal business practice or to be pro-competitive:

- (a) AWS said that discounts are pro-competitive and directly benefit customers and that 'a closer look' at its pricing and discounts makes it clear that they do not raise barriers to entry and expansion.²⁶³⁵ AWS submitted that CSDs benefit customers who exercise their bargaining power to pay lower prices.²⁶³⁶
- (b) Microsoft said that its CSDs are a key aspect of price competition for new and existing customers – and lead to lower prices for customers in the UK.²⁶³⁷ Microsoft also said CSAs facilitate investment in UK cloud infrastructure and customer switching.²⁶³⁸
- (c) Google said that CSDs are mutually beneficial for customers and cloud providers and that the prevalence of discounting practices is typically one indicator of a market that is functioning well.²⁶³⁹
- (d) Oracle submitted that volume discounts are standard practice in the industry and are generally unproblematic, but the contract's duration might be an issue.²⁶⁴⁰
- (e) OkTik (an ISV) submitted that it relies upon Microsoft's MACC,²⁶⁴¹ for sales of its product and could be 'heavily impacted' if MACCs were ever to be restricted in the UK.²⁶⁴²

²⁶³² OVHcloud's response to the CMA's PDR dated 28 January 2025, paragraph 14.

²⁶³³ Company A's response to the Issues Statement dated 17 October, paragraph 2.1.

²⁶³⁴ Cloudflare's submission to the CMA [§]. Cloudflare's response to the CMA's PDR dated 28 January 2025, page 3-5.

²⁶³⁵ AWS' response to the Issues Statement dated 17 October 2023, paragraphs 30-32.

²⁶³⁶ AWS' response to the CMA's PDR dated 28 January 2025, paragraph 62.

²⁶³⁷ Microsoft's response to the Issues Statement 17 October 2023, section 6.2. Microsoft response to the Provisional Decision Report, paragraph 133-139.

²⁶³⁸ Microsoft response to the Provisional Decision Report, paragraph, paragraph 138-139.

²⁶³⁹ Google's response to the Issues Statement dated 17 October 2023, paragraph 32.

²⁶⁴⁰ Oracle's response to the Issues Statement dated 17 October 2023, page 5.

²⁶⁴¹ MACC stands for 'Microsoft Azure Consumption Commitment' and is Microsoft's main CSD programme.

²⁶⁴² OkTik's response to the Committed spend agreements working paper.

- (f) Civo submitted that there are foreclosed markets where there is intense concentration on one or two cloud providers and increasing reliance on CSAs to control costs.²⁶⁴³

Industry bodies

- 8.26 We have received submissions from ACT/The App Association, Computer and Communications Industry Association (CCIA) and The Startup Coalition that considered CSDs to be beneficial to customers and investment in cloud services.²⁶⁴⁴
- 8.27 We note that some cloud providers are 'members' or 'supporters' of these industry bodies. In particular, Microsoft is a sponsor of ACT/The App Association,²⁶⁴⁵ AWS and Google are members of the CCIA²⁶⁴⁶ and Google is a supporter of the Startup Coalition.²⁶⁴⁷
- 8.28 Open Cloud Coalition submitted that multi-year CSA contracts with the dominant players can foreclose competition at the procurement stage, making it prohibitively expensive to move workloads elsewhere. It submitted that the assumption that businesses can choose to switch providers 'ignores real-world technical, financial and contractual barriers'.²⁶⁴⁸

Customers

- 8.29 We have also received submissions from a couple of customers to our Committed spend agreement working paper:
- (a) A customer submitted that AWS' current pricing practice, ie linking discounts to spend commitments, means it has no confidence that in a 'steady state' AWS would honour its current pricing structures. The customer submitted that this gives it 'no certainty or ability to properly forecast, let alone shop around'.²⁶⁴⁹
- (b) A banking provider submitted that overall, 'regulatory intervention should be carefully evaluated to avoid jeopardizing the broader business and economic

²⁶⁴³ [Civo's response to the CMA's PDR dated 28 January 2025](#), page 2.

²⁶⁴⁴ We note that some of these submissions referred to discounts more generally and might not necessarily be relevant to CSDs. [ACT's response to the Issues Statement](#) dated 17 October 2023, page 5. [CCIA's response to the Issues Statement](#) dated 17 October 2023, page 3. [Startup Coalition's response to the Issues Statement](#) dated 17 October 2023, page 2. [Startup Coalition's response to the CMA's PDR](#) dated 28 January 2025, page 1.

²⁶⁴⁵ [ACT/The App Association](#), accessed on 19 April 2024.

²⁶⁴⁶ [CCIA Members](#), accessed on 19 April 2024.

²⁶⁴⁷ [Startup Coalition](#), accessed on 19 April 2024.

²⁶⁴⁸ [Open Cloud Coalition's response to the CMA's PDR dated 28 January 2025](#), pages 2-3.

²⁶⁴⁹ [X] submission to the CMA in [X].

advantages associated with cloud adoption and the commercial deals that have been struck to date absent clearly abusive terms and conditions'.²⁶⁵⁰

- 8.30 We note that many customers are unlikely to identify any harm brought about by CSDs. That is because customers might not appreciate the effect of CSDs on the overall competitive structure and might rather focus on the benefits they are receiving from their CSD. Therefore, it is likely that most customers would see CSDs in a positive light, even if CSDs were having an anti-competitive effect.

Other parties

- 8.31 We have also had submissions from two academics, Dr George R Barker and R. Parisi and a consultancy, Rayo, who were generally positive about CSDs.²⁶⁵¹
- 8.32 Simon Hansford, UKCloud former CEO, submitted that CSAs are wholly capable of foreclosing entire market verticals from challengers.²⁶⁵²

Our assessment

- 8.33 We considered first the prevalence of CSDs and after that the impact of CSDs on customers' choices. We then assessed the impact of CSDs on competition, before discussing any potential benefits of CSDs.

Prevalence of CSDs

- 8.34 We considered the prevalence of CSDs. This provides an indication of the size of the segment that could be affected by these agreements and whether this is material enough for any concern to arise.
- 8.35 In relation to the prevalence of CSDs, we have set out below:
- (a) cloud provider views on the prevalence of CSDs; and
 - (b) evidence from data on AWS' and Microsoft's CSDs.

Cloud provider views

- 8.36 AWS said that the market coverage of AWS' CSDs is the relevant metric for a competitive assessment. It submitted that, according to its estimates, the coverage of AWS' CSDs is [X]% of total revenue from cloud infrastructure services in the

²⁶⁵⁰ [Banking Provider 1's response to working papers and updated issues statement](#), pages 1-2.

²⁶⁵¹ Dr George R Barker, [response](#) to the CMA's Competitive landscape working paper, Egress fees working paper and Committed spend agreements working paper, page 74. We note that Dr George Barker is a member of the Oxford Cross Disciplinary Machine Learning Research Cluster (OXML). The OXML is supported by Microsoft (see page 1). R. Parisi, [response to the CMA's Competitive landscape working paper](#), page 14. [Rayo's response to the CMA's PDR](#) dated 28 January 2025, page 1.

²⁶⁵² [Simon Hansford's \(former CEO, UK Cloud Ltd.\) response to the CMA's PDR](#) dated 28 January 2025, page 2.

UK and that this is so low that it is implausible that any rivals could be foreclosed.²⁶⁵³

- 8.37 AWS further submitted that a high market coverage at the level of individual cloud providers is a necessary (but not a sufficient) condition to establish that cloud providers have a ‘unilateral incentive to offer discounts that are aimed at foreclosing rivals’. It submitted that, aggregating the CSD market coverage of different cloud providers ignores the competition between the cloud providers and that the CSD market coverage of all cloud providers together provides no basis to conclude that CSDs can lead to an adverse effect on competition.²⁶⁵⁴ AWS also submitted that treating AWS and Microsoft ‘as a single entity with singular aims and discount structures can only be predicated on an implicit assumption of collective dominance or tacit collusion’, which the CMA had not articulated nor proved.²⁶⁵⁵

Our assessment

- 8.38 AWS’ estimates of market coverage for its CSDs ([X]%) of revenues from IaaS and PaaS in the UK) are in line with our own estimates and indicate that this is a material level of coverage.
- 8.39 Assessing the market coverage of CSDs across AWS and Microsoft, rather than looking at each of their CSDs individually, is the right framework for assessing whether CSDs are a feature of the market that is harming competition:
- (a) Our assessment of market power within our framework for assessing CSDs is appropriate in light of the relevant statutory framework. We are required to decide whether any feature, or combination of features, of each relevant market prevents, restricts or distorts competition in connection with the supply or acquisition of any goods or services in the UK or a part of the UK (an ‘AEC’).²⁶⁵⁶ In determining whether there is an AEC, the CMA will consider a range of factors, which may include the levels of firms’ market power.²⁶⁵⁷
 - (b) Additionally, market power may also be exerted by a single firm, or by several firms acting independently (ie unilateral market power).²⁶⁵⁸ In this case, significant market power is likely to translate into a larger demand for a cloud provider’s services that is sticky, both in absolute terms (by virtue of a

²⁶⁵³ AWS’ submission to CMA [X].

²⁶⁵⁴ AWS’ submission to CMA [X].

²⁶⁵⁵ [AWS’ response to the Committed Spend Agreements working paper](#) dated 31 July 2024, paragraph, 21.

²⁶⁵⁶ See Our task chapter, section ‘our statutory duty’.

²⁶⁵⁷ See [CC3 \(Revised\)](#), Guidelines for market investigations: Their role, procedures, assessment and remedies, paragraphs 9,14 and 101. Assessing whether an undertaking holds a dominant position (be it individually or through collective dominance) is specific to an investigation under the Competition Act 1998.

²⁶⁵⁸ See [CC3 \(Revised\)](#), paragraph 178.

higher market share of that provider) and as a share of total demand for that provider's services.

- (c) Further, smaller cloud providers may still end up being weakened or marginalised by AWS' and Microsoft's CSDs taken together, even if there is competition on some dimensions between AWS and Microsoft. In a similar vein, any competition between cloud providers, including AWS and Microsoft, may actually be softened or be lower than it could be, consequent to their CSDs reducing the incentive of cloud providers to compete for the whole contestable demand.

8.40 In addition, AWS submitted that its estimate does not include UK customers whose non-UK parent entity holds a CSD with AWS and does not include [REDACTED].²⁶⁵⁹ Therefore, we consider that some relevant CSD agreements may be missing from the data set, which might lead to the actual 'market coverage' of AWS' CSDs being underestimated. In particular, we have understood that where a CSD is not held by the UK part of the business, the UK revenue from these customers would be counted in AWS' total UK revenue (the denominator) but would be excluded from the UK revenue that is subject to CSDs (the numerator).

Data analysis

8.41 We have assessed the prevalence of CSDs using two metrics for each of AWS and Microsoft:

- (a) the proportion of customers that have a CSD in place (number of customers with a CSD divided by the total number of customers); and
- (b) the proportion of total annual public cloud spend these customers represent (total public cloud spend by customers with a CSD divided by total public cloud spend by all customers).

8.42 We have calculated these metrics using data provided by AWS and Microsoft.²⁶⁶⁰ We note that this data reflects cloud providers' estimates that are inevitably affected by the complexity of identifying UK customers and the revenue associated with them.

8.43 We have found that the proportion of total annual cloud spend accounted for by customers with a CSD in 2022 was [REDACTED]% for AWS and [REDACTED]% for Microsoft.

8.44 Both of those figures have grown substantially over the years – from [REDACTED]% in 2018 to [REDACTED]% in 2022 for AWS and from [REDACTED]% in 2019 to [REDACTED]% in 2022 for Microsoft.

²⁶⁵⁹ AWS' submission to CMA [REDACTED].

²⁶⁶⁰ AWS' response to the CMA's information requests [REDACTED].

- 8.45 Given AWS and Microsoft had a combined revenue share of supply in IaaS and PaaS in the UK of [60-70]% in 2022, this means their CSDs covered a significant proportion of total annual cloud spend in that year.
- 8.46 These figures represented a very small proportion of the providers' total number of UK customers in 2022 ([X]% for AWS and [X]% for Microsoft). This is mainly driven by the large number of smaller customers (ie those spending less than \$500,000 yearly),²⁶⁶¹ that accounted for around [X]% of total customer base for AWS and for around [X]% for Microsoft.
- 8.47 When looking at customers in higher spend bands, the proportion of customers with CSDs is much higher. For example, in 2022, [X]% of customers spending more than \$500,000 per year with AWS had a CSD in place. The equivalent figure was [X]% for Microsoft.

Impact of CSDs on customers' choices

- 8.48 We considered the extent to which CSDs affect customers' choices in relation to the allocation of workloads on public cloud. If these agreements play a limited role in customers' choices of where to allocate workloads, then they would be less likely to harm competition.
- 8.49 In order to examine the impact of CSDs on customers' choices, we considered:
- (a) cloud provider views on the impact of CSDs on customers' choices; and
 - (b) customer views on the impact of CSDs on their allocation of workloads.

Cloud provider views

- 8.50 In relation to the allocation of new workloads, Microsoft submitted that customers indicated they chose Microsoft for a diverse range of reasons, including quality of service, product range and geographic reach and not just purely price.²⁶⁶²
- 8.51 AWS submitted that customers responding to price differentials is evidence of effective price competition.²⁶⁶³

²⁶⁶¹ As set out in the chapter 3, there are a relatively small number of high-spend customers responsible for a significant proportion of providers' UK revenue and a relatively large number of low-spend customers responsible for a small proportion of their revenue.

²⁶⁶² [Microsoft's response to the Competitive Landscape, Committed Spend Agreements and Egress Fees Working Papers](#), paragraph 69.

²⁶⁶³ AWS' submission to the CMA [X].

- 8.52 Google submitted that, in its experience, CSAs do not stop customers from switching away from AWS and Microsoft to an alternative cloud provider and/or pursuing a multi-cloud strategy for incremental workloads.²⁶⁶⁴

Our assessment

- 8.53 With respect to these submissions, our view is that:
- (a) even though customers indicated they chose Microsoft for a diverse range of reasons (including quality of service, product range and geographic reach and not just purely price), or that they do not proactively mention CSDs as a reason, these submissions are not inconsistent with our view that CSDs have an influence on customer choices in workload allocation. What we are interested in is whether, all other things being equal, CSDs influence such choices. The way we have assessed this is to ask customers directly the extent to which they think CSDs impact their allocation choices, which we did in our customer request for information, as discussed in the Customer evidence section below; and
 - (b) merely observing that CSDs have an impact on customer workload allocation choices does not necessarily mean that CSDs are pro-competitive. We consider the key question to be whether CSDs reduce the ability and incentive of rivals to compete.

Customer evidence

- 8.54 We asked large customers for their views on whether their decisions on switching and the allocation of workloads would be affected by CSDs. We have looked in turn at:
- (a) the impact of CSDs on the allocation of new workloads; and
 - (b) the impact of CSDs on the allocation of existing workloads.
- 8.55 All but one customer that answered these questions had either AWS or Microsoft as their main provider.

New workloads

- 8.56 We asked large UK customers whether CSDs offered by their main cloud providers would influence the allocation of new workloads in the future. The customers had to choose one option on a scale from 'very unlikely' to 'very likely'.

²⁶⁶⁴ [Google's response to the Committed Spend Agreement Working Paper](#), paragraph 6.

- 8.57 In response to this question, most customers we spoke to said that the existence of any CSDs with their main cloud provider was ‘somewhat likely’ or ‘very likely’ to have an impact on their choice of where to allocate new workloads.²⁶⁶⁵
- 8.58 A handful of customers said it is ‘neither likely nor unlikely’ to have an impact on their choice of where to allocate new workloads, whilst a few customers said it is ‘somewhat unlikely’ or ‘very unlikely’ that their choice of new workload is impacted by existence of CSDs with their main cloud provider.²⁶⁶⁶

Existing workloads

- 8.59 We also asked large customers about the effects of CSDs on the allocation of existing workloads. We asked customers whether CSDs offered by their main cloud providers would influence their approach to switching workloads from their main cloud providers to other cloud providers, in the event they wanted to switch any such workloads. The customers had to choose one option on a scale from ‘very unlikely’ to ‘very likely’.
- 8.60 Most customers we spoke to said that it was ‘somewhat likely’ or ‘very likely’ that the existence of any CSDs with their main public cloud provider would influence their decision to move existing workloads away from their main public cloud provider.²⁶⁶⁷
- 8.61 A handful of customers said that the existence of CSDs with their main provider would be ‘neither likely’ nor ‘unlikely’ to influence this decision, whilst a few customers said that it would be ‘somewhat unlikely’ or ‘very unlikely’ to influence this decision.²⁶⁶⁸

Evidence from the Jigsaw report

- 8.62 The Jigsaw report suggested that there are two main ways in which discounts, including CSDs, have at least some impact on cloud customers’ behaviour:
- (a) first, participants said discounts do influence their propensity to stay with their current cloud provider – at least to some extent;²⁶⁶⁹ and
 - (b) second, there is evidence from across the sample that the existence of behaviour-based discounts influences how companies actually use the services offered by their cloud provider. In particular, some research participants described how their companies use certain cloud services, not

²⁶⁶⁵ Responses to the CMA’s information requests [X].

²⁶⁶⁶ Responses to the CMA’s information requests [X].

²⁶⁶⁷ Responses to the CMA’s information requests [X].

²⁶⁶⁸ Responses to the CMA’s information requests [X].

²⁶⁶⁹ [Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw \(2024\)](#), paragraph 6.2, page 78.

because there is a business or an IT need, but for the sole purpose of meeting committed spend targets. Other participants described migrating more workloads than they might otherwise need to, simply to meet committed spend targets.²⁶⁷⁰

Impact of CSDs on competition

8.63 In this section, we have set out our assessment on the extent to which the pricing structure of CSDs has characteristics that may lead to harm to competition.

- (a) First, we have presented a quantitative assessment of whether a rival would have the ability and incentive to compete against the incumbent's CSDs for the customers' contestable demand. We have done so by drawing together the following key factors: the share of demand for AWS' and Microsoft's services that is sticky, the relationship between the discount rate and commitment in AWS' and Microsoft's CSDs and the rival's margins.
- (b) We have then complemented this analysis with a qualitative assessment of two additional factors which, despite not featuring directly in the quantitative analysis, might exacerbate the impact of CSDs on competition.

Cloud provider views

8.64 Microsoft said that over time, the value of discounts in percentage terms that Azure customers received has [REDACTED]. Microsoft also said that its regression analysis of UK CSD contracts shows that [REDACTED].²⁶⁷¹

8.65 AWS submitted the following points, which it said are inconsistent with the claim of customer lock-in or rivals' foreclosure from AWS' CSDs:

- (a) [REDACTED].²⁶⁷² [REDACTED].²⁶⁷³
- (b) The structure and level of AWS' discounts is such that efficient rivals can compete on incremental demand.²⁶⁷⁴
- (c) There is a large body of evidence of entry and growth by many rivals of all sizes.²⁶⁷⁵

²⁶⁷⁰ Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraphs 6.2.8-6.2.9, pages 79-80.

²⁶⁷¹ Microsoft's submissions to the CMA [REDACTED].

²⁶⁷² AWS' submission to the CMA [REDACTED].

²⁶⁷³ AWS' submission to the CMA [REDACTED].

²⁶⁷⁴ AWS' submission to the CMA [REDACTED]. AWS' submissions to the CMA [REDACTED].

²⁶⁷⁵ AWS' submission to the CMA [REDACTED].

(d) Market outcomes (repeated innovation, higher quality of services, decreasing prices, decreasing margins and efficiencies passed on to customers) are consistent with effective competition.²⁶⁷⁶

- 8.66 AWS also submitted that the notion of ‘sticky demand is ‘conceptually inconsistent’ and is not supported by empirical evidence. It submitted that demand which is apparently ‘sticky’ due to switching costs is in truth contestable, since switching costs can be offset, eg, by migration credits, higher discounts, or by offering higher quality services. It submitted that a partition of the demand space should be between demand that is non-contestable and demand that is contestable, as opposed to the CMA’s adopted partition between ‘sticky’ and contestable.²⁶⁷⁷
- 8.67 Google submitted that Microsoft offers customers a discounted bundle of products across cloud and non-cloud services, which acts as a further incentive for customers to continue to place all their demand with Microsoft. Google also said that those discounts notably extend to Microsoft’s most popular non-cloud products, including its ‘must-have’ Office software productivity suite and client PC and work group server operating systems.²⁶⁷⁸
- 8.68 In response to that point, Microsoft said that it does not offer bundled discounts across Azure and its non-Azure services.²⁶⁷⁹ In particular, Microsoft submitted that, while there are many customers that consume Azure services as well as Microsoft 365 services, ‘it is just not possible’ for Microsoft to offer such bundled discounts.²⁶⁸⁰
- 8.69 We asked smaller cloud providers to explain whether CSDs offered by larger cloud providers have had any impact on their business and to provide any supporting evidence. Only one of these providers submitted examples of how CSDs offered by larger cloud providers have had an impact on its business.²⁶⁸¹
- 8.70 We also asked Google and a number of smaller cloud providers whether, when negotiating with a customer that has a standing offer for a CSA from its incumbent cloud provider, they would typically compete for the whole of the customer’s expected spend for new workloads or just the spend for new workloads which will not be covered by the incumbent’s CSA. Overall, the view from these cloud providers was either that the hypothetical situation we described does not typically arise in practice – and therefore, it is unlikely that a decision needs to be made regarding how much demand to compete for – or that when the situation arises, it

²⁶⁷⁶ AWS’ submissions to the CMA [§].

²⁶⁷⁷ AWS’ submission to the CMA [§].

²⁶⁷⁸ [Google’s response to the Committed Spend Discounts Working Paper](#), paragraph 17.

²⁶⁷⁹ [Summary of hearing with Microsoft on Tuesday 16 July 2024](#), paragraph 24.

²⁶⁸⁰ Transcript of hearing with Microsoft [§].

²⁶⁸¹ [§] response to the CMA’s information request [§].

will not generally impact the scope of the workloads for which cloud providers are trying to compete.²⁶⁸²

Our assessment

- 8.71 With respect to AWS' and Microsoft's analyses of any differences between discounts received by customers renewing their CSA and discounts received by customers signing a CSA for the first time – we consider that the model used for these analyses may have led to biased results. In particular, their models might conflate the effect of interest (whether the discount *for a given contract* changes at renewal) with a 'composition effect' (contracts with different levels of discount might be more or less likely to renew). For example, if AWS/Microsoft reduced the discount on renewal, but CSDs with higher levels of discounts were more likely to renew, then the regressions would not pick up the effect of interest.
- 8.72 We have assessed the extent to which effective prices have increased, decreased, or remained stable over time in our assessment of market outcomes in chapter 3. As discussed more in detail there, we have had some reservations about some of the methodology used by both Microsoft and AWS in their pricing analysis, which in some cases overstate price decreases.
- 8.73 Also, even if prices were shown to be stable, this would not be evidence of effective competition if, for example, costs were falling (in the context of expanding demand and economies of scale), as stable prices would be consistent with rising margins.
- 8.74 Finally, even if prices were stable and margins were also stable or declining, any impact of CSDs on competition may not be observed in the short run and there would likely be several other factors that might have affected prices and margins alongside CSDs. The same line of reasoning also applies whether customers receive higher or lower discounts at renewal, which we covered in the paragraph above.
- 8.75 Irrespective of whether there has been some entry and growth by rivals, our analysis indicates that there is concentration in the supply of public cloud infrastructure services in the UK.
- 8.76 We note that the 'sticky' demand, as defined above, is effectively 'non-contestable' for the purpose of the quantitative analysis. That said, we do acknowledge that the assumption made in our quantitative analysis that the existing demand for AWS' and Microsoft's services is all sticky – ie effectively non-contestable – is conservative. That is, it might be the case that at least some of that existing demand is in fact contestable, eg because some services are not must-have or

²⁶⁸² Responses to the CMA's information requests [X].

because the cost of switching some of these services would be modest. We have discussed evidence on sticky demand more in detail in the ‘Share of demand for the incumbent’s services that is sticky’ section below.

- 8.77 In terms of whether Microsoft’s CSDs have non-Azure service such as SPLA and Microsoft 365 among the eligible services, we have not seen evidence of this happening.

Quantitative assessment

- 8.78 We now discuss our quantitative assessment of whether a rival would have the ability and incentive to compete against the incumbent’s CSDs for the customers’ contestable demand. We have drawn together the following factors in our analysis: the share of demand for the incumbent’s services that is sticky, the relationship between the discount rate and commitment in the incumbent’s CSDs and the rival’s margin.
- 8.79 Appendix U CSA – quantitative provides a detailed explanation of our quantitative assessment. This includes a summary of AWS’ economic analysis on which our assessment was based, the data we have used and the methodology we have adopted, the results and the potential limitations of the analysis.
- 8.80 In the appendix, we have also explained why we have not done this quantitative analysis on Microsoft’s CSD data, that is because of: the way Microsoft’s discounts are granted vis-à-vis commitments in its CSDs and the quality of Microsoft’s data which appears to be generally less reliable. Nevertheless, it is reasonable to assume that the main conclusions and results obtained from the analysis of AWS’ data would be broadly similar if extended to Microsoft’s data. [X].
- 8.81 Below, we first explain our approach to the analysis. We then discuss one of the key assumptions behind the analysis – that customers’ existing demand is sticky and finally presented the results of our analysis.

Our approach

- 8.82 As a starting point to our analysis, we have drawn on the quantitative analysis submitted by AWS. AWS’ analysis looked at whether a competitor, which is at least as profitable as AWS itself, can profitably compete for the incremental demand (assumed to be contestable) of customers who have existing demand (assumed to be non-contestable) and an existing CSD with AWS.²⁶⁸³

²⁶⁸³ AWS’ submissions to the CMA [X].

- 8.83 Similar to AWS' approach, we have performed a systematic analysis of the discount that a competitor would have to apply to win the customer's incremental demand in a range of hypothetical scenarios. We have then compared that discount with the AWS' estimated margin after interest and tax on each of its CSDs. Further, we considered hypothetical increases in the size of the incremental demand relative to the existing demand (from 10% to 150%, in 1 percentage point increments), which is in line with AWS' approach.
- 8.84 Our analysis compared the discount that competitor would have to apply to AWS' margin after interest and tax on a specific CSD before the existing discount from that CSD is applied – what AWS calls 'gross margin'. For each CSD, this gross margin is calculated by adding up the existing discount rate from that CSD to AWS' estimate of the expected margin after interest and tax on that specific CSD after the discount is applied – what AWS calls 'net margin'.²⁶⁸⁴ In the following paragraphs, we have adopted AWS' terminology on margins.
- 8.85 We have adjusted AWS' analysis in relation to two dimensions of its modelling:
- (a) **Rivals' efficiency:** AWS focuses on rivals that are at least as profitable as AWS. Instead, we have allowed for varying degrees of profitability. We have done so because we would be concerned about the impact of CSAs on rivals who have the potential to achieve the same margin as AWS and Microsoft if they had similar scale. These rivals could still exert a competitive constraint on the incumbents dynamically. In practice, this has translated into using lower values for the rival's net margin as a benchmark. As a lower bound, we used a value of net margin which is [significantly lower than] [X] AWS' average of [X]% – this translates to a net margin of around [X]%.²⁶⁸⁵ [X].²⁶⁸⁶
 - (b) **The proportion of incremental demand which is covered by commitment:** we have allowed for only a fraction of the customer's incremental demand to be taken up by an increase in the customer's commitment. In other words, we have allowed for the incremental demand to be at least partially free from any extra commitment. This allowed us to look at whether rivals have the incentives to compete for the whole incremental demand or just for the uncommitted incremental demand. [X]. However, one possible effect of CSDs is that rivals experience a substantial drop in profitability if they encroach on the rival's commitment and they may therefore find it more profitable to compete only for the uncommitted demand.

²⁶⁸⁴ AWS' submissions to the CMA [X]; AWS' response to the CMA's information requests [X].

²⁶⁸⁵ [X] AWS' responses to the CMA's information requests [X].

²⁶⁸⁶ See chapter 3, section 'Our analysis of profitability' and appendix K Profitability analysis, sections 'Cost of capital' and 'Our profitability analysis'. See also [CC3 \(Revised\)](#), paragraph 116.

- 8.86 In practice, our analysis works in the following way. Starting from a data set with [X] AWS CSDs, we have built a range of scenarios for each CSD. In these scenarios:
- (a) we allowed for the rival to be less profitable than AWS. We assumed net margins equal to 100%, [X]%, [X]% and [X]% of AWS' average net margin of [X]%;
 - (b) we varied the incremental demand by one percentage point from 10% to 150%; and
 - (c) we varied the proportion of the incremental demand left uncommitted over five values: 20%, 40%, 60%, 80% and 100%.
- 8.87 This meant that each scenario looked like a version of the chart presented in Figure 7.2 above.
- 8.88 This resulted in [X] scenarios.²⁶⁸⁷ For each of these scenarios, we checked if the rival would have the ability and incentive to compete against the incumbent:
- (a) **Ability** – we computed the total profits the rival would make if it won the customer's incremental demand. These were calculated by the difference between the net margin achieved by the rival on that CSA on one hand and the discount that a competitor would have to apply to win the customer's incremental demand, all multiplied by the incremental demand on the other hand. If the total profits were positive or equal to zero, the rival was deemed as having the ability to compete.
 - (b) **Incentive** – we estimated the difference between the total profits the rival would make if winning only the uncommitted incremental demand on one hand and the total profits the rival would make if winning the whole incremental demand on the other hand. If this difference was negative, the rival was deemed as having the incentive to compete for the whole incremental demand.

Share of demand for the incumbent's services that is sticky

- 8.89 As explained above, our quantitative analysis assumed that customers' existing demand for the incumbent's services was all sticky. This was a key assumption underpinning the analysis – the lower the share of customers' existing demand for

²⁶⁸⁷ $[X] \times [X] \times [X] = [X]$; where [X] is the number of AWS' CSAs in the starting data set, [X] is the number of 1 percentage points increments in the range 10% - 150%, [X] is the number of values we consider for the proportion of the incremental demand left uncommitted and [X] is the number of values we consider for the rival's net margin relative to AWS' net margin.

the incumbent's services that is sticky, the higher the ability and incentive or rivals to compete for the incremental demand.

8.90 However, our analysis has taken into account, albeit indirectly and only to a certain extent, different assumptions on the share of existing demand that is sticky. This was through the hypothetical increases in the size of the incremental demand relative to the existing demand (from 10% to 150%). For example, let us assume that only 80% of the share of customers' existing demand for the incumbent's services was sticky, in a hypothetical scenario where incremental demand increased by 50%. This would be similar to assuming that all of the customers' existing demand for the incumbent's services was sticky – ie our default assumption in the analysis, in a hypothetical scenario where incremental demand increased by around 100%. That is because the overall share of demand that is sticky – across the customer's existing and incremental demand – would be the same in both scenarios.²⁶⁸⁸

8.91 In this section, we assessed our default assumption in more detail by looking at evidence on the share of customers' existing demand that is sticky.

Cloud provider views

8.92 AWS submitted that its services are not a 'must-have' (ie unique to AWS) and therefore cannot be leveraged to foreclose competitors.²⁶⁸⁹

8.93 AWS also submitted that the customer evidence presented in paragraphs 2.57-2.64 of the Committed spend agreement working paper does not distinguish between customers who are satisfied with their current cloud provider (and therefore unwilling to switch) and customers who cannot switch part of their demand because it is 'sticky'. It also submitted that, taking these results at face value, they do not suggest that AWS profits from a material share of 'sticky' demand.²⁶⁹⁰ In response to the PDR, AWS submitted that AWS' rivals have the ability and incentive to compete and that it considered the PDR to be 'conservative' and 'incorrect'.²⁶⁹¹

8.94 Microsoft submitted that there will be constant and significant growth in contestable demand. According to Microsoft, this means that, all things being equal, the alleged 'contestable' demand should be significantly larger than any alleged 'sticky' demand at any given time.²⁶⁹²

²⁶⁸⁸ $80/(80+20+50) \sim 100/(100+100) = 50\%$.

²⁶⁸⁹ AWS' submissions to the CMA [36].

²⁶⁹⁰ AWS' submissions to the CMA [36].

²⁶⁹¹ AWS' response to the CMA's PDR dated 28 January 2025, paragraph 64

²⁶⁹² Microsoft's response to the competitive landscape, committed spend agreements and egress fees working papers, paragraph 62a.

- 8.95 Microsoft also submitted that the CMA’s approach to gathering customer evidence on the share of sticky demand is flawed. It submitted that ‘the CMA asked customers what share of demand they are willing to switch, not what they are able to switch’. Microsoft submitted this approach is incorrect, ‘as what matters in the CMA’s concept of sticky demand is the extent to which the customer ‘can exercise effective choice over that demand [which] is limited by factors such as lack of suitable alternatives or barriers to switching’.²⁶⁹³

Our assessment

- 8.96 We focused our evidence gathering to assess the extent to which our assumption of existing demand being all sticky is reasonable. We did so by looking at evidence collected from other workstreams as well as customer evidence specific to CSAs.

Evidence from other workstreams

- 8.97 The analysis on the presence of any other barriers to switching is key to the analysis of the share of existing demand for AWS’ and Microsoft’s services that is sticky:
- (a) We found that there are barriers to switching from factors such as technical differentiation between clouds.²⁶⁹⁴
 - (b) We also found evidence that the additional costs for customers brought about by the presence of egress fees act as barrier to switching.²⁶⁹⁵
 - (c) We found evidence on barriers to switching from Microsoft’s software licensing practices.²⁶⁹⁶
- 8.98 The Jigsaw report noted that a few of the customers interviewed chose to switch between cloud providers.²⁶⁹⁷ It noted that, overall, switching had not only brought cost and operational risk, but it also took IT staff away from the customer’s core work and typically ended up being more challenging and time consuming than anticipated.²⁶⁹⁸ The report also noted that switching cloud providers is seen as the equivalent of moving other kinds of infrastructure, such as ‘moving house’ or moving a business from one country to another. The report further noted that switching cloud providers ‘is not something to undertake lightly or consider at all unless it leads to significant business benefits long term that override the inherent cost and risk of changing. To an extent, dependency on a current cloud

²⁶⁹³ [Microsoft’s response to the competitive landscape, committed spend agreements and egress fees working papers](#), paragraph 70.

²⁶⁹⁴ See chapter 6, Technical barriers.

²⁶⁹⁵ See chapter 6, Egress fees.

²⁶⁹⁶ See chapter 7.

²⁶⁹⁷ [Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw \(2024\)](#), paragraph 1.3.14.

²⁶⁹⁸ [Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw \(2024\)](#), paragraph 1.3.16.

provider(s), or a sense of 'lock in', is a factor across all cloud providers as change brings cost and risk'.²⁶⁹⁹

- 8.99 This evidence on barriers to switching and multi-cloud suggested that the share of existing demand for AWS' and Microsoft's services that is sticky is high. Any additional evidence, such as the customer evidence discussed below, should be read in that context.

Customer evidence

- 8.100 We asked large customers to list any cloud services they were getting from their main cloud provider which they would not be willing to switch to alternative cloud providers, including the proportion of spend that these services accounted for.
- 8.101 This evidence gave an indication of the share of existing demand for AWS' and Microsoft's services that is sticky for this specific group of CSD customers. We considered this evidence in the round and alongside other evidence received and discussed above.
- 8.102 Based on our analysis, there was significant variation in customer responses to this question.²⁷⁰⁰ Among customers having AWS as their main cloud provider, two customers said they would not be willing to switch any or very little of their demand placed with AWS to alternative cloud providers.²⁷⁰¹ A few customers said that they would not be willing to switch a proportion of their AWS demand, ranging from 21-66%.²⁷⁰² Finally, a few customers said that they were willing to switch all of their AWS demand.²⁷⁰³
- 8.103 Among customers having Microsoft as their main cloud provider, a few said that they would not be willing to switch any or virtually any of their demand placed with Microsoft.²⁷⁰⁴ Some customers said that they would be unwilling to switch a proportion of their Microsoft demand, ranging 25-75% and two customers said that they were willing to switch a small proportion of their Microsoft demand, ranging 3-5%.²⁷⁰⁵ Finally, some customers said they were willing to switch all of their Microsoft demand.²⁷⁰⁶
- 8.104 In relation to the cloud provider views set out above, our view is that:

²⁶⁹⁹ [Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw \(2024\)](#), paragraph 1.3.13.

²⁷⁰⁰ Responses to the CMA's information requests [X].

²⁷⁰¹ Responses to the CMA's information requests [X].

²⁷⁰² Responses to the CMA's information requests [X].

²⁷⁰³ Responses to the CMA's information requests [X].

²⁷⁰⁴ Responses to the CMA's information requests [X].

²⁷⁰⁵ Responses to the CMA's information requests [X].

²⁷⁰⁶ Responses to the CMA's information requests [X].

- (a) a customer being unwilling to switch because it is satisfied with its current cloud provider fits our definition of sticky demand, as it likely means that the customer does not think there are suitable alternatives;
- (b) on the use of the term ‘willing’ instead of ‘able’ to switch, use of the term ‘willing’ is reasonable under the given context. Using the term ‘willing’ allowed us to capture instances where a customer would be able to switch from AWS (eg because barriers to switching in its case are low) but not willing to do so because AWS’ services are a must-have for that customer. However, we recognise that some customers said they would not be willing to switch because of their existing CSD (more specifically, because they have a commitment to meet), but they would be willing to otherwise. This means that, were it not for the CSD, that demand would be contestable. This would cause a circularity issue: our framework is such that the larger the share of existing demand for AWS’ services that is sticky, the lower the ability and/or incentive of rivals to compete and therefore the higher the potential harm to competition from CSDs. In other words, any estimation of the share of existing demand for AWS’ services that is sticky should be depurated from any impact of the CSDs. As a consequence, it might be that the estimates for those customers are overestimated.


Results

- 8.105 Generating a chart like the one in Figure 8.2 for a specific AWS’ CSD yielded in the graph set out in Figure 8.3 below: we took a CSD with commitment value equal to a [representative] [%] commitment value for all first-time AWS’ CSDs in our data set.²⁷⁰⁷ We also assumed the incremental demand to be [significantly lower than] [%] the existing demand/commitment, split into a 10% left uncommitted and 40% increase in commitment.²⁷⁰⁸

²⁷⁰⁷ We identified [%] first-time CSAs out of total [%] in AWS’ data set used for our analysis.

²⁷⁰⁸ More details on our assumptions and calculations can be found in appendix U: CSA – quantitative analysis.

Figure 8.3: Total profit to rival for different units of contestable demand it wins for a representative ([X]) AWS' CSD



Source: CMA analysis.

- 8.106 Figure 8.3 above shows that, for this illustrative CSD and scenario, the rival would have both the ability and the incentive to compete for the whole incremental demand.
- 8.107 However, we may have still been concerned if the assessment failed, ie a rival would not have the ability or the incentive to compete, for a subset of CSDs and scenarios. This is why it was important to replicate the analysis for the whole data set of customers. We note that the scenarios used in the assessment were meant to reproduce a range of hypothetical although plausible customer demand increments. In choosing to align with AWS' approach, we were aware that the results presented below are conditional upon the specific range selected (ie, 10%–150%). Although we made some high-level assumptions on what demand increment scenarios may more likely materialise (see further below in this section), the chosen range may not in itself be representative of expected demand growth patterns (eg, scenarios in the higher end of the range may become ever less relevant if demand becomes more sticky).
- 8.108 Our analysis produced the following results.

- (a) When taking rivals as profitable as AWS as a benchmark, we found that:
 - (i) with respect to ability – the assessment failed in 0.1% of cases (both by count and commitment value). This means a rival as profitable as AWS would have the ability to profitably compete for the incremental demand in 99.9% of cases (both by count and by commitment value).
 - (ii) with respect to incentive – the assessment failed in 0.4% of cases (by count) and 0.8% cases (by commitment value). This means a rival as profitable as AWS would have the incentive to compete for the whole incremental demand, as opposed to just the uncommitted demand, in 99.6% of cases (by count) and 99.2% cases (by commitment value).
- (b) When taking rivals [significantly less profitable] [✂] profitable as AWS as a benchmark, we found that:
 - (i) with respect to ability – the assessment failed in 1.6% of cases (by count) and 2.3% cases (by commitment value). This means a rival [significantly less profitable than] [✂] AWS would have the ability to profitably compete, in 98.4% of cases (by count) and 97.7% cases (by commitment value).
 - (ii) with respect to incentive – the assessment failed in 2.3% cases (by count) and 5.3% cases (by commitment value). This means a rival [significantly less profitable than] [✂] AWS would have the incentive to compete for the incremental demand in 97.7% of cases (by count) and 94.7% cases (by commitment value).

8.109 In general, the less profitable the rival is, the more instances in which the ability/incentive assessment failed. We also found that, all else being equal, the higher the share of incremental demand which is uncommitted the more frequently the incentive assessment failed.²⁷⁰⁹ Finally, both assessments tended to fail more frequently for low demand increments. For example, by restricting the set of demand increments to 10% to 50% (instead of the full set 10% to 150%), the instances where the rival would have the ability and incentive to compete for the incremental demand would be lower.

8.110 We considered what demand increment scenarios may be more likely. If the growth in customer spend on cloud services tends to slow down to low levels fairly quickly, then scenarios where the demand increments are lower might be more relevant to assess the impact of CSDs on competition.

²⁷⁰⁹ This is as expected because, all else being equal, the higher the uncommitted share of incremental demand the more profitable it may be for the rival to only focus on this segment of customer's demand.

- 8.111 AWS submitted that our view that CSAs may lead to competitive concerns in the future is ‘speculative’ and ‘based on flawed assumptions’.²⁷¹⁰ Similarly, Microsoft submitted that the view that CSAs could potentially become detrimental to competition in the future is purely theoretical and the CMA, in its PDR, did not assess or calibrate how likely either of scenarios are or what the impact on its analysis would be.²⁷¹¹ We have not sought to assess the likelihood of harm arising from either of the scenarios given our finding that CSDs are currently not a feature giving rise to harm to competition. However, we consider it appropriate to observe that CSDs could, in other scenarios, lead to competition concerns.
- 8.112 To assess this, we looked at two sets of metrics. First, we looked at the average and median increment of AWS’ and Microsoft’s customers’ commitments when they renew their CSDs. We found that, [REDACTED].²⁷¹²
- 8.113 We also looked at average customer growth rates in cloud services spend on AWS from year-to-year since joining AWS for the first time. We found that on average, customers maintained a high growth rate even when customer spend decreased after joining the cloud. This is particularly the case for large customers, which are more relevant than smaller ones in this analysis, because they are more likely to have CSDs.¹
- 8.114 Based on this evidence, we have found that scenarios where demand increments are very low are not likely to materialise in the foreseeable future, though they could materialise at some point if demand for cloud infrastructure services markets becomes more sticky.

Qualitative assessment

- 8.115 The quantitative modelling presented above was driven by three key factors: the share of demand for the incumbent’s services that is sticky, the relationship between the discount rate and the commitment in an incumbent’s CSDs and how profitable the rival is.
- 8.116 We have complemented that assessment by looking at the proportion of customer demand covered by the commitment and the length of CSDs as, despite not featuring directly in our quantitative analysis, these are also relevant factors that could lead CSDs to harm competition.
- 8.117 Overall, we have found that neither the proportion of customer demand which is covered by the commitment, nor the length of AWS and Microsoft’s CSDs, would exacerbate any impact of CSDs on competition. In particular, we have found that customers typically have enough demand left uncommitted from their CSDs which

²⁷¹⁰ AWS’ response to the CMA’s PDR dated 28 January 2025, paragraph 67.

²⁷¹¹ Microsoft response to the CMA’S PDR dated 24 February 2025, paragraph 135-136.

²⁷¹² [REDACTED].

could potentially be allocated to rivals. Although we have found some evidence showing that some customers feel compelled to spend on irrelevant services in order to meet their commitment, we have not considered that evidence was sufficient to lead to a concern. Similarly, we have not found the fact that AWS' and Microsoft's CSDs are typically several years long would, in itself, lead to a concern or alter the results of our quantitative analysis.

Potential benefits of CSDs

8.118 In assessing the potential impact of CSDs on competition, we have also considered whether they may benefit competition and operate to the benefit of customers.

8.119 Therefore, in this section, we considered whether there are any potential benefits arising from efficiencies on competition associated with CSDs (otherwise known as 'rivalry enhancing efficiencies'). This included considering any potential benefits:

- (a) accruing to customers in terms of lower prices; and
- (b) accruing to cloud providers that allow them to operate more efficiently, which may in turn benefit customers.

8.120 In this section we have set out submissions from cloud providers on potential benefits and our assessment of these submissions.

Stakeholder views

Benefits accruing to customers

8.121 Some industry bodies and one customer considered CSDs to be part of normal business practice that should not be restrained. Some of these industry bodies also said that they directly benefit customers, for example, by leading to lower prices.²⁷¹³

8.122 Microsoft and Google also stated that removing or restricting CSDs would reduce benefits accruing to customers.²⁷¹⁴

²⁷¹³ [ACT's response to the Issues Statement](#) dated 17 October 2023, page 5; [CCIA's response to the Issues Statement](#) dated 17 October 2023, page 3; [Startup Coalition's response to the Issues Statement](#), dated 17 October 2023, page 2; and [Banking Provider 1's response to the CMA's working papers](#). As noted above, we note that some cloud providers are 'members' or 'supporters' of these industry bodies.

²⁷¹⁴ Microsoft's submission to the CMA [§]; [Microsoft's response to the Competitive Landscape, Committed Spend Agreements and Egress Fees Working Papers](#), paragraph 59; [Google's response to the Committed Spend Agreements Working Paper](#), paragraph 11.

Benefits accruing to cloud providers and incentives to invest

- 8.123 We have received submissions from cloud providers on how CSDs factor into their investment decisions. Based on this evidence, we have found that cloud providers differ in the extent to which they use CSDs to inform decisions to invest in further infrastructure.
- 8.124 AWS and Microsoft said that CSDs are one of several factors they consider when deciding the timing of investment in infrastructure.²⁷¹⁵
- 8.125 A cloud provider said that its decisions to invest in infrastructure are not directly linked to committed volumes.²⁷¹⁶ It also said that, given that CSAs apply cross-product and cross-region, it cannot reliably use them to predict future demand in specific locations.²⁷¹⁷

Our assessment

- 8.126 Based on the evidence above, we have found that:
- (a) in general, the provision of discounts can be beneficial, at least in the short-term, to those customers who qualify for them.²⁷¹⁸ CSDs, as a form of discounting, allow these customers to pay less than they would if they paid list prices;
 - (b) CSDs can help cloud providers with investment decisions, though it is not clear that CSDs are used extensively for this purpose or that they are the only means of achieving this, based on evidence we have seen. We have not received any empirical evidence of any benefits derived from any surety of demand being passed on to customers through lower prices;
 - (c) even if CSDs were beneficial to customers and investment decisions short term, such benefits would have to be balanced against any potential longer term harm to competition.

Conclusions

- 8.127 AWS' and Microsoft's CSAs cover a significant portion of the cloud infrastructure services markets and influence customers' choices in relation to workload allocation.

²⁷¹⁵ AWS' and Microsoft's responses to the CMA's information requests [X].

²⁷¹⁶ [X] response to the CMA's information request [X].

²⁷¹⁷ [X] response to the CMA's information request [X].

²⁷¹⁸ Discounts used as part of a foreclosure strategy can be harmful for customers and competition. See for example chapter 7.

- 8.128 However, in relation to the impact of CSAs on competition, our quantitative analysis shows that AWS' and Microsoft's CSAs, as currently applied, do not adversely affect rivals' ability or incentive to compete against AWS and Microsoft.
- 8.129 Our qualitative analysis supports these results and further showed that neither the proportion of customer demand covered by the commitment, nor the length of CSAs are likely to be consistent with CSAs leading to harm to competition.
- 8.130 In view of the above, the way in which AWS and Microsoft currently apply CSAs does not constitute a feature leading to harm to competition in the provision of public cloud infrastructure services in the UK.
- 8.131 Although we have not identified competition concerns, this does not mean that harm to competition might not materialise in the future. In particular, we consider that the link between sticky and contestable demand in AWS' and Microsoft's CSAs could harm competition if, for example:
- (a) the share of sticky demand increases significantly; and/or
 - (b) AWS and/or Microsoft change the way their CSA discounts are applied by increasing the incentive of customers to concentrate their spend with them. For example, this could happen by significantly reducing the number and/or increasing the width of the discount bands, therefore increasing the coverage of the commitment and/or the effect of discounts on customers' incentives.

9. Our decision on competition in cloud services markets

- 9.1 We have found, pursuant to section 134(1) EA02, that there are relevant features of the relevant cloud services markets in the UK, which alone or in combination, prevent, restrict or distort competition in the supply of cloud services in the UK such that there are adverse effects on competition (AECs) within the meaning of section 134(2) EA02.²⁷¹⁹
- 9.2 We have found a product market for IaaS based on standard compute infrastructure which is separate from IaaS based on accelerated compute²⁷²⁰ and a product market for PaaS in aggregate: these are the relevant cloud services markets.²⁷²¹
- 9.3 Our AEC findings would not change based on the precise boundaries of the relevant product markets. This is because we have considered competitive constraints coming from inside and outside of the relevant product markets, to the extent relevant, when conducting the competitive assessment. If the market was wider, the additional constraints that would then fall within the market would be relatively weak compared to those we have assessed as in-market constraints. If the market was narrower, the constraints that would then fall outside the market would be relatively strong out-of-market constraints (albeit weaker than the in-market constraints). As a result, any reasonable difference in how we define the relevant product markets would not affect our decision.
- 9.4 We found that the geographic scope of the IaaS and PaaS markets is Europe-wide (that is, UK and EEA).²⁷²²
- 9.5 We set out below our view of what the cloud services markets would look like, were they to be functioning well, an overview of our competitive assessment and our decision on the AECs that we have identified in this investigation.

Our view of well-functioning cloud services markets

- 9.6 In identifying relevant features or a combination of features of the market that prevent, restrict or distort competition thereby giving rise to an AEC(s), we have to

²⁷¹⁹ The statutory task of this market investigation is to decide whether any relevant feature, or combination of features, prevents, restricts or distorts competition, in connection with the supply or acquisition of goods or services in the UK or a part of the UK pursuant to the EA02 and not whether there has been an infringement of the Competition Act 1998 (CA98). The process is 'investigative and inquisitorial, not accusatorial', so involvement in a market investigation or being subject to remedies does not imply that participants have infringed the law or are suspected of wrongdoing. See [CC3 \(Revised\), Guidelines for market investigations: Their role, procedures, assessment and remedies](#) (CC3(Revised)) paragraph 21.

²⁷²⁰ For the purposes of this investigation, when not otherwise specified, we use the term IaaS to refer to IaaS based on standard compute.

²⁷²¹ Our product market definition analysis is set out in chapter 3.

²⁷²² Our geographic market definition analysis is set out in chapter 3.

find a benchmark against which to determine how the market may be judged to be performing.²⁷²³

- 9.7 In the absence of a statutory benchmark, we use the benchmark of ‘a well-functioning market’²⁷²⁴ to mean one that displays the beneficial aspects of competition, rather than an idealised, perfectly competitive market. The benchmark will generally be how we envision the market without the features that are identified as harming competition.
- 9.8 There may sometimes be reasons to depart from that general concept, for example, if features are intrinsic to the market but nevertheless have anticompetitive effects (as in the case of a natural monopoly).²⁷²⁵
- 9.9 In summary, we use the term ‘well-functioning market’ to mean a market without the features, or, where they are intrinsic, the effect of these features, causing the AEC(s).²⁷²⁶
- 9.10 In our view, well-functioning cloud services markets are those in which customers do not become locked into their cloud providers but rather are able to switch cloud provider, if and when they wish to and use multi-cloud architectures to the degree that they find beneficial, with materially fewer impediments (including costs) than they currently face. We would also expect there to be more entry and expansion by cloud providers, and cloud providers being more able to compete effectively for a wider range of workloads and across the range of cloud services, including those services that use Microsoft software.
- 9.11 If customers were able to switch and multi-cloud more easily and if they had more choice of cloud provider, they could choose the best products and services for their needs, whether by switching cloud provider, by combining services best suited for their needs across cloud providers through multi-cloud architectures, or both. Greater entry and expansion by cloud providers²⁷²⁷ would also increase the incentive of incumbents to compete effectively to retain customers throughout all stages of a customer’s or a workload’s ‘life-cycle’ in the cloud.²⁷²⁸
- 9.12 Therefore, in well-functioning cloud services markets we would expect there to be more switching and more multi-cloud than we observe at present. This may ultimately be expected to lower prices and improve the quality of cloud services.

²⁷²³ CC3 (Revised), paragraph 320.

²⁷²⁴ CC3 (Revised), paragraph 30. See also paragraphs 10-12, for a description of the beneficial aspects of competition that we would typically expect to see in a market.

²⁷²⁵ CC3 (Revised), paragraphs 154 and 320.

²⁷²⁶ CC3 (Revised), paragraph 320.

²⁷²⁷ We recognise that some of the existing barriers to entry and expansion that have been identified (along with some of the identified technical barriers) may be intrinsic to some extent. We consider that a well-functioning market is one that is not experiencing the adverse effects of any intrinsic aspects of these barriers.

²⁷²⁸ CC3 (Revised), paragraphs 205 and 207.

Competitive assessment

- 9.13 In the following paragraphs we set out an overview of our competitive assessment.²⁷²⁹ We then set out our decision on the AECs we have identified and the customer detriment arising from them.

Market concentration

- 9.14 We found that cloud services markets are highly concentrated in the case of IaaS and concentrated in the case of PaaS, and forward-looking metrics suggest that these levels of concentration are likely to endure. All metrics on shares of supply in IaaS, PaaS and combined IaaS and PaaS, show that Microsoft and AWS are the largest cloud providers with Google, the third largest cloud provider, with a much smaller share. Over the period from 2020 to 2024, Microsoft's share has increased while AWS' share has decreased slightly.
- 9.15 The IaaS market is particularly highly concentrated. Microsoft and AWS' combined market share remains high and stable. Despite a small increase in its share from 2020-2024, Google remains some way behind Microsoft and AWS and, on the basis of forward looking metrics, is likely to remain so for the foreseeable future. Other smaller cloud providers have lower and/or declining shares, suggesting that they are likely to remain significantly smaller than Microsoft and AWS.
- 9.16 The PaaS market is also concentrated but much less so than IaaS and with a longer tail of small providers, many of which are ISVs. However, we note that most ISVs do not also offer IaaS as well as PaaS, so any PaaS sale they make implies a sale for an IaaS provider, as IaaS is an input to their services. This weakens the constraint that they exert on vertically integrated firms supplying both IaaS and PaaS, such as Microsoft, AWS and Google.
- 9.17 Across both the IaaS and PaaS markets, there are some smaller providers, including Oracle and IBM, whose share of supply is smaller than Google's and which do not supply as wide a range of cloud services as Microsoft, AWS and Google.
- 9.18 On the basis of this assessment, we found that levels of concentration in the cloud services markets are high, particularly in the case of IaaS. This has contributed to weakened competition in these markets, as explained further below.

Market outcomes

- 9.19 We have considered the outcomes of the competitive process in the cloud services markets in terms of profitability, pricing, quality and innovation and the

²⁷²⁹ Our competitive assessment is set out in chapters 3 to 8 of this report.

prevalence of switching and multi-cloud in order to determine how well the markets are currently functioning. While we did not consider any indicator to be solely determinative of a particular outcome, we placed more weight on observable and measurable indicators (in this case, profitability) given the inconclusiveness and/or limitations of other indicators we assessed. When considered in the round, and in light of our findings on market concentration described above, we found that they point to competition in the cloud services markets not working as effectively as we would expect.

- 9.20 In particular, we found that Microsoft's and AWS' cloud businesses have generated sustained profits (as measured by their ROCE) in excess of their cost of capital in cloud services for several years. Together with their large market shares, this indicates that Microsoft and AWS, both individually and collectively, hold strong positions in the IaaS and PaaS markets and that other cloud providers present only a weak constraint on their activities.
- 9.21 We found that pricing trends presented a mixed picture. Some services have fallen in price, while others have seen their prices increase, which is consistent with markets with the potential to be more competitive.
- 9.22 We found that there has been innovation in the market, including by the three main cloud providers. Cloud providers innovate and invest to win business and this is recognised by customers. However, limitations in the measures for quality and innovation made it difficult to evaluate the significance of these metrics and levels. Given this, and the clear and established method for assessing profitability, we placed more weight on our assessment of the cloud providers' profitability and market shares than the analysis of quality and innovation.
- 9.23 When examining the prevalence of switching, we found that full switching between cloud providers is extremely rare in these markets. Although this could reflect customers' satisfaction with their current cloud provider, we considered that very low levels of switching in the context of high levels of profitability among the larger cloud providers more likely emanates from the existence of high barriers to switching. This is reflected in customer evidence and suggests that switching costs outweigh the benefits of changing cloud provider for many customers. We would expect to see greater levels of switching in a well-functioning market.
- 9.24 The use of multi-cloud architectures is more common than switching throughout the cloud markets, particularly for larger customers. However, we found that barriers (as explained further below) still exist especially where customers are seeking to integrate workloads across different clouds. These barriers may be harder to overcome for smaller customers than for larger ones. We have found that even in circumstances where customers have opted for a multi-cloud strategy, they generally concentrate a large proportion of their workloads on a 'primary' cloud provider.

- 9.25 In summary, we found that each of Microsoft and AWS holds a strong market position in cloud services. Google remains significantly smaller than them and other cloud providers are considerably smaller. The competitive constraint posed by Google and other cloud providers collectively is limited.
- 9.26 We would expect that a more competitive market would have sustained better market outcomes, including more consistently competitive prices, potentially further improvements in quality and innovation and greater prevalence of switching and multi-cloud. We considered that certain features in the cloud services markets are having an impact on competition as indicated by these market outcomes and we set out more detail about those below.

Impact of AI on cloud services

- 9.27 Some cloud providers have told us that growth in AI is having a significant impact on competition in cloud and that Google is a stronger competitor in AI-related services than in cloud services more broadly.
- 9.28 We found that Microsoft, AWS and Google are all strong vertically-integrated cloud providers of AI-related cloud services, including through their partnerships with FM developers, and that other cloud providers are significantly weaker in this respect.
- 9.29 Whilst AI-related cloud services are likely to be an increasingly important part of the offer from cloud providers, current evidence from customers and providers suggests that AI is yet to materially alter the competitive dynamics in cloud services.
- 9.30 Cloud providers' supply of IaaS based on accelerated compute is not currently significantly impacting competition in cloud infrastructure services. This is because there is a lack of demand and supply side substitutability between standard and accelerated compute services and significant barriers to entry and expansion in IaaS. Any impact is further limited by the fact that IaaS based on accelerated compute is being used by a small number of customers who are primarily AI model developers.
- 9.31 Revenue from the provision of access to FMs is rising for Microsoft, AWS and Google, but it is still a relatively small part of their businesses. This means that the scope for demand for access to FMs to translate into greater demand for other cloud services is currently limited, although we expect demand to grow over time.
- 9.32 AI capabilities currently play a limited role in driving customers' choice of cloud provider. However, these capabilities are likely to become more important to customers over time and so AI could impact competition in cloud services to a greater degree in the future.

- 9.33 The precise nature and timing of any impact of AI, including on the relative positions of Microsoft, AWS and Google in cloud services, is currently unclear but is likely to depend on how customers' demand for AI capabilities across IaaS, PaaS and SaaS evolves and the extent to which increased customer demand for AI capabilities leads to increased demand for other cloud services from the same cloud provider.

Barriers to entry and expansion

- 9.34 Market entry and expansion in the supply of IaaS based on standard compute requires significant capital investment in fixed assets, which for many asset types would be largely irrecoverable upon exit. This combines with economies of scale, whereby the larger cloud providers have comparatively lower associated ongoing costs. Unless a new entrant (or company seeking to expand) is willing to make investments of a similar magnitude to those of the largest suppliers, it is likely to face higher ongoing costs to provide an equivalent level of service and so may struggle to compete effectively. We considered it likely that this disincentivises IaaS market entry and expansion.
- 9.35 Furthermore, given the scale of investment and expansion that the three main cloud providers have made in IaaS, any new entrant (or company seeking to expand) would need to invest substantially more over the next few years than the large existing suppliers in order to gain competitively significant market share. The levels of investment that Microsoft, AWS and Google are expecting to make in IaaS and in AI infrastructure in the coming years may raise these barriers to entry and expansion even higher.
- 9.36 We note that there have been recent announcements of significant investments in energy infrastructure and data centres by a range of investment firms and that these investments are focused on increasing the supply of data centres to support AI.
- 9.37 Investment in infrastructure is a requirement for entry and expansion into IaaS, but not PaaS. However, competition in PaaS using third party IaaS often involves sourcing IaaS from firms that are also competing in the supply of PaaS. This limits the scope for PaaS providers to compete effectively.
- 9.38 While we recognise that investments by cloud providers may have pro-competitive effects and that they can benefit customers, this does not preclude them also having the effect of deterring market entry.
- 9.39 The wide product portfolios of the three main cloud providers also contribute to the barriers to entry and expansion in both IaaS and PaaS markets because range of services is an important consideration for customers selecting a cloud supplier. ISVs also value access to the wider customer bases of the three main cloud

providers and prioritise developing services compatible with cloud services offered by these cloud providers. In addition, customers place some importance on a cloud provider's reputation and this contributes to overall barriers to entry and expansion in cloud services markets.

- 9.40 In summary, we have found substantial barriers to entry and expansion in the cloud markets, in particular for IaaS, and the indications are that these are unlikely to reduce significantly over time, if at all. The deterrent effect of these barriers impacts effective competition in the cloud services markets.

Barriers to switching and multi-cloud

- 9.41 In the context of extremely low switching rates and limited use of multi-cloud, as set out above, we found that customers face both technical and commercial barriers (the latter in the form of egress fees) that affect the ease with which they are able to switch cloud provider or use multiple clouds.
- 9.42 Customers' ability to exercise choice of cloud provider is important to exert pressure on providers and drive effective competition. We have found that the existence of technical and commercial barriers means that customers' ability to switch and use multiple clouds is more limited than we would expect in a well-functioning market, and that cloud providers' incentives to make switching and multi-cloud easier for customers are mixed.
- 9.43 Cloud services have developed in different ways, with different functionalities and features evolving over time. As a result, technical barriers have arisen from the differentiation of features and interfaces across cloud services. This affects the interoperability of cloud services and means that customers cannot easily compare or substitute products without technical effort and other costs thereby impacting their ability to multi-cloud and/or switch clouds. Further disincentives for customers to multi-cloud include higher latency (the time it takes to transfer data between public clouds), a lack of transferable skills across clouds and insufficient transparency about how to mitigate or overcome technical barriers.
- 9.44 Some customers may be able to multi-cloud where the need for integration between particular cloud workloads is minimal or by adopting a workaround. Equally, some customers may be able to switch or multi-cloud if they are willing to invest sufficient time and money to overcome the technical barriers they face. However, for many customers, the costs associated with overcoming these barriers outweigh the benefits.
- 9.45 We found that customers have some ways to mitigate the impact of technical barriers to a greater or lesser extent in specific use cases, such as by using containers or cloud-agnostic services and tools. Some customers recognised some positive aspects of existing technical mitigations, but also thought that cloud

providers could do more to improve them or offer additional mitigations. We found that, in general, the incentives for cloud providers to develop or further improve mitigations for certain technical barriers are mixed, and in some cases are not apparent. Existing mitigations at best lessen the effects of specific technical barriers on customers' ability to switch and multi-cloud rather than fully addressing the impact of technical barriers as a whole.

- 9.46 In terms of commercial barriers, we considered a number of commercial practices in cloud services and found that egress fees paid by customers for transferring data from one cloud to another are a barrier to switching and/or use of multi-cloud. These costs can be a factor deterring customers from exercising effective choice and from responding to better offers from other cloud providers, leading to weaker competition between providers.
- 9.47 We found that egress fees are not typically the most important factor in customer decisions relating to switching and multi-cloud. Nonetheless, we found that some customers identify them as a challenge and in some instances a disincentive when considering switching and multi-cloud. We found that the significance of egress fees as a barrier varies among customers, with the highest costs in egress fees relative to their overall cloud spend more likely to be paid by smaller customers and customers egressing larger proportions of stored data.
- 9.48 Although Microsoft, AWS and Google have introduced free switching programmes, these have limitations and restrictions including that they do not cover multi-cloud use, are limited to a 60-day switching period and afford cloud providers a substantial degree of discretion over the application of the programmes. These programmes do not comprehensively address our concerns regarding switching and multi-cloud egress fees.
- 9.49 We have not seen sufficient evidence that investments and innovations cited by cloud providers as being funded by egress fees are dependent on or are solely driven by these fees, nor is it clear that they result in passed-on cost savings to customers or deter customers from inefficient egress usage.
- 9.50 In summary, where the expected costs of switching and/or multi-cloud exceed the expected benefits, customers are likely to choose not to switch or to multi-cloud, including when alternative cloud providers would otherwise have a better offer (in terms of price, quality, range of features/capabilities, or other considerations).
- 9.51 Technical and commercial barriers, particularly where customers feel disincentivised to switch, enable cloud providers to exercise a degree of power over their existing customer base, which reduces their incentives to compete for rivals' customers on price, service quality, and/or innovation. Technical and commercial barriers also deter entry and expansion by making it harder for cloud providers to win customers from their rivals.

- 9.52 As such, these technical and commercial barriers exacerbate the barriers to entry and expansion arising in these markets, weaken competition between cloud providers and contribute to Microsoft and AWS being able to maintain high and stable market shares.

Microsoft's licensing practices

- 9.53 Microsoft provides a range of software that customers use on cloud services, either on Microsoft's own cloud Azure or through licensing to other cloud providers. We have examined whether Microsoft's licensing practices are influencing customers' choice of cloud provider, and whether this harms competition by preventing AWS and Google from competing as effectively for customers that use Microsoft software on cloud.²⁷³⁰
- 9.54 We found that Microsoft's software licensing practices have the effect of reducing competition in cloud services markets by materially disadvantaging AWS and Google in the supply of cloud services, and that they consequently compete less effectively.
- 9.55 Microsoft has significant market power in relation to each of Windows Server, SQL Server, Windows 10/11, Visual Studio and its productivity suites. This is because customers are generally unable or unwilling to switch away from these products, there are limited alternatives and Microsoft has moderate to high market shares. AWS and Google have had limited success in moving customers away from the Microsoft software products through modernisation programmes and these programmes have not significantly affected Microsoft's market power in relation to these software products.
- 9.56 Windows Server and SQL Server are important inputs to cloud services for a range of customers. They account for a material proportion of AWS' and Google's costs in providing different bundles of cloud services to a number of customers.
- 9.57 The input price paid to Microsoft by AWS and Google for Windows Server and SQL Server is higher than Microsoft's customer-facing price for some cloud customers. This price disparity shows the potential for Microsoft's licensing practices to disadvantage AWS and Google in competing for customers.
- 9.58 A comparison of AWS and Google's profit margins is not determinative of partial foreclosure, due to both conceptual difficulties in interpreting margin levels and practical difficulties in estimating them in this particular case, notably the uncertainty identified above around how customers purchase sets of cloud services. We have therefore considered ranges of estimates of AWS' and

²⁷³⁰ The relevant Microsoft software products are: Microsoft's Windows Server (which includes Active Directory functionality), SQL Server, Windows 10/11, Visual Studio and the productivity suites.

Google's margins and we consider these ranges to be consistent with partial foreclosure.

- 9.59 VDI allows customers to use Windows 10/11, Visual Studio and the Microsoft productivity suites on cloud. VDI represents a small proportion of customers' spend on public cloud services. However, Windows 10/11 and the Microsoft productivity suites are particularly important to VDI and Visual Studio is important to VDI for some customers. As such, in combination with Windows Server and SQL Server workloads, VDI workloads add to the importance of Microsoft software as an input to cloud services.
- 9.60 Most customers are not aware of quality differences between using Microsoft software on Azure compared to AWS or Google. However, some customers identified quality differences and, for these customers, Windows Server, SQL Server and/or Windows 10/11 may be particularly important in terms of affecting the overall quality or attractiveness of AWS and Google as alternatives to Azure.
- 9.61 Microsoft does not make Windows 10/11 and Microsoft 365 available to AWS and Google through their respective SPLAs, and customers with existing licences cannot bring these to AWS and Google (except for specific Microsoft 365 licences to Amazon Workspaces). These restrictions mean that AWS' and Google's competitive offerings are directly affected by Microsoft's practices in relation to these software products.
- 9.62 Microsoft's licensing practices with respect to the productivity suites, Windows 10/11 and Visual Studio are likely to affect competition for VDI workloads. While VDI represents a relatively small proportion of usage of all cloud services, we found that Microsoft's licensing practices relating to these software products contribute to an adverse effect on the competitiveness of AWS' and Google's offerings.
- 9.63 Windows Server and SQL Server are used disproportionately on Azure compared to on AWS and Google. We considered that, given the very large difference, this is at least partly because some customers' choice of cloud is influenced by Microsoft's licensing practices.
- 9.64 Considering all the evidence in the round, we found that Microsoft's licensing practices are adversely impacting the competitiveness of AWS and Google in the supply of cloud services, particularly in competing for customers that purchase cloud services which use the relevant Microsoft software as an input. As a result, Microsoft faces weaker competitive constraints from AWS and Google, its most significant competitors, which is reducing competition in cloud services markets.

Our decision on adverse effects on competition

- 9.65 Taking all of the above considerations into account, there are relevant features of the markets for the supply of cloud services which, individually and in combination, prevent, restrict or distort competition in connection with the supply of cloud services in the UK such that there are AECs within the meaning of section 134(2) EA02.
- 9.66 Customers of cloud services markets are varied in terms of their spending and use of cloud (for example, in terms of number of services used and overall reliance). This means that the individual features of the cloud services markets we have identified may be experienced by particular customers to a greater or lesser degree. Even if a feature affects only a proportion of cloud services customers, or affects customers only at certain periods in time such as when they are considering whether to switch provider, this does not preclude a finding that the feature(s) prevent, restrict or distort competition in the cloud services markets. Where one or more features are affecting competition in combination with each other, those features may not contribute equally to the harm identified.²⁷³¹
- 9.67 Market concentration, which is particularly high in the case of IaaS, and barriers to entry and expansion are features that have enabled each of the two main cloud providers, Microsoft and AWS, to hold significant unilateral market power in the cloud services markets and to earn returns above the cost of their capital over a sustained period. Forward-looking metrics indicate that these levels of concentration are likely to endure. These features give rise to an AEC in cloud services in the UK because it is harder for alternative cloud suppliers to enter and grow in these markets and customers face a limited choice of suppliers and products. This harm is exacerbated by the features arising from technical and commercial barriers to switching and multi-cloud.
- 9.68 These technical and commercial barriers are features that give rise to an AEC in cloud services in the UK. These barriers lock customers into their initial choice of cloud provider, which may not reflect their evolving needs and limits their ability to exercise choice of cloud provider. These barriers can restrict customers from responding to attractive offers or accessing innovative new services from another cloud provider, leading to weaker competition between cloud providers.
- 9.69 Microsoft's licensing practices have the effect of reducing competition in cloud services markets by adversely impacting the competitiveness of AWS and Google in the supply of cloud services, particularly in competing for customers that purchase cloud services which use the relevant Microsoft software as an input. This gives rise to an AEC in cloud services in the UK. These licensing practices are a feature that, in combination with the other features we have identified,

²⁷³¹ CC3 (Revised), paragraphs 155 and 162.

including Microsoft's large and increasing market share in these markets, further restricts the already limited choice and attractiveness of alternative products and suppliers.

Customer detriment

- 9.70 The AECs that we have found may be expected to result in substantial customer detriment in cloud services in the UK.²⁷³²
- 9.71 Our guidance sets out that prices and costs are among the more observable and measurable outcomes of how well competition is working in a market, and an analysis of these may be useful in quantifying the extent and nature of competition and in measuring customer detriment. In addition, less quantifiable factors, such as choice, quality and innovation are also important to customers.²⁷³³
- 9.72 The AECs that we found have had a material impact on customers' ability to switch, multi-cloud and exercise choice over their cloud services provider, which may ultimately be expected to impact the price and quality (including access to innovative new services) of cloud services.
- 9.73 Detriment may manifest itself in terms of UK customers facing more limited choice of cloud services and paying higher prices or benefiting from lower levels of quality or innovation for these services than they would if the markets were more competitive.
- 9.74 We expect that prices are higher than they would be in the absence of the harms to competition that we have identified. In the context of the £10.5 billion spent annually on cloud services by UK businesses, the public sector and other customers, even a small impact on competition in the markets could lead to significant customer detriment. By way of illustration, if prices were on average 5% above those in a well-functioning market, this would in aggregate lead to UK customers paying around £500 million more per year to purchase the same services than they would in a more competitive market.²⁷³⁴
- 9.75 We have also found that the quality of cloud services offerings would be expected to improve in certain respects absent the harms to competition we have identified (for example, in respect of services supporting multi-cloud). A reduction in the quality of cloud services also represents a detriment to UK customers.

²⁷³² A detrimental effect on customers is defined as one taking the form of (a) higher prices, lower quality or less choice of goods or services in any market in the UK (whether or not the market or markets to which the feature or features concerned relate); or (b) less innovation in relation to such goods or services. EA02, section 134(5).

²⁷³³ CC3 (Revised), paragraph 104.

²⁷³⁴ £10.5billion was the value of IaaS and PaaS UK revenue in 2024. Calculations are as follows: 10.5billion - (10.5billion/1.05) = 10.5billion - 10billion = 500million. See our full estimates for UK IaaS and PaaS revenues from 2020 to 2024, in chapter 2.

10. Our remedy decision

- We considered whether the recently introduced DMCC Act powers are better suited to address the concerns that we have identified than the powers available to us under EA02. DMCC Act powers were specifically designed to address concerns arising in digital markets and would allow the CMA to take a targeted and iterative approach to its remedies. These include new powers designed to enhance the effectiveness of remedies, and better provisions for ongoing monitoring and oversight. These powers also enable a more flexible approach that allows the CMA to take into account new developments and their potential impact on any interventions addressing issues arising in markets in the UK on an ongoing basis.
- We have assessed potential remedies that we could implement through our remedy-making powers under EA02, consisting of order-making or accepting undertakings from parties. We considered potential remedies in relation to technical and commercial barriers and Microsoft's licensing practices, and while, in principle, these may be capable of addressing the features, we have identified a number of material risks to their effectiveness if taken forward under the remedy-making powers in EA02.
- Taking this into account, we have concluded that the following remedies would represent as comprehensive a solution as is reasonable and practicable to the AECs and resulting consumer detriment that we have found:
- Remedy 1: a recommendation to the CMA Board to prioritise commencing an SMS investigation of AWS' digital activities in respect of cloud services using its DMCC Act powers taking into account all relevant factors at the time, and, if an SMS designation is made, to consider imposing appropriate interventions; and
- Remedy 2: a recommendation to the CMA Board to prioritise commencing an SMS investigation of Microsoft's digital activities in respect of cloud services using its DMCC Act powers taking into account all relevant factors at the time, and if an SMS designation is made, to consider imposing appropriate interventions.
- Focusing any interventions on Microsoft and AWS as the two largest providers of cloud services would directly benefit most UK customers as well as altering the wider competitive conditions and commercial pressures on other cloud providers which, in combination, would effectively address the AECs we have identified.
- The CMA has indicated that no decision will be taken by the CMA Board on future SMS designation investigations in 2025 but that it will keep under review possible options and it anticipates that these will be considered in early 2026.
- While we recognise that there is some uncertainty around the implementation of our recommendations, we believe that they are the only effective remedy capable of comprehensively addressing the AECs we have found in this investigation. We

recommend that, in the intervening period pending the CMA Board's decision on prioritisation, developments in the market, including any actions taken by cloud providers, are kept under review and that the CMA Board factors these into its decision.

Our decision on remedies

- 10.1 This chapter sets out the actions we have decided to take for the purposes of remedying, mitigating or preventing the AECs we have identified, and any detrimental effects on customers that result, or may be expected to result, from the AECs.
- 10.2 We have identified competition concerns in the cloud services markets arising from high concentration, barriers to entry and expansion, technical and commercial barriers to customers switching and use of multi-cloud and Microsoft's licensing practices.
- 10.3 We consider that the CMA's new digital markets powers are an effective way to remedy the harms to competition we have found. The remedies we have decided on are therefore as follows:
- (a) Remedy 1: a recommendation to the CMA Board to prioritise commencing an SMS investigation of AWS' digital activities in respect of cloud services using its DMCC Act powers taking into account all relevant factors at the time, and, if an SMS designation is made, to consider imposing appropriate interventions; and
 - (b) Remedy 2: a recommendation to the CMA Board to prioritise commencing an SMS investigation of Microsoft's digital activities in respect of cloud services using its DMCC Act powers taking into account all relevant factors at the time, and if an SMS designation is made, to consider imposing appropriate interventions.
- 10.4 In arriving at our decision, we have taken into account the considerations relevant to making a recommendation and to determining its scope set out in EA02²⁷³⁵ and in relevant guidance applicable to market investigations.²⁷³⁶ These key considerations are:
- (a) Effectiveness: the extent to which recommendations to the CMA Board would be a comprehensive solution to the AECs we have found, including the likelihood of the recommendations being acted upon and over what time period; and

²⁷³⁵ EA02, sections 134(6) and (7).

²⁷³⁶ These are the considerations set out in [CC3 \(revised\)](#) at paragraphs 379, 380, 390, 391 and Annex B, paragraphs 94 to 102.

- (b) Proportionality: whether recommendations to the CMA Board are a proportionate remedy.

- 10.5 We have found that our recommendations to the CMA Board would effectively address the competition concerns we have identified in relation to the significant unilateral market power held by AWS and Microsoft arising from high levels of market concentration, barriers to entry and expansion, barriers to switching and multi-cloud and Microsoft's licensing practices.
- 10.6 We have also found that remedial action targeted at AWS and Microsoft would be effective in addressing concerns relating to barriers to switching and multi-cloud on a market-wide basis: such action would directly benefit most UK customers, and any changes made by the two largest suppliers would be expected to result in other providers reflecting any interventions in relation to technical and commercial barriers.

Structure and context

- 10.7 The chapter is structured as follows:
- (a) first, we set out the relevant framework, guidance and process we have followed as well as how the CMA's new powers under the DMCC Act factor into our remedy assessment.
 - (b) second, we set out the recommendations we have decided to make in this market investigation, together with our reasons for making them. This includes explaining the relevant background, summarising stakeholder views, describing the remedy and its intended effects on the AECs, explaining our design considerations, and setting out our assessment on effectiveness and proportionality.
 - (c) third, we include an explanation of potential remedies that we have decided not to pursue with further details provided in appendix W, where we explain why these remedies, if implemented through our remedy-making powers under EA02, would not effectively address the AECs that we have identified.

Framework for our remedies assessment

Our duties

- 10.8 Where we identify an AEC, we are required to determine:
- (a) whether we should take action ourselves, or whether we should recommend others to take action for the purpose of remedying, mitigating or preventing

the AEC or any detrimental effect on customers, so far as it has resulted from, or may be expected to result from, the AEC;

- (b) where we consider that we should take action ourselves, whether that should be through exercising our order-making powers under EA02 or through accepting undertakings from parties;
- (c) what action should be taken, including whether a single remedy or a package of two or more remedies is required.²⁷³⁷

10.9 In coming to a view on remedies, EA02 requires the CMA to ‘in particular have regard to the need to achieve as comprehensive a solution as is reasonable and practicable to the adverse effect on competition and any detrimental effects on customers so far as resulting from the adverse effect on competition’.²⁷³⁸ In satisfying this requirement, the CMA considers how comprehensively its proposed remedies (or a package of remedies) address the AEC and/or resulting detrimental effects on customers, as well as whether the remedies are effective and proportionate.²⁷³⁹

10.10 A detrimental effect on customers is one that results, or may be expected to result, from any AECs and takes the form of:²⁷⁴⁰

- (a) higher prices, lower quality, or less choice of goods or services in any market in the UK (whether or not the market(s) to which the feature or features concerned relate); or
- (b) less innovation in relation to such goods and services.

10.11 Where more than one measure is introduced, the CMA considers the way in which the measures are expected to interact with each other,²⁷⁴¹ which may be complementary in their effectiveness and costs, or may be in tension in some areas.

10.12 The CMA’s interventions seek to remedy, mitigate or prevent the AEC or its detrimental effects on customers. The CMA’s clear preference is to deal comprehensively with the cause or causes of AECs wherever possible, and by this means to significantly improve competitive conditions in a market within a reasonable period of time. However, while generally preferring to address the causes of the AEC, the CMA will consider introducing measures which mitigate the harm to customers created by competition problems, for example if other

²⁷³⁷ EA02, section 134(4); [CC3 \(revised\)](#), paragraphs 325 to 328.

²⁷³⁸ EA02 section 134(6).

²⁷³⁹ [CC3 \(revised\)](#), paragraph 329.

²⁷⁴⁰ EA02, section 134(5).

²⁷⁴¹ [CC3 \(revised\)](#), paragraph 393.

measures are not available, or as an interim solution while other measures take effect.²⁷⁴²

10.13 In assessing remedies, the CMA considers the extent to which different remedy options are likely to be effective in achieving their aims, including their practicability.²⁷⁴³ The effect of any remedy is always uncertain to some degree and in assessing effectiveness:

- (a) the CMA considers the risks associated with different remedies and will tend to favour remedies that have a higher likelihood of achieving their intended effect;²⁷⁴⁴
- (b) a remedy should be capable of effective implementation, monitoring and enforcement. To facilitate this, the operation and implications of the remedy need to be clear to the parties to whom it is directed and also to other interested persons, such as customers, other businesses that may be affected by the remedy, sectoral regulators, and/or any other body which has responsibility for monitoring compliance. The effectiveness of any remedy may be reduced if elaborate monitoring and compliance programmes are required;²⁷⁴⁵
- (c) the CMA generally looks for remedies that prevent an AEC by extinguishing its causes, or that can otherwise be sustained for as long as the AEC is expected to endure.²⁷⁴⁶ It also tends to favour remedies that are expected to show results within a relatively short time,²⁷⁴⁷ and
- (d) where more than one measure is being introduced as part of a package of remedies, the CMA considers the way in which the measures are expected to interact with each other.²⁷⁴⁸

10.14 Where the CMA has identified an effective remedy, it considers its proportionality. In assessing proportionality, the CMA is guided by the following principles. A proportionate remedy is one that:

- (a) is effective in achieving its legitimate aim;
- (b) is no more onerous than needed to achieve its aim;

²⁷⁴² CC3 (revised), paragraphs 330 to 333.

²⁷⁴³ CMA3, Market Studies and Market Investigations: Supplemental guidance on the CMA's approach, paragraph 4.15.

²⁷⁴⁴ CC3 (revised), paragraph 335 and CMA3, paragraph 4.16.

²⁷⁴⁵ CC3 (revised), paragraph 336 and CMA3, paragraph 4.17.

²⁷⁴⁶ CC3 (revised), paragraph 337 and CMA3, paragraph 4.18.

²⁷⁴⁷ The CMA may also consider, when designing remedies, whether to specify a finite duration, for example by means of a sunset clause. A sunset clause will generally specify when individual measures cease to have effect, whether by reference to a specific date or a clearly defined future event. CMA3, paragraph 4.19.

²⁷⁴⁸ CC3 (revised), paragraph 341.

(c) is the least onerous if there is a choice between several effective measures;
and

(d) does not produce disadvantages which are disproportionate to the aim.²⁷⁴⁹

10.15 In judging whether to proceed with a particular remedy, the CMA considers its potential effects – both positive and negative – on those parties most likely to be affected by it, with particular regard to the impact of potential remedies on customers, as well as on those businesses subject to them.²⁷⁵⁰

10.16 Beneficial effects might include lower prices, greater choice, higher quality products/services and/or greater innovation, while the potential negative effects of a remedy may arise in various forms, for example:

- (a) unintended distortions to market outcomes, which may reduce economic efficiency (including dynamic incentives to invest and innovate) and adversely affect the economic interests of customers over the longer term;
- (b) implementation costs, ongoing compliance costs, and monitoring costs (for example, the costs to the CMA or other agencies in monitoring compliance);
and
- (c) if remedies extinguish Relevant Consumer Benefits (RCBs), the amount of RCBs foregone may be considered to be a relevant cost of the remedy.²⁷⁵¹

10.17 However, where businesses have been found to be earning profits persistently in excess of their cost of capital as a direct result of a feature of the market and are likely to continue to do so in the absence of intervention, the CMA will not usually give any significant weight to the anticipated reduction of such profits as a negative effect of a remedy.²⁷⁵²

²⁷⁴⁹ CC3 (revised), paragraph 344.

²⁷⁵⁰ CC3 (revised), paragraph 348.

²⁷⁵¹ The CMA may have regard to the effect of any remedial action any RCBs of the feature(s) of the market(s) concerned (EA02, section 134(7)). For these purposes, a benefit is an RCB if: (a) it is a benefit to customers or future customers in the form of lower prices, higher quality or greater choice of goods or services in any market in the UK, or greater innovation in relation to such goods or services; and (b) the CMA believes that the benefit has accrued, or may be expected to accrue within a reasonable period, as a result of the feature(s) concerned and the benefit was or is unlikely to accrue without the feature(s) concerned (EA02, section 134(8)).

²⁷⁵² CC3 (revised), paragraphs 348 to 353. The CMA will normally collect information from parties about the potential cost of implementing and complying with its remedies. In evaluating such information, the CMA will bear in mind that it has less information than the parties have about how such potential costs have been estimated and that there might be incentives for parties to overstate the cost of those remedies that they do not support. The CMA is likely to place most weight on estimates of implementation and compliance costs where parties have provided a clear explanation of how the estimate was reached, together with supporting evidence as to the assumptions used to derive those estimates.

The CMA's digital markets regulatory powers

- 10.18 We have considered whether the nature and extent of the DMCC Act powers²⁷⁵³ are better suited to addressing the concerns we have identified than the powers available to us under EA02.
- 10.19 When designing the new digital markets competition regime, the UK government stated that 'the existing market investigation tools are designed for one-off interventions and are not well suited to tackling entrenched market power in digital markets, where the market characteristics mean that competition problems are expected to persist over time and require ongoing and proactive oversight.'²⁷⁵⁴ The DMCC Act powers were created with this concern in mind and, in addition to strengthening the existing competition rules, enable the CMA to take a targeted and iterative approach to address issues often seen in digital markets.²⁷⁵⁵
- 10.20 For companies designated as having Strategic Market Status (SMS) in a particular digital activity, the DMCC Act empowers the CMA to impose interventions. This includes all of the interventions available to us under EA02,²⁷⁵⁶ as well as forward-looking conduct requirements (CRs) which are not available to us under EA02.²⁷⁵⁷ The DMCC Act also allows for greater flexibility than is available to us via our order-making powers (eg it is easier to iterate remedies), as well as powers designed to enhance the effectiveness of remedies (eg the ability to test and trial remedies ahead of implementation). As the new regime entrusts the CMA with relevant functions to investigate and monitor digital markets, we consider it better suited to provide an ongoing monitoring and oversight function, particularly where this is required for an extended time period.
- 10.21 The digital markets competition regime applies to firms designated by the CMA as having SMS in relation to one or more digital activities. The DMCC Act provides that a digital activity is the provision of a service by means of the internet, the provision of digital content (which includes software), or any activity which is being carried out for the purposes of providing an internet service or digital content.²⁷⁵⁸
- 10.22 Under the DMCC Act, for a firm to have SMS in respect of a digital activity it must have:

²⁷⁵³ [The Digital Markets, Competition and Consumers Act 2024 \(Commencement No. 1 and Savings and Transitional Provisions\) Regulations 2024](#) (SI 2024/1226) (the 'DMCCA Commencement Regulations'), sections 1(2), 2(1) and 2(2).

²⁷⁵⁴ A new pro-competition regime for digital markets, [Consultation document](#).

²⁷⁵⁵ [Sarah Cardell: Ensuring digital market outcomes that benefit people, businesses and the wider UK economy](#).

²⁷⁵⁶ See DMCC Act, section 51(1), which references Schedule 8 to EA02. (Schedule 8 sets out the permitted remedial measures that can be included in an order remedying any AECs found in a market investigation).

²⁷⁵⁷ For example, conduct requirements can be imposed ex ante to prevent harm from arising.

²⁷⁵⁸ DMCC Act, section 3(1).

- (a) substantial and entrenched market power in a digital activity which is linked to the UK;²⁷⁵⁹
- (b) a position of strategic significance in respect of that activity;²⁷⁶⁰ and
- (c) global turnover of more than £25 billion or UK turnover of more than £1 billion.²⁷⁶¹

10.23 Certain decisions in respect of the new digital markets competition regime are the responsibility of the CMA Board,²⁷⁶² including whether to begin an initial SMS investigation and whether to make an SMS designation.²⁷⁶³

10.24 The DMCC Act sets out that, once the CMA designates a firm with SMS in respect of a digital activity, it may impose CRs on the designated firm to specify how that firm must conduct itself in relation to that digital activity. The CMA may only impose CRs if it considers that it would be proportionate to do so for the purposes of one of the following objectives: fair dealing, open choices, and trust and transparency, having regard to what the CRs are intended to achieve.²⁷⁶⁴

10.25 The CMA can also impose pro-competition interventions (PCIs), following a designation, if the CMA finds that a factor or combination of factors relating to a designated digital activity is having an adverse effect on competition and it would be proportionate to do so. A PCI may take the form of one or both of an order imposing requirements as to how the firm must conduct itself, and/or recommendations to other persons exercising functions of a public nature.²⁷⁶⁵ PCI investigations are subject to a nine-month statutory deadline.²⁷⁶⁶

10.26 The CMA's procedure for SMS investigation and consideration of CRs and PCIs would provide opportunities for affected parties and other stakeholders to make representations on effectiveness and proportionality of potential interventions.

Stakeholder views on our remedies

10.27 We received mixed views from stakeholders on our proposed recommendations in the PDR:

²⁷⁵⁹ DMCC Act, sections 2(1) and (2).

²⁷⁶⁰ DMCC Act, section 2(b).

²⁷⁶¹ DMCC Act, section 7.

²⁷⁶² DMCC Act, section 106.

²⁷⁶³ DMCC Act, section 106. See further [CMA194: Digital markets competition regime guidance](#), paragraphs 9.39-9.40.

²⁷⁶⁴ DMCC Act, section 19.

²⁷⁶⁵ DMCC Act, section 46.

²⁷⁶⁶ DMCC Act, section 50(1). This is subject to the CMA's power to extend the deadline in the circumstances set out in section 104.

- (a) AWS, Microsoft and some third parties were against the proposed recommendations.^{2767 2768 2769} Some of their responses related to a purported absence of evidence of harm, which we address elsewhere in this report. AWS and Rayo questioned the benefits of the DMCC Act powers and the ability of the digital markets competition regime to engage with the issues identified, for example in relation to technical barriers, which we discuss below.²⁷⁷⁰ Some responses queried the legitimacy of an approach focussed on only two providers, which again is discussed below.²⁷⁷¹
- (b) Microsoft submitted that we should recommend that the CMA monitors market outcomes and developments in lieu of an SMS investigation and consult on any further action to the extent required.²⁷⁷² [REDACTED].²⁷⁷³

10.28 Other stakeholders were in favour of the proposed recommendations, noting the benefits of the DMCC Act regime, provided that intervention was ‘urgent’, ‘rapid’, ‘fast-tracked’ or ‘expedited’.²⁷⁷⁴

10.29 In relation to Microsoft’s licensing practices:

- (a) AWS²⁷⁷⁵ and CFSL²⁷⁷⁶ submitted that MI orders, rather than recommendations to prioritise SMS investigations, were more appropriate.
- (b) Google acknowledged the benefits of the DMCC Act regime, such as the ability to iterate remedies,²⁷⁷⁷ but added that the flexibility of the digital markets competition regime to monitor remedies and prevent circumvention must be weighed up against concerns about timing and other uncertainties of addressing the AEC through DMCC Act powers.²⁷⁷⁸ Google said that its

²⁷⁶⁷ [AWS' response to the CMA's PDR](#) dated 28 January 2025.

²⁷⁶⁸ [Microsoft's response to the CMA's PDR](#) dated 28 January 2025.

²⁷⁶⁹ [International Center for Law and Economics' response to the CMA's PDR](#) dated 28 January 2025. [Rayo's response to the CMA's PDR](#) dated 28 January 2025.

[CCIA's response to the CMA's PDR](#) dated 28 January 2025.

²⁷⁷⁰ [AWS' response to the CMA's PDR](#) dated 28 January 2025, paragraph 43.

[Rayo's response to the CMA's PDR](#) dated 28 January 2025, page 1.

²⁷⁷¹ [AWS' response to the CMA's PDR](#) dated 28 January 2025, paragraph 47; [International Center for Law and Economics's response to the CMA's PDR](#) dated 28 January 2025, page 12.

²⁷⁷² [Microsoft's response to the CMA's PDR](#) dated 28 January 2025, paragraph 61.

²⁷⁷³ [REDACTED].

²⁷⁷⁴ [OVHcloud's response to the CMA's PDR](#) dated 28 January 2025, paragraphs 16 and 17; [Blackbox Hosting's response to the CMA's PDR](#) dated 28 January 2025, paragraph 22; [Prolinx's response to the CMA's PDR](#) dated 28 January 2025; [Civo's response to the CMA's PDR](#) dated 28 January 2025, page 1; [Open Cloud Coalition's response to the CMA's PDR](#) dated 28 January 2025; [Simon Hansford's \(former CEO, UK Cloud Ltd.\) response to the CMA's PDR](#) dated 28 January 2025.

²⁷⁷⁵ [Summary of hearing with Amazon Web services \(AWS\)](#), 12 March 2025, paragraphs 44-45.

²⁷⁷⁶ CFSL's submission to the CMA [REDACTED]. CFSL also submitted that if the CMA was not minded to use its MI powers, another option, albeit a less suitable one, was to impose remedies as part of interim measures under the CA98. As noted in chapters 1 and 7 we are not investigating a breach of the CA98 and therefore interim measures under this legislative provision are not available to us. We consider an effective proportionate remedy to be a recommendation to the CMA Board to consider using its DMCC Act powers and have therefore not recommended action under CA98.

²⁷⁷⁷ [Google's response to the CMA's PDR dated 28 January 2025](#), paragraphs 15 and 18.

²⁷⁷⁸ Google's submission to the CMA [REDACTED].

primary position is that remedies can and should be imposed on Microsoft using the CMA's market investigation order-making powers.²⁷⁷⁹

- (c) Google also submitted that, were we to remedy Microsoft's licensing practices through a recommendation to the CMA Board, it would be necessary to also put in place certain 'interim' measures, intended to prevent the situation from worsening pending any interventions under DMCC Act powers.²⁷⁸⁰ This is discussed in more detail below.
- (d) Two other stakeholders we spoke to suggested that Microsoft's position in related software markets would be appropriate for consideration under the SMS regime.²⁷⁸¹

10.30 Our views on these points are set out below.

Our assessment of effectiveness of these remedies

10.31 As set out in our guidance, the effect of any remedy is always uncertain to some degree, and in evaluating the effectiveness of potential remedies we will consider the risks associated with different remedy options and will tend to favour remedies that have a higher likelihood of achieving their intended effect.²⁷⁸²

10.32 We recognise that recommendations are not binding on the party to which they are addressed, and that this represents an intrinsic risk to their effectiveness as a remedy. A recommendation may not be accepted, may not be implemented in a way that is consistent with our intentions, or may become redundant following a change of policy.²⁷⁸³ We have reflected these risks in our effectiveness assessment below, in particular by focusing on:

10.33 how the recommendations address the identified harms; and

10.34 the likelihood of the recommendations being acted upon and, if so, over what time period.

How the recommendations are effective in addressing the harm

10.35 We have identified certain features that in combination are giving rise to AECs in cloud services in the UK. These features are high levels of market concentration,

²⁷⁷⁹ Google's submission to the CMA [🔗].

²⁷⁸⁰ Google's submission to the CMA [🔗]. See further [Google's response to the CMA's PDR](#) dated 28 January 2025, paragraphs 10-11, 20-21.

²⁷⁸¹ [CCIA's response to the CMA's PDR](#) dated 28 January 2025, page 3. [🔗].

²⁷⁸² [CC3 \(revised\)](#), paragraph 335 and [CMA3](#), paragraph 4.16.

²⁷⁸³ [CC3 \(revised\)](#), Annex B, paragraph 97.

barriers to entry and expansion, technical and commercial barriers to switching and multi-cloud and Microsoft's licensing practices.

10.36 We consider that these features and the AECs to which they give rise represent the types of issue that the digital markets competition regime was designed to address.²⁷⁸⁴ The most relevant aspects of these powers for our consideration of recommendations to the CMA Board include:

- (a) a wide range of intervention powers;
- (b) flexibility to allow for future variation and iteration;
- (c) powers to test and trial potential interventions; and
- (d) ongoing regulatory oversight including a wide range of investigatory powers.

10.37 We consider each of these below.

Range of intervention powers

10.38 In terms of the range of intervention powers, if a firm is designated as having SMS in a digital activity, the CMA can impose competition requirements in the form of CRs or PCIs on the firm.²⁷⁸⁵ A PCI may implement a similar set of remedy options as those available to us in this market investigation via pro-competition orders (PCOs) or making recommendations to another body.²⁷⁸⁶

10.39 In the event that the CMA designated AWS and Microsoft with SMS, it would be for the CMA Board or its delegates to consider the most appropriate interventions to impose, as well as the sequence in which any interventions may be put into place.²⁷⁸⁷

²⁷⁸⁴ For example, 'the Digital Markets Unit should prioritise designation assessments for activities where there are likely to be significant network effects, economies of scale and scope, and/or there are high fixed costs to entering the market. These are the activities most likely to tend towards concentration and entrenched market power, and therefore where potential SMS firms could emerge and which are most likely to require the tools of the new regime' and 'Digital firms can entrench, and take advantage of, their strategic position by creating an 'ecosystem' of accompanying products and services that expands into new markets and undermines their competitors.' A new pro-competition regime for digital markets, [Consultation document](#).

²⁷⁸⁵ Whilst CRs are largely limited to the digital activity itself, Section 20(3)(c) of the DMCC Act allows the CMA to impose a CR that applies to an SMS firm's conduct in an activity other than the relevant digital activity. Under this permitted type, the CMA may impose CRs for the purpose of preventing an SMS firm from carrying on activities other than the relevant digital activity in a way that is likely to materially increase the SMS firm's market power or materially strengthen its position of strategic significance in relation to the relevant digital activity. The CMA can implement a PCI in relation to any part of the SMS firm's business, where the CMA decides that doing so would address the competition problem.

²⁷⁸⁶ The DMCC Act expressly references Schedule 8 of EA02 which sets out a list of structural and behavioural remedies. However, there are some additional aspects to PCIs, which we describe in more detail above.

²⁷⁸⁷ We also note that our guidance states that 'the remedy that is ultimately selected need not directly address every feature identified, if for example, tackling a subset of features directly would be sufficient to generate effective competition and thereby remedy the AEC.' (CC3, paragraph 332.) As explained below, we consider the CMA's DMCC Act powers would be well-suited to consideration of the sequencing of any interventions.

Flexibility for future variation and iteration

- 10.40 The CMA can impose CRs to guide the practices of SMS firms, and it has the ability to iterate CRs over time to ensure they are effective. CRs are also well-suited to a principles-based approach. The CMA has set out the conditions in which it is more likely to impose outcome-focused CRs compared with action-focused CRs, as well as where it is more likely to rely on higher-level requirements compared with more detailed requirements.²⁷⁸⁸
- 10.41 The CMA also has a duty to consider on an ongoing basis whether to impose, vary or revoke CRs,²⁷⁸⁹ as well as a duty to review any PCO at a pre-identified date to determine whether to revoke or replace it thereby allowing the CMA to proactively iterate.²⁷⁹⁰ This flexibility to respond to changing circumstances, and iterate to respond to firm behaviour, is well-suited to cloud services.
- 10.42 In our PDR, we discussed interactions with other laws and regulations, primarily the EU Data Act's provisions regarding egress fees and interoperability, and acknowledged the need to consider how they might interact with any remedies we might impose. We received submissions in response to this:
- (a) AWS and Microsoft submitted that it would be more appropriate to monitor how the developments with respect to the EU Data Act unfold before taking any further action under the DMCC Act.²⁷⁹¹
 - (b) AWS also said that the 'host of problems and risks associated with the proposed standardisation remedies in the PDR [...] would equally apply under the digital markets competition regime'.²⁷⁹²
 - (c) Rayo said that challenges with interoperability requirements 'would not suddenly be eliminated by pushing implementation towards the CMA's new DMU'.²⁷⁹³
 - (d) in relation to egress fee remedies, Google stated that 'any remedies should align with and be no more onerous than the equivalent requirements under the EU Data Act'.²⁷⁹⁴
- 10.43 We agree that it would be important to consider how regulations elsewhere might interact with any action taken by the CMA. The flexibility of the DMCC Act powers would allow the CMA to consider regulatory developments and their potential

²⁷⁸⁸ [CMA194: Digital markets competition regime guidance](#), paragraph 3.29.

²⁷⁸⁹ DMCC Act, section 25.

²⁷⁹⁰ DMCC Act, section 55(1).

²⁷⁹¹ [AWS' response to the CMA's PDR](#) dated 28 January 2025, paragraph 48; [Microsoft's response to the CMA's PDR](#) dated 28 January 2025, paragraph 61.

²⁷⁹² [AWS' response to the CMA's PDR](#) dated 28 January 2025, paragraph 44.

²⁷⁹³ [Rayo's response to the CMA's PDR](#) dated 28 January 2025, page 1.

²⁷⁹⁴ [Google's response to the CMA's PDR](#) dated 28 January 2025, paragraph 22b.

impact on any interventions addressing issues arising in markets in the UK on an ongoing basis and to react where appropriate by iterating or varying those interventions. This is consistent with the CMA's guidance on the implementation of the digital markets competition regime²⁷⁹⁵ and the UK government's recent Strategic Steer to the CMA to 'where appropriate, seek to ensure parallel regulatory action is timely, coherent and avoids duplication'.²⁷⁹⁶

- 10.44 We recognise that some risks to implementing effective and proportionate remedies could remain even where the CMA acts using its DMCC Act powers, for example with certain specification challenges in relation to technical barriers. However, we consider the DMCC Act powers are better suited to addressing these risks particularly where there is the ability for an intervention to be tested and trialled, iterated over time and where consideration can be given to the sequence in which to put interventions into place.²⁷⁹⁷ This means that an effective solution to the competition concerns we have identified is more likely to be developed using these powers.

Ability to test and trial

- 10.45 Regarding powers to test and trial potential interventions in the digital markets competition regime:
- (a) the CMA can require SMS firms to perform a specified demonstration or test.²⁷⁹⁸ This can include requiring a firm to demonstrate a technical process, such as how a cloud provider identifies and authenticates a user, and can require a firm to vary its usual conduct, for example by making it easier for third party IAM providers to authenticate users of a cloud service. For PCIs, the CMA may impose requirements in a PCO to test and trial different remedies or remedy design options to gain practical evidence on their effectiveness, including for specific user or customer groups.²⁷⁹⁹
 - (b) the testing and trialling powers would assist the assessment of potential remedies for certain technical changes to reduce frictions, eg any standardisation in relation to APIs. Such powers would assist in properly understanding how customer and user behaviour is likely to be impacted by any changes ahead of any remedial action being imposed, which would be

²⁷⁹⁵ [CMA194: Digital markets competition regime guidance](#), paragraph 3.32. 'The CMA will also have regard to the extent to which the CR or combination of CRs is coherent with any other CRs imposed on the SMS firm or on other SMS firms operating in the same or similar digital activities. The CMA may also consider, as appropriate, coherence with interventions under other CMA tools, such as cases under the Competition Act 1998 ('CA98') and requirements imposed by other UK regulators. In addition, the CMA may take into account actions and experiences of other regulators or legislators internationally.'

²⁷⁹⁶ [Strategic steer to the Competition and Markets Authority - GOV.UK](#)

²⁷⁹⁷ For example, in the case of the CMA's provisional decision to designate Google with SMS in relation to search, the CMA has published a 'Roadmap' of how it would propose to prioritise actions taken during the first half of any designation period. See [CMA press release dated 24 June 2025](#).

²⁷⁹⁸ DMCC Act, section 69(5).

²⁷⁹⁹ DMCC Act, see sections 51(3) and 51(4).

particularly important where the remedies would be costly to develop on a large scale, difficult to reverse, or could result in longer term distortion or disruption.

Ongoing regulatory oversight

- 10.46 The DMCC Act provides the CMA with a range of investigatory powers which it can use when administering and enforcing the digital markets competition regime. Monitoring is a key part of the CMA's role in overseeing the digital markets competition regime. It would allow the CMA to respond quickly where firms fail to comply with competition requirements, considering whether enforcement action is warranted for non-compliance and informing whether new or varied competition requirements may be necessary, including removing or narrowing particular interventions. This flexibility would enable the CMA to take account of and react to relevant changes in relation to cloud services. These ongoing monitoring and enforcement powers would be important should the CMA implement remedies in relation to technical and commercial barriers and Microsoft's licensing practices.
- 10.47 These powers would support an effective and comprehensive solution to the competition concerns we have identified:
- (a) any remedies to address technical barriers under DMCC Act powers²⁸⁰⁰ would need to be sufficiently well-specified but also sufficiently flexible to remain appropriate for a wide range of uncertain future changes in cloud services, which is technically complex. These remedies would also need to adapt as new technical mitigations are developed, or the incentives on SMS firms change as the market matures. Therefore, a key element in the design of any remedies would be the ability to iterate them so that they would continue to be effective, adapting to changing market conditions and preventing any harmful unintended consequences.
 - (b) any remedies to address Microsoft's licensing practices under DMCC Act powers²⁸⁰¹ would need to be implemented as a combination of principles-based and rules-based interventions to comprehensively address the relevant licensing practices. Given the complexity of the potential remedies, a key element in the design would be the ability to refine, calibrate or amend them over time.

²⁸⁰⁰ We note that a number of the permitted types of CRs appear to be relevant to the potential remedies we have considered. For example, requirements preventing an SMS firm from restricting interoperability between the relevant service or digital content and products offered by other undertakings, section 20 DMCC Act.

²⁸⁰¹ We note that a number of the permitted types of CRs appear to be relevant to the potential remedies we have considered. For example, requirements obliging an SMS firm to trade on fair and reasonable terms. DMCC Act, section 20.

Effectiveness of remedies in addressing market-wide harms

- 10.48 We have considered whether our recommendations to the CMA Board would effectively address the competition concerns we have identified arising from technical and commercial barriers on a market-wide level. Our view is that interventions aimed at influencing the behaviour of AWS and Microsoft would be an effective remedy because:
- (a) AWS and Microsoft collectively account for 30-50% of PaaS and 60-80% of IaaS revenues.²⁸⁰² Therefore, any measures addressed to AWS and Microsoft would directly benefit the majority of UK customers.
 - (b) there would be an indirect effect where the competitive conditions and commercial pressures would be expected to result in other providers reflecting any interventions imposed on the larger suppliers, further reinforcing the impact of these interventions across the relevant markets.
- 10.49 AWS said that the ‘indirect effect’ of other providers following suit is ‘an unsubstantiated prediction which cannot be assumed’.²⁸⁰³
- 10.50 In contrast, Google submitted that a remedy directly imposed on AWS and Microsoft is likely to change the commercial conditions in the wider sector, and that other cloud providers such as Google will have little option commercially but to respond to any changes imposed upon AWS and Microsoft.²⁸⁰⁴
- 10.51 We have considered evidence showing that cloud providers have been influenced by the actions of other providers, details of which are provided in appendix W.
- 10.52 Our view is that competitive and commercial pressure would result in other providers reflecting interventions imposed on the largest suppliers.

Likelihood of the recommendations being acted upon

- 10.53 Our guidance states that when evaluating the effectiveness of a recommendation as a proposed remedy, the CMA must consider:
- (a) the likelihood that the recommendation will be acted on; and
 - (b) the time period over which this might be expected to occur.²⁸⁰⁵
- 10.54 In reaching its view, the CMA must have regard to:

²⁸⁰² Shares of supply are by reference to revenue, see chapter 3 (market shares).

²⁸⁰³ [AWS' response to the CMA's PDR](#) dated 28 January 2025, paragraph 58.

²⁸⁰⁴ [Google's response to the CMA's PDR](#) dated 28 January 2025, paragraphs 4 and 22.

²⁸⁰⁵ [CC3 \(revised\)](#), paragraph 391.

- (a) the stated policy of the body to which the recommendation is to be directed; and
- (b) the possibility that the stated policy may change, either in light of our recommendation or subsequent events.²⁸⁰⁶

10.55 In order to assess the likelihood that our recommendations will be acted upon, we set out below:

- (a) the criteria that the CMA Board would use when deciding whether to designate a firm as having SMS; and
- (b) any factors affecting when our recommendations may be acted on.

10.56 We have not formally assessed whether AWS and/or Microsoft meet the specific SMS designation criteria because these are for the CMA Board to consider when deciding whether to prioritise a firm's digital activity for SMS investigation. However, there is good evidence indicating that each of the SMS criteria may be met in respect of each of AWS' and Microsoft's digital activities in relation to cloud services:²⁸⁰⁷

- (a) each of AWS and Microsoft has significant market power in the supply of cloud services in the UK and that is likely to endure.²⁸⁰⁸
- (b) AWS' and Microsoft's respective positions in the provision of cloud services appear to be strategically significant.
- (c) the provision of cloud services itself is a strategically important digital activity²⁸⁰⁹ which places each of AWS and Microsoft in a position where they play an important role in the business of a significant number of other firms.²⁸¹⁰
- (d) each of Amazon and Microsoft have a global turnover in excess of £25 billion.²⁸¹¹

²⁸⁰⁶ CC3 (revised), paragraph 391.

²⁸⁰⁷ The likelihood of a recommendation being implemented is one of the criteria of we consider for the specification of any recommendation; see CC3 (revised), Annex B, paragraph 100.

²⁸⁰⁸ See chapter 9.

²⁸⁰⁹ The importance of cloud services has been recognised by a number of sector regulators (see Regulatory Barriers section of the chapter 5), and the UK government has recently classed data centres as Critical National Infrastructure.

²⁸¹⁰ To assess whether an undertaking has a position of strategic significance in respect of the digital activity, the CMA must assess whether one or more of the four following conditions has been met: i) the firm has achieved a position of significant size or scale in respect of the digital activity, ii) a significant number of other firms use the digital activity as carried out by the firm in carrying on their business, iii) the firm's position in respect of the digital activity would allow it to extend its market power to a range of other activities, iv) the firm's position in respect of the digital activity allows it to determine or substantially influence the ways in which other firms conduct themselves, in respect of the digital activity or otherwise.

²⁸¹¹ Amazon and Microsoft's public accounts. We note that the DMCC Act, section 7(2), states that the £25 billion global turnover condition is met if the CMA estimates that 'the total value of the global turnover of an undertaking or, where the undertaking is part of a group, the global turnover of that group in the relevant period exceeds £25 billion.'

- 10.57 In terms of the stated policy, the CMA has said that it is focused on progressing current SMS investigations for the remainder of 2025. The CMA has also said that it will keep under review possible options for further designation investigations and anticipates that these will be considered by the CMA Board in early 2026.²⁸¹² We understand that its future investigation decisions would take into account relevant information and evidence on the UK's digital markets, as set out in its digital markets guidance,²⁸¹³ and that this would include the findings and recommendations from this inquiry.
- 10.58 In considering SMS investigations into AWS and Microsoft's cloud services activities, we anticipate that the CMA board would also need to assess the actions being taken by the European Commission in relation to Data Act requirements regarding cloud egress fees and interoperability, in order to assess whether any CMA interventions, if appropriate, would be timely, coherent and avoid any unnecessary duplication with requirements under EU law. This is in line with the UK government's strategic steer to the CMA.²⁸¹⁴
- 10.59 We also recognise that the CMA Board would need to take into account any market developments prior to initiating an SMS investigation, as well as factors beyond the scope of this report when making its prioritisation decision, including the urgency of any other potential SMS investigations. Therefore, we recommend that in the period pending the CMA's decision on future designation investigations, developments in the relevant markets, including any actions taken by cloud providers, are kept under review and that the CMA Board factors these into its prioritisation decision.
- 10.60 In terms of timing, in addition to the CMA's recent announcement, we note that once launched, the CMA has to complete SMS investigations within a statutory timeframe of nine months.²⁸¹⁵ The CMA would be able to impose an initial set of CRs on designation or as soon as practicable afterwards,²⁸¹⁶ while any PCIs would follow a further nine-month investigation.²⁸¹⁷
- 10.61 Our guidance states that we can make recommendations that may not be implemented immediately where we judge that these are nonetheless likely to be more effective than other possible remedies.²⁸¹⁸ While we recognise that there is some uncertainty around their implementation, we believe that recommendations

²⁸¹² [CMA takes first steps to improve competition in search services in the UK - GOV.UK](#)

²⁸¹³ [CMA194: Digital markets competition regime guidance](#)

²⁸¹⁴ [Strategic steer to the Competition and Markets Authority - GOV.UK](#)

²⁸¹⁵ DMCC Act, section 14(2). This is subject to the CMA's power to extend the deadline in the circumstances set out in section 104.

²⁸¹⁶ [CMA194: Digital markets competition regime guidance](#), paragraph 3.40.

²⁸¹⁷ DMCC Act, section 50(1). This is subject to the CMA's power to extend the deadline in the circumstances set out in section 104.

²⁸¹⁸ [CC3 \(Revised\)](#), paragraph 391.

are the only effective remedy capable of comprehensively addressing the AECs we have found in this investigation.

10.62 In view of the above, we are satisfied that our recommendations to the CMA Board are an effective remedy.

Proportionality assessment of these remedies

10.63 In this section we set out our assessment of whether the remedies would be proportionate to achieve their aim. We do so by considering whether each remedy would:²⁸¹⁹

- (a) be effective in achieving its legitimate aim;
- (b) not be more onerous than needed to achieve its aim;
- (c) be the least onerous if there were a choice between several effective measures; and
- (d) not produce disadvantages which are disproportionate to the aim.

Effective in achieving its legitimate aim

10.64 We discussed above the effectiveness of the recommendations, and conclude that they would be effective in achieving their aim.

No more onerous than needed

10.65 As noted above, given that Microsoft and AWS are the two largest providers of cloud services, focusing any interventions on them would directly benefit the majority of UK customers and alter the wider competitive conditions, thus achieving an effective remedy without being more onerous than is necessary to do so.

10.66 In addition, were the recommendations to be acted upon, key aspects of the digital markets competition regime (such as the ability to enact tailored interventions and the explicit proportionality requirements for any intervention under the DMCC Act)²⁸²⁰ would further ensure that the remedy would be no more onerous than needed.

²⁸¹⁹ CC3 (Revised), paragraph 344.

²⁸²⁰ See discussion below on whether the remedy would produce disadvantages which are disproportionate to the aim.

Least onerous if there is a choice between several effective measures

- 10.67 If the CMA is choosing between two remedy measures which are both effective, it should choose the remedy measure that imposes the least cost or is least restrictive. We have not identified any other effective remedy to the AECs that we have found, so this choice does not arise.

Does not produce disadvantages which are disproportionate to the aim

- 10.68 We consider that the recommendations would not produce disadvantages which are disproportionate to their aims. The recommendations do not impose any immediate obligations on AWS or Microsoft. Whilst we accept that potentially being subject to an SMS designation and if so, subject to CRs and/or PCIs if appropriate, could place certain costs on AWS and Microsoft, there are proportionality requirements built into the new digital markets competition regime. The CMA can only impose CRs or PCIs if it would be proportionate to do so.^{2821,2822}
- 10.69 Since the publication of our PDR, AWS has said that our recommendation to open SMS investigations would ‘distort competition by creating asymmetric obligations for AWS and Microsoft’.²⁸²³ It also said that SMS designation of providers under the DMCC regime would create an uncertain regulatory environment for the UK cloud sector.²⁸²⁴
- 10.70 Microsoft submitted that it would be ‘highly distortive to competition in the UK cloud market’.²⁸²⁵
- 10.71 ICLE submitted that ‘[e]ven if one accepts that there are competition concerns in the UK cloud-services market, the remedies must be narrowly tailored and evidence-based. Imposing broad behavioural obligations or regulatory oversight via an SMS designation is a far-reaching step that could reshape the incentives in this dynamic sector.’²⁸²⁶
- 10.72 CCIA submitted that SMS designation creates ‘significant uncertainty for regulated entities’ and that the lack of evidence of specific harms ‘means there is limited

²⁸²¹ The CMA may only impose a conduct requirement or a combination of conduct requirements on a designated undertaking if it considers that it would be proportionate to do so for the purposes of one or more of the following objectives — (a) the fair dealing objective, (b) the open choices objective, and (c) the trust and transparency objective, having regard to what the conduct requirement or combination of conduct requirements is intended to achieve; DMCC Act, sections 19(5)-(8).

²⁸²² The CMA may make a PCI in relation to a designated undertaking where, following a PCI investigation, the CMA considers that [...] it would be proportionate to make the PCI for the purposes of remedying, mitigating or preventing the adverse effect on competition; DMCC Act, section 46.

²⁸²³ [AWS' response to the CMA's PDR](#) dated 28 January 2025, paragraph 7.

²⁸²⁴ AWS' submission to the CMA [§].

²⁸²⁵ [Microsoft's response to the CMA's PDR dated 28 January 2025](#), paragraph 116.

²⁸²⁶ [International Center for Law and Economics's response to the CMA's PDR](#) dated 28 January 2025, page 12.

upside to compensate for uncertainty that would diminish the case for investment in UK digital infrastructure.’²⁸²⁷

- 10.73 We have found AECs in the cloud services markets in the UK and our view is that recommendations to the CMA Board to use its DMCC Act powers would not produce disadvantages that are disproportionate to the aim of addressing these AECs.
- 10.74 Targeted interventions in the form of recommendations to the CMA Board in relation to AWS and Microsoft is commensurate with the significant unilateral market power we have found both to hold. As above, we also think that competitive and commercial pressure would result in other providers reflecting any interventions imposed on the largest suppliers.
- 10.75 The DMCC Act powers enable tailored interventions, and impose safeguards against damaging or distortive approaches, notably via an explicit proportionality requirement on any potential intervention. Accordingly, the CMA Board would need to evidence the proportionality of any intervention before it could be implemented on either AWS or Microsoft. When considering the proportionality of CRs, the CMA takes account of potential effects – both positive and negative – on those affected, which may include the SMS firm, third parties, consumers and business users, as well as wider effects such as the risks of market distortions or unintended consequences.²⁸²⁸ It also has a duty to consider on an ongoing basis whether to impose, vary or revoke CRs.²⁸²⁹
- 10.76 We also received submissions stating that particular interventions would be disproportionate. Insofar as they relate to the implementation of such actions through our MI powers, these are addressed in appendix W. Insofar as they relate to potential future interventions under DMCC Act powers, these would be considered by the CMA Board should it consider imposing such interventions following any SMS designation.

Other remedy options that we have decided not to pursue

- 10.77 As noted above, we have considered a number of potential remedies which we have decided not to pursue. These are:
- (a) measures we have decided not to pursue under the remedy-making powers provided by EA02 for:
 - (i) technical barriers;

²⁸²⁷ CCIA's response to the CMA's PDR dated 28 January 2025, page 3.

²⁸²⁸ CMA194: Digital markets competition regime guidance, paragraph 3.34.

²⁸²⁹ DMCC Act, section 25.

- (ii) egress fees;
 - (iii) Microsoft licensing practices; and
 - (iv) time-limited interventions relating to Microsoft licensing practices.
- (b) structural remedies; and
 - (c) utility-style regulation.

10.78 We consider each of these to explain our reasoning.

Measures we have decided not to pursue under the remedy-making powers provided by EA02

- 10.79 We turned our minds to possible remedies that might address the AECs and features we have found using the powers available to us under EA02. While implementing the remedies using these powers could, in principle, be capable of addressing the features we have found are giving rise to AECs, these would not result in as comprehensive a solution as is reasonable and practicable.
- 10.80 As noted above, our remedy-making powers under EA02 are more suited to one-off interventions and not as well-suited to addressing certain issues arising in evolving digital markets. This is the case in relation to the features we have found in this investigation.
- 10.81 In particular, implementing the remedies we have assessed would be challenging as these remedies would suffer from limitations such as the inability to iterate the remedy design on an ongoing basis,²⁸³⁰ the inability to test and trial remedies of this nature,²⁸³¹ and the inability to provide ongoing regulatory oversight of the remedy, which would be important aspects of an effective and comprehensive remedy. In addition, implementing remedies using our powers under EA02 in

²⁸³⁰ We note that EA02, section 162A, introduced through section 139 of the DMCC Act, gives the CMA new powers to review the effectiveness of orders and undertakings implemented following the determination of market investigation references under EA02. Under this provision, the CMA will have jurisdiction to amend an undertaking or order if found ineffective within 10 years of the publication of the final report and where the CMA last took relevant action no less than two years ago. This provision applies to undertakings accepted or orders made after 1 January 2025 and therefore would apply to any market investigation order implementing remedies following the publication of our final report in this investigation ([DMCCA Commencement Regulations](#), Schedule, paragraph 11). However, as has been set out above and in appendix W, we consider that a remedy for the AECs identified in this investigation would require greater flexibility to iterate and refine the remedy design than afforded by this change.

²⁸³¹ EA02, sections 161B to 161E and 162B, introduced through section 138 of the DMCC Act, give the CMA new powers to conduct trials to assess the likely effectiveness of final undertakings and orders that the CMA is minded to accept or impose following the determination of a market investigation reference under EA02. This provision applies in cases where a market investigation report is issued after 1 January 2025 and therefore would potentially apply to any undertakings accepted or order made in this market investigation ([DMCCA Commencement Regulations](#), Schedule, paragraph 10). However, as set out above and in appendix W, we consider that the ability to conduct trials of any order or undertakings intended to address the AECs identified would not resolve concerns in this case as such trials are intended to be used, unless and until the Secretary of State makes orders to broaden their scope, only for remedies concerning the provision or publication of information to consumers, which is not currently applicable to this market investigation.

addition to the proposed recommendations to the CMA Board would likely add complexity and the risk of regulatory divergence which would adversely impact the effectiveness of any package of remedies. Further details on why we are not pursuing these potential remedies are provided in appendix W and are summarised below.

Technical barriers

- 10.82 In appendix W, we have set out our assessment and views on each of the potential remedies addressing technical barriers that we have identified during the course of this investigation. In particular:
- (a) we have examined the case for standardisation of IaaS, PaaS, ancillary services, and APIs. We consider that in designing any standardisation remedies there would be a need for a careful assessment of the benefits of a standard against its risks, in order to correctly identify — especially in light of the danger of distortion risks — the services for which the risks would be outweighed by the benefits.
 - (b) for more differentiated services, such as some PaaS services, the distortion risks of standardisation could be particularly high. We consider that even for ancillary services and IaaS and also for APIs, where in principle there is greater scope for standardisation, there is still benefit from flexibility and iteration in the implementation of any remedy. This would be important in addressing residual distortion risks, and technological or service design changes that the standards might need to adapt to. It would also benefit from testing and trialling of standards so that they could be refined or recalibrated efficiently.
 - (c) accordingly, a key element in the design of any standardisation remedy would be the ability to iterate the remedy in response to technical changes, new innovations, the re-introduction of technical frictions in an attempt to circumvent the aim of the remedy or changes arising from regulatory requirements elsewhere. This would require a level of flexibility to the design, implementation, monitoring and enforcement of remedies that would be challenging to achieve through a market investigation order. In particular, we note that the EU Data Act includes provisions which allow the EU to introduce standards that require cloud providers to ensure compatibility with common specifications based on open interoperability specifications or harmonised standards for interoperability within 12 months of a standard being introduced.²⁸³² It is too early to assess the types of standards that may be required or to what extent any standards implemented in response to the

²⁸³² EU Data Act: Article 30: 3. [Regulation - EU - 2023/2854 - EN - EUR-Lex](#), last accessed on 24 October 2024.

EU Data Act would impact cloud services in the UK. To mitigate the potential for regulatory fragmentation, effective monitoring and iteration of any remedies to address any inconsistencies and/or contradictions that might arise would be essential.

- (d) we consider that a transparency remedy could in principle be implemented through a market investigation order but we are of the view that this would not be an effective remedy on its own (even on an interim basis), as information transparency remedies would not reduce the prevalence or impact of technical barriers that are currently present. There would also be limited benefit to customers who have already migrated to the cloud or to more sophisticated customers, who have some pre-existing awareness of lock-in risks.

- 10.83 In light of these assessments, there are likely to be material risks associated with these remedies such that it would be difficult to achieve a comprehensive solution through the use of the remedy-making powers under EA02.
- 10.84 We consider that if the CMA were to designate AWS and Microsoft with SMS in respect of cloud services, it would have the ability to test and trial remedies with a range of market participants, as well as to iterate remedies over time, which we consider would likely address many (if not all) of the major risks we have identified in our assessment.

Egress fees

- 10.85 In appendix W, we set out our assessment of a potential price control remedy, in the form of a ban or cap, for egress fees for switching and multi-cloud (with supplemental information transparency requirements).
- 10.86 We have considered submissions on the potential remedies, including those that highlighted the potential for regulatory divergence with the EU Data Act. It is relevant whether any remedy addressing egress fees diverged across jurisdictions as this may impact the effectiveness of such a remedy and we have therefore taken this into account.
- 10.87 Divergent egress fee pricing requirements in different jurisdictions (eg divergence from the EU Data Act's provisions relating to egress fees) may cause complexity for customers and compliance issues for providers. We note that the approach to setting an at-cost cap for multi-cloud egress under the EU Data Act's provisions (coming into effect in September 2025) is not yet clear and cloud providers have not yet determined how they will comply with these provisions and/or the egress fee rates that will be applied. We also received submissions highlighting the

complexity of determining a cost rate for switching and multi-cloud egress data transfers.²⁸³³

- 10.88 We consider that a potential price control remedy in the form of a ban on egress fees for switching and multi-cloud use would be analytically straightforward, with minimal specification or circumvention risk, compared to the complexity of a price control capping egress fees at cost. However, a ban on egress fees would diverge with the EU Data Act,²⁸³⁴ whereas an at-cost cap would have greater alignment across jurisdictions. Therefore, for either approach for a price control remedy for egress fees, implementation via our remedy-making powers under EA02 would not provide sufficient flexibility to enable the CMA to take into account developments in the EU when determining the level of fees allowed.
- 10.89 We consider that the use of the CMA's DMCC Act powers would provide the flexibility to iterate interventions in relation to egress fees if issues materialised due to developments elsewhere including possible regulatory divergence.
- 10.90 We also consider that there would be risks to the effectiveness of a remedies package that involved a process of implementing, monitoring and maintaining any remedies implemented under the DMCC Act in parallel with an egress fees remedy under EA02, and the coherence of any substantive obligations being placed on AWS and Microsoft were the CMA to designate these parties with SMS status. In particular:
- (a) the implementation and enforcement of interventions through the market investigation could reduce the flexibility available to the CMA to design and implement a complete and holistic set of effective interventions following any SMS designation.
 - (b) although we consider the risks associated with monitoring and enforcement of a well-specified egress fee remedy to be relatively low, the ongoing monitoring and enforcement of such a remedy alongside parallel monitoring and enforcement of any potential obligations imposed on AWS and Microsoft under the DMCC Act, would increase ongoing complexity. This approach would also raise a risk of contradictory or conflicting obligations and/or approaches to monitoring and enforcement by the CMA, as well as

²⁸³³ We also received a submission from Google that the CMA would be able to draw on the expertise built up by providers in anticipation of the EU Data Act to determine a cost rate for egress data transfers, however as noted a rate-setting methodology has not yet been commonly established. [Google's response to the CMA's PDR dated 28 January 2025](#), paragraph 29-30. We found AWS, Google, Microsoft and IBM to have indicated that there are substantial difficulties in determining an 'at cost' rate to apply to egress data transfers for compliance with the EU Data Act provisions in relation to 'in parallel use'. Note of meeting with [REDACTED]; [AWS' response to the Egress fees working paper](#), paragraphs 26 and 28; [Google's response to Egress fees working paper](#), Annex response (f); [REDACTED]; Note of meeting with Google [REDACTED]; Note of meeting with Microsoft [REDACTED]; IBM's response to the CMA's information request [REDACTED].

²⁸³⁴ A ban would diverge from the EU Data Act's provisions in relation to egress fees for in parallel use data transfers once in effect, and for egress fees for switching data transfers, up until the EU Data Act's phased ban comes into effect in 2027 (although we note that Microsoft, AWS and Google have voluntarily introduced free switching programmes globally, which we assessed in chapter 6 of this report).

potentially resulting in inefficient use of CMA resources through separately implementing and maintaining remedies under both EA02 and the DMCC Act.

- 10.91 We consider that implementing an egress fees remedy pending the CMA acting on our recommendations would not be an effective remedy given the concerns we have set out above about the risk of regulatory divergence and the need for any remedy to be able to take account of developments arising from the EU Data Act in particular.
- 10.92 We consider that these risks would not arise following any SMS designation of AWS and Microsoft in respect of cloud services, given the advantages discussed above of using the DMCC Act powers.

Microsoft licensing practices

- 10.93 In appendix W, we have considered stakeholder submissions that while there may be benefits to addressing Microsoft's licensing practices using DMCC Act powers, it would still be possible to implement effective remedies under the MI, to ensure a swift resolution for customers.
- 10.94 We consider that in order to be effective, a package of licensing remedies would need to comprise a combination of rules-based and principles-based measures. We have identified that this could potentially include:
- (a) an obligation on Microsoft to apply a fair, reasonable and non-discriminatory (FRAND) approach in relation to pricing its software products regardless of which cloud they are hosted on.²⁸³⁵ It would also be obliged to provide greater information transparency;
 - (b) restrictions on Microsoft's ability to favour its own cloud through licensing practices which grant unequal access to software products and product functionality depending on which cloud the software products are deployed on; and
 - (c) changes to contractual licensing practices relating to the transfer and/or deployment by end customers of previously purchased software products on the cloud of their choice.
- 10.95 We have identified risks with implementing these forms of measures using our remedy-making powers under EA02:

²⁸³⁵ A FRAND approach would require Microsoft to provide access to its software products on fair, reasonable and non-discriminatory pricing terms, where different fees and commercial terms are charged to different customers only where objectively justified.

- (a) Impact and risk profile: factors such as Microsoft's position of significant market power, the nature of the practices, the potential for change in Microsoft's software products or supply arrangements, and connections to wider elements of Microsoft's business mean that it would be challenging to address the potential for circumvention without the ability to iteratively develop these measures.
- (b) Effects of complexity: we do not consider that any of these measures are straightforward to specify, monitor or enforce. Given their complexity, there is a specification risk and any initial design would likely require refinement, calibration, or correction, particularly over time.
- (c) Monitoring and enforcement: these measures would require a monitoring and compliance programme, which would involve establishing and maintaining a framework akin to the CMA's digital markets competition regime and resulting in potentially overlapping regulatory regimes in relation to some digital services.
- (d) Timescales: long timescales increase the likelihood of circumvention and distortion, particularly as a result of overriding market signals for an extended period. The lack of flexibility to adapt and iterate the measures over time increases this concern.

10.96 We have identified that while, in principle, it might be possible to address some of these risks in the design of a market investigation order, there are likely to be material risks associated with these remedies such that it would be difficult to achieve an effective remedy through the use of our remedy-making powers under EA02. For example, in appendix W we discuss (and dismiss) the potential for establishing a new implementation and oversight body.

10.97 We also considered whether a 'simpler' form of remedy design implemented under our EA02 powers could be effective in addressing Microsoft's licensing practices.

10.98 However, we do not consider this would be effective in addressing the AEC we have found, and we consider that it would carry even higher circumvention risks than a comprehensive package of remedies. For example, a remedy implemented using our EA02 powers which simply removed the concept of Listed Provider, and/or allowed customers to transfer their licences to any platform would not address:

- (a) the variety of mechanisms through which Microsoft is able to directly and indirectly influence the pricing of its products;
- (b) the variety of ways through which Microsoft would be able to respond, noting among other things that Microsoft is a vertically integrated provider which

could be incentivised to implement a licence transfer in a way that favours its own business, for example through direct or indirect pricing mechanisms; or

- (c) future contractual changes that could potentially be introduced which have similar effects to those currently in place.

10.99 In contrast, we consider that if the CMA were to designate Microsoft with SMS in respect of cloud services, it would have the ability to test and trial remedies, as well as to iterate remedies over time, which we consider would likely sufficiently address the risks we have identified in our assessment.

Time-limited interventions relating to Microsoft licensing practices under our order-making powers in combination with a recommendation to the CMA Board

10.100 We have considered stakeholder submissions that any recommendation to the CMA Board to prioritise commencing an SMS investigation of Microsoft's digital activities in respect of cloud services should be accompanied by time-limited remedies pending any implementation of the recommendation.

10.101 In particular, Google has said that were we to remedy Microsoft's licensing practices through a recommendation to the CMA Board, it would be necessary to also put in place certain time-limited interventions, intended to prevent the situation from worsening pending any interventions under the CMA's DMCC Act powers.²⁸³⁶

10.102 Our recommendations are a comprehensive solution to the AECs we have identified, and because these are an effective and proportionate remedy, we do not think it necessary to also consider the need for mitigating measures. However, having received these submissions from Google, we have assessed them below.

10.103 Google's proposed time-limited interventions are:

- (a) a requirement for Microsoft to reverse its planned restrictions on the use of software on Listed Providers' infrastructure by ISVs and MSPs;²⁸³⁷
- (b) a prohibition on Microsoft taking any steps to worsen the position of Listed Providers [redacted] or imposing further discriminatory measures against Listed Providers pending the imposition of any remedies under the CMA's DMCC Act powers;²⁸³⁸ and

²⁸³⁶ [redacted].

²⁸³⁷ [Google's response to the CMA's PDR](#) dated 28 January 2025, paragraph 21.

²⁸³⁸ [Google's response to the CMA's PDR](#) dated 28 January 2025, paragraph 21.

- (c) a prohibition on Microsoft locking in new customers (eg through very long term customer agreements) pending the outcome of any SMS investigation.²⁸³⁹

10.104 In response to the CMA's announcement that it will not consider further SMS designations until early 2026,²⁸⁴⁰ Google has requested a reassessment of our proposal to defer remedies to the DMU rather than to impose remedies using the powers available to us under EA02. It asked that we use powers under EA02 to reverse all licensing restrictions introduced by Microsoft since 2019 ([REDACTED]) and ensure that, at a minimum, the first requirement set out above is introduced.²⁸⁴¹

10.105 We do not consider that the three proposed measures listed above,²⁸⁴² even if implemented on an interim basis and accompanying our recommendations, would be effective, given that they are also subject to the risks set out above of using EA02 powers to address Microsoft's licensing practices.²⁸⁴³

10.106 In the context of identifying measures that remedy, mitigate or prevent the AECs we have found, we do not consider these proposals, individually or collectively to be effective (on either a standalone basis or supporting our recommendations) for the following reasons:

- (a) first, the proposed prohibitions on Microsoft worsening the position of Listed Providers or locking in customers lack the necessary precision to be effective in the absence of the ability to iterate swiftly. As discussed above, the ability to iterate and refine remedies is more limited using our market investigation powers than it is under the digital markets regime. This poses a challenge to any remedy relating to licensing practices implemented under our market investigation powers, but would be particularly challenging for a remedy intended to apply on a time-limited basis. For instance, a measure prohibiting Microsoft from worsening the position of Listed Providers or imposing discriminatory measures would be subject to similar specification concerns as those discussed above for licensing remedies more generally (such as prescribing what would constitute discrimination or a worsening of position). As a result, these measures could not, on their own, provide much certainty to parties affected by the practices.
- (b) related to the previous point, and with the exception of the proposed requirement for Microsoft to reverse its planned restriction on the use of software on Listed Providers' infrastructure (which could potentially be

²⁸³⁹ [REDACTED].

²⁸⁴⁰ [CMA takes first steps to improve competition in search services in the UK - GOV.UK](#)

²⁸⁴¹ [REDACTED].

²⁸⁴² For the purposes of our assessment of this proposal, we understand these to be proposals for final orders issued under EA02, section 161 EA02 with the intention for them to operate on a time-limited basis and with the ability to be varied or revoked should the CMA take action in relation to Microsoft's licensing practices under the DMCC Act.

²⁸⁴³ See risks set out elsewhere in this chapter and in appendix W.

implemented through targeted contractual changes), these measures primarily consist of principles-based proposals and would require substantial resources to implement and monitor. This makes them especially unsuitable as a short-term solution pending further action. They would require a framework for ongoing implementation and oversight in the context of a complex digital market, functions to which the new digital markets regime is better-suited. For example, disputes about the implementation of these measures and any alleged circumvention would require adjudication through such a framework. As a result, these measures could not, on their own, provide much certainty to parties affected by the practices.

- (c) in addition, these proposed measures are insufficiently targeted at addressing the AEC we have found in relation to Microsoft's licensing practices. This is particularly true of the proposed measures relating to ISVs and MSPs²⁸⁴⁴ and those relating to long-term customer agreements. We have not identified either of these to be currently harming competition in this investigation.

10.107 We consider these proposals further in appendix W.

Structural remedies

10.108 In principle, structural remedies could reduce barriers to entry and expansion and, by extension, reduce the resulting significant unilateral market power held by AWS and Microsoft. For example, the largest cloud providers could divest parts of or their entire cloud services operations, or these could be operationally separated.^{2845 2846}

10.109 In our working paper,²⁸⁴⁷ we explained that structural remedies and operational separation raised particular risks and practical considerations and that we were minded not to prioritise further consideration of structural or operational separation remedies. Since then, no parties to this inquiry told us that structural remedies would be an effective and proportionate remedy, or provided evidence to support such a view. Hence, we have concluded that structural remedies and operational separation would not represent an effective and proportionate remedy to address the AECs we have identified.

Utility-style regulation

10.110 One customer said that the CMA should prioritise using its DMCC Act powers, in what it termed 'utility-style regulation' in cloud services now. This could include

²⁸⁴⁴ Further details on why we would not consider this measure to be effective in particular are set out in appendix W.

²⁸⁴⁵ [Potential remedies working paper](#), section 5.

²⁸⁴⁶ [Potential remedies working paper](#), paragraphs 5.11 to 5.13.

²⁸⁴⁷ [Potential remedies working paper](#), paragraph 5.14.

direct price regulation of some IaaS services, guaranteed minimum service standards and FRAND conditions on access to services.²⁸⁴⁸

- 10.111 However, such an approach would have the effect of overriding many (if not all) market signals within the relevant market(s), and so would likely largely forego the benefits of competition, instead relying wholly on the regulatory system to address any detrimental effects. Given that we do see a degree of existing competitive pressure, we would be concerned about the distortion risks associated with some of the specific measures identified.
- 10.112 Therefore, while we recognise that this approach reflects one method of seeking to remedy significant unilateral market power (and/or its detrimental effects), we consider that these forms of measures are often used in regulated sectors where it may not be feasible to introduce effective competition.
- 10.113 We consider that the recommendations described above would represent a comprehensive remedy to the AECs. We note that, if the CMA designated AWS and Microsoft as having SMS in respect of their cloud activities, it would have flexibility in its approach to imposing regulatory obligations on these companies and so could reflect any future changes in these markets.

²⁸⁴⁸ [X] submission to the CMA [X].