CLOUD INFRASTRUCTURE SERVICES

Provisional decision report

28 January 2025

Competition & Markets Authority

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 [✂]. Some numbers have been replaced by a range.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
Entry and expansion in cloud services 13
Customers' ability to switch cloud provider and multi-cloud 14
Microsoft's software licensing practices 15
Committed spend agreements 15
Our provisional decision on competition 15
Customer detriment 16
Our proposed remedies 16
Next steps 17

Findings **18**
1. Our task 18
Introduction 18
Statutory duty 18
This market investigation reference 19
The role of public cloud infrastructure services in the UK today 20
Scope and focus of the investigation 20
Our approach to evidence gathering 22
Approach to our assessment and the structure of this document 23
Next steps 24
Consultation 24
Response hearings 25
Final decision report 25
2. Industry background 26
The demand for cloud services 27
Customers of cloud services 27
Growth in demand for cloud services 28
Providers' market positions 29
AWS 29
Microsoft 30
Google 31
IBM 32
Oracle 33
Other providers 34
Our assessment of the main providers 35
How customers purchase cloud services 35
Parameters of competition 37
Evidence from providers 37

Evidence from customers 38
Switching and multi-cloud 38
Switching process 39
Types of multi-cloud 41
Benefits and disadvantages of using multi-cloud for customers 43
Provisional conclusions 45
3. Competitive landscape 46
Market definition 47
Product market 48
Geographic market 77
Market structure and concentration 82
Framework for our assessment 82
Shares of supply by revenue 83
Shares of supply by capacity 91
Shares of supply by flows of new business 95
Provisional conclusions 101
Market outcomes 103
Introduction 103
How we use evidence on market outcomes 103
Profitability 105
Pricing trends 116
Quality and innovation 119
Prevalence of switching and multi-cloud 122
Provisional conclusions 139
The impact of AI on cloud services 141
The supply of accelerated compute by cloud providers to FM developers 150
Cloud providers' supply of access to FMs to their customers 159
Provisional conclusions 172
4. Barriers to entry and expansion 175
Sunk investment costs 176
Stakeholders' views 177
Our assessment 177
Our assessment 181
Economies of scale 183
Stakeholders' views 183
Purchasing efficiencies 184
Operating efficiencies 185
Investment in research and development 188
Our assessment 190
Size of cloud providers' product portfolio 190
Importance of range of services 191
Economies of scope 192
Network effects 194

Other potential barriers to entry and expansion 196
Cloud credits 196
Reputational barriers 201
Regulatory barriers 201
Public sector procurement 203
Provisional conclusions 205
5. Barriers to switching and multi-cloud 206
Introduction 206
Providers' incentives to facilitate switching and multi-cloud 208
Evidence from providers 208
Our assessment 211
Technical barriers 212
Introduction 212
Conceptual framework 215
Technical barriers to switching and multi-cloud 217
Core services 230
Ancillary services and tools 235
Other technical factors 251
Technical mitigations 258
Provisional conclusions 270
Egress fees 270
Introduction 270
Background 271
Conceptual framework 273
Relevance of egress fees on customers' switching and multi-cloud decisions 274
Our assessment of the relevance of egress fees to customers' switching and multi-cloud decisions 298
Potential benefits of egress fees 301
Our assessment of the potential benefits of egress fees 318
Provisional conclusions 318
Provisional conclusions on barriers to switching and multi-cloud 319
6. Licensing 321
Background 322
Description of the software licensing practices 324
Timeline of licensing practices 324
Using Microsoft software products on public cloud 327
Differences between using Microsoft software products on Azure compared to on non-Azure clouds via SPLA 330
Conceptual framework 331
Microsoft's market power in related software markets 333
Introduction 333
Framework 334
Product background 338

Market definition and market power assessments 345
Provisional conclusions on Microsoft’s market power in related software markets 376
The importance of Microsoft software inputs 377
Significance in the cost base of rival providers...378
Significance in shaping downstream competition...391
Our assessment 400
Provisional conclusions 402
Microsoft’s conduct..402
Our comparison of the price that Microsoft charges Azure customers for Windows Server and SQL Server and the price it charges AWS and Google 403
A cloud provider’s analysis of Windows Server licensing cost versus infrastructure price..413
Price and non-price differences for the client-side products...............................415
Our assessment 420
Provisional conclusions 421
The impact on rivals’ competitive offerings from Microsoft’s conduct 421
Pass-through of input prices..422
List prices 424
Margins and mark-ups 425
Relative usage 428
Customers’ choice of cloud providers 430
Customer examples from AWS and Google...437
Submissions from cloud providers 438
Our assessment 441
Provisional conclusions 443
Summary of our assessment and provisional conclusions 443
Ability 443
Incentive..444
Effect 445
Provisional conclusion...446
7. Committed spend agreements 447
Introduction...447
Conceptual framework and analysis structure..448
Stakeholder submissions...454
Cloud providers and ISVs 454
Industry bodies..455
Customers...456
Other parties 456
Our assessment 457
Prevalence of CSDs..457
Impact of CSDs on customers’ choices...460
Impact of CSDs on competition...462

Potential benefits of CSDs 476
Provisional conclusions 477
8. Provisional decision on competition 479
Introduction....479
Our view of well-functioning cloud services markets 479
Market definition 480
Market concentration 481
Impact of AI on cloud services....481
Market outcomes 482
Barriers to entry and expansion....482
Barriers to switching and multi-cloud 483
Microsoft licensing practices....484
Our provisional decision 486
Customer detriment 486
9. Our proposed remedies....488
Introduction....489
Structure and context 490
Framework for our assessment of proposed remedies 491
Proposed remedies: recommendations to the CMA Board 494
Background 494
Stakeholder views 496
Description of the remedies....496
Effectiveness....497
Proportionality 503
Other remedy options that we are minded not to progress....505
Measures we are not minded to progress under the remedy-making powers provided by the Act but recommend the CMA consider if SMS designations are made 505
Structural remedies 511
Utility-style regulation 512
Provisional decision on remedies 513

Tables

Table 2.1 : UK IaaS and PaaS revenues, 2020-2023 (£bn)....28
Table 3.1 : UK shares of supply for IaaS, 2020 – 2023 85
Table 3.2 : UK shares of supply for PaaS, 2020 – 2023 87
Table 3.3 : UK shares of supply for IaaS and PaaS combined, 2020 – 2023 88
Table 3.4 : UK and European shares of supply by data centre capacity, 2020 – 2026....92
Table 3.5 : Global shares of supply by data centre capacity, 2020 – 2026....93
Table 3.6 : UK shares of supply by year-on-year IaaS revenue growth, 2021 – 2023....97
Table 3.7 : UK shares of supply by year-on-year IaaS and PaaS revenue growth, 2021 – 2023....98
Table 3.8 : UK shares of supply by newly acquired customers, 2021 – 2023....99

Table 3.9 : UK shares of supply by revenue from newly acquired customers, 2021 – 2023 100
Table 3.10 : prevalence of multi-cloud, unweighted and weighted by spend, 2020-2023 137
Table 3.11 : Average proportion of spend on primary cloud for customers that multi-cloud, by spend band, 2020-2023 (%)...138
Table 3.12 : FM developers as customers of public cloud services. 144
Table 3.13 : Cloud providers' managed platforms..148
Table 3.14 : Non-exhaustive list of FM models available on managed platforms 149
Table 3.15 : Partnerships between cloud providers and AI developers 157
Table 6.1 : Microsoft's market share in server OS, variety of measures...........................348
Table 6.2 : Types of open-source server OSs (paid and free) 349
Table 6.3 : Windows Server input costs as a proportion of customer spend on all Azure services and Windows Server VMs on Azure, Google SPLA prices 386
Table 6.4 : 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 386
Table 6.5 : 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 387
Table 6.6 : 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 387
Table 6.7 : Average percentage difference between the wholesale prices that Google pays for Windows Server and Microsoft's customer-facing prices...406
Table 6.8 : Median percentage difference between the wholesale prices that Google pays for Windows Server and Microsoft's customer-facing prices...407
Table 6.9 : Average difference between the wholesale 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 407
Table 6.10 : Median difference between the wholesale 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 408
Table 6.11 : Average percentage difference between the wholesale prices that Google pays for SQL Server and Microsoft's customer-facing prices relating to SQL Server PAYG usage 408
Table 6.12 : Median percentage difference between the wholesale prices that Google pays for SQL Server and Microsoft's customer-facing prices relating to SQL Server PAYG usage 409
Table 6.13 : Average difference between the wholesale 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 409
Table 6.14 : Median difference between the wholesale 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..410

Table 6.15 : Average combined difference between the wholesale 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 n Windows Server VMs on Azure..411
Table 6.16 : Median combined difference between the wholesale 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 o 411
Table 6.17 : The cloud provider's analysis of differences in licensing costs paid by that cloud provider compared to customer facing price for Azure customers..........................413
Table 6.18 : Availability of the Microsoft products for use in the public cloud 415
Table 6.19 : Microsoft's submission on list prices for Windows Server VM PAYG prices across AWS, Azure and GCP 424

Figures

Figure 1.1 : Sources of evidence - Overview 23
Figure 3.1 : Prevalence of multi-cloud, weighted by spend, split by spend band, 2020-2023 138
Figure 3.2 : Illustration of the value chain for FM development and deployment 143
Figure 5.1 :The costs associated with different customer activities 214
Figure 5.2 : Types of data transfers 272
Figure 6.1 : [✂] 386
Figure 6.2 : [✂] 386
Figure 6.3 : [✂] 387
Figure 6.4 : [✂] 387
Figure 6.5 : Distribution of Azure cloud use by spend on each service............................391
Figure 6.6 : Shares of Azure VM usage by operating system 2020-2023. 395
Figure 6.7 : [✂] 406
Figure 6.8 : [✂] 407
Figure 6.9 : [✂] 407
Figure 6.10 : [✂]...408
Figure 6.11 [✂] 408
Figure 6.12 : [✂]...409
Figure 6.13 : [✂]...409
Figure 6.14 : [✂]...410
Figure 6.15 : [✂]...411
Figure 6.16 : [✂]...411
Figure 6.17 : Average Windows Server usage on Azure, AWS, and GCP in each Year .429
Figure 6.18 : Total average usage of SQL Server on Azure compared with average usage of SQL Server on AWS and GCP as licensed through the SPLAs in each Year 429
Figure 6.19 : SQL Server average usage in 2022..430
Figure 7.1 : Discount that the rival must apply to win units of contestable demand 451
Figure 7.2 : Total profit to rival for different units of contestable demand it wins.............452

Figure 7.3 : Total profit to rival for different units of contestable demand it wins for a representative ([✂]) AWS' CSD 473

Appendix

A. Demand for public cloud
B. Parameters of competition
C. Assessment of quantitative survey evidence
D. Market structure and concentration methodology
E. Profitability
F. Pricing
G. Quality and innovation
H. Prevalence of switching
I. Prevalence of multi-cloud
J. Barriers to entry and expansion [Redacted]
K. Public sector procurement
L. Egress fees - Background
M. Egress fees - Hypothetical scenarios
N. Egress fees - Free switching programmes
O. Egress fees - Customers' views
P. Egress fees - Internal documents
Q. Egress fees - Analysis of cloud providers' cost of egress
R. Licensing - Market power evidence and market shares
S. Licensing - Google's data analysis and customer examples
T. Licensing - Data analysis
U. CSA - Quantitative analysis
V. CSA - Qualitative analysis
W. Remedies appendix

Summary

Introduction

1. Cloud services are now a vital input to businesses and organisations across the UK economy, with £9 billion spent on them in 2023. This spend has been growing by over 30% a year. Cloud services support many sectors' contribution to the UK's 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.

2. Cloud services underpin UK businesses and organisations' main activities: for example, they enable banking technology, track courier deliveries and help retailers manage their stock. Healthy competition in cloud services markets can enable innovation, investment and improved productivity amongst all customers for the benefit of people, businesses and the UK economy.

3. The purpose of our investigation is to decide whether any feature or combination of features of the cloud services markets in the UK 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 'adverse effect on competition' or AEC). Should we find an AEC, we are required to decide whether we should take any remedial action or whether we should recommend the taking of action by others to remedy, mitigate or prevent the AEC(s) we have found.

4. We have provisionally found that there are AECs arising from certain features in the cloud services markets in the UK, and we are proposing remedies to address the harms to competition that we have identified. We are now consulting on these provisional findings.

The nature of competition in cloud services markets

5. Cloud services allow customers to have remote access to technology resources, on demand over a network. We define cloud services as infrastructure as a service (IaaS) and platform as a service (PaaS). IaaS includes services such as compute, networking and storage. We define IaaS based on standard compute as a separate market to IaaS based on accelerated compute.[1] PaaS includes platforms based on this infrastructure which enable customers to develop and run applications in the cloud.

6. We have provisionally found that cloud services markets in the UK are highly concentrated and each of the two largest providers, Amazon Web Services (AWS) and Microsoft, has a high share of supply, particularly in IaaS where they had

[1] Further references to IaaS in this chapter refer to IaaS based on standard compute only, unless stated otherwise

shares of [40-50]% and [30-40]% respectively in 2023. Across IaaS and PaaS markets together, they each had a share of [30-40]%.

7. The third largest provider, Google, has much lower share of supply in UK cloud services markets, and there are also other providers, including Oracle and IBM, whose share of supply is smaller and who do not supply as wide a range of cloud services as AWS, Microsoft and Google.

8. We have provisionally found that AWS and Microsoft 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, with some services and products increasing in price over time, while others are falling. Customers say that cloud services offer both quality and innovation to them. However we consider that a more competitive market would have sustained better market outcomes, including more consistently competitive prices, as well as further improvements in quality and innovation.

10. We have considered whether the growth of AI is affecting competition in cloud services as these products and services rely on cloud computing. Partnerships between larger cloud providers and FM developers are widespread and may play an important role in shaping the competitive conditions in the supply of accelerated compute to FM developers and in the supply of access to FMs to other customers. Access to FMs has emerged as a potential future driver of customers' choice and the competitive conditions in this area are not fixed.

11. This area of cloud services has been developing during the course of our inquiry and so evidence available now on the impact of AI on competition in cloud services is mixed. Overall, it is unclear if or how cloud service providers' relative strengths in the supply of IaaS based on accelerated compute will affect competition in the supply of IaaS based on standard compute. This is because the supply of accelerated compute is not currently substitutable for IaaS based on standard compute due to their different technical specifications and use cases. In that context, while we have provisionally found that AWS, Microsoft and Google each has a strong position in the supply of IaaS based on accelerated compute, we have provisionally found that there is currently no significant direct impact from this on competition in cloud services.

Entry and expansion in cloud services

12. We have provisionally found significant 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.

13. There are also economies of scale, whereby larger cloud providers have lower ongoing costs. The largest cloud providers are making very large investments 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.

14. The broad product portfolios of AWS, Microsoft 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.

15. We have considered whether procurement of cloud services by public sector customers affects competition in cloud services markets. AWS and Microsoft 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 we consider that greater competition in cloud services would create greater choice for public sector customers. We will suggest that UK government should continue to collect data on the outcomes of public procurement and drive best practice in the application of its procurement frameworks.

Customers’ ability to switch cloud provider and multi-cloud

16. We have looked at whether customers can switch cloud provider and/or use multiple clouds as their ability to exercise choice can drive competition in a market, including by lowering barriers to entry and expansion.

17. Large cloud customers are more likely than smaller ones to use multiple cloud providers, although their spending generally remains concentrated with one main provider. 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.

18. Some customers can and do successfully multi-cloud but we have found that 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.

19. We have considered whether the charging of egress fees for transferring data between cloud providers for the purposes of switching and/or multi-cloud harms competition. We have provisionally found that the presence and magnitude of egress fees reduces the ability of, and/or incentives for, customers to switch and/or multi-cloud to other cloud providers; they also reduce the incentives of suppliers to compete for their rivals’ customers.

Microsoft’s software licensing practices

20. We have investigated whether Microsoft's software licensing practices may partially foreclose its rivals in cloud services.

21. We have provisionally found that 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 unable or unwilling to switch away from these products, there are limited alternatives and Microsoft has high market shares in respect of each of these products.

22. We have also provisionally found that the Microsoft products are important inputs to cloud services, such that Microsoft has the potential to harm its rivals in cloud services when customers purchase cloud services that incorporate these products.

23. We have found differences relating to price and/or quality factors when customers use these software products on Microsoft’s cloud compared to its main rivals, AWS and Google: in fact, the price that Microsoft charges these rivals for some of these products can be higher than the retail price it charges its own customers.

24. As Microsoft has a significant market share in the concentrated markets of IaaS and PaaS, cloud customers who switch away from AWS and Google, or those that do not choose them in the first place, as a result of these licensing practices, are more likely to be captured by Microsoft.

25. We have provisionally found that Microsoft has the ability and incentive to partially foreclose AWS and Google using the relevant Microsoft software products and that its conduct is harming competition in cloud services.

Committed spend agreements

26. We investigated whether the use of committed spend agreements for customers of AWS and Microsoft harms competition in the cloud services markets.

27. We found that these agreements are widespread and can influence customers’ choices in relation to workload allocation, but we have provisionally found that rivals can profitably compete against these and so in their current form and application, they do not harm competition in cloud services markets.

Our provisional decision on competition

28. Our task is to examine whether there are any feature(s) of the UK cloud services markets that lead to an adverse effect on competition.

29. We have provisionally found that high levels of market concentration and barriers to entry and expansion have enabled each of the two largest providers, AWS and

Microsoft, to hold significant unilateral market power in these markets. This harms competition 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. This harm is exacerbated by the features we have found arising from technical and commercial barriers.

30. We have also provisionally found that there are technical barriers and commercial barriers in the form of egress fees to switching and multi-cloud that harm competition in cloud services in the UK by locking customers into their initial choice of provider which may not reflect their evolving needs and limiting 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.

31. We have provisionally found that Microsoft's licensing practices are partially foreclosing AWS and Google which is having an impact on their competitive positions, and that this harms competition in cloud services in the UK. It also exacerbates the harm we have provisionally found arising from high market concentration and barriers to entry and expansion in relation to Microsoft's significant unilateral market power.

Customer detriment

32. We consider that the AECs we have provisionally found may be expected to result in substantial customer detriment in cloud services in the UK, in terms of a material impact on customers' ability to switch, multi-cloud and exercise choice over their provider, which may ultimately be expected to impact the price and quality (including access to innovative new services) of cloud services.

33. In cloud services markets, we consider that detriment may manifest itself in terms of UK customers paying higher prices for these services than they would if the markets were more competitive. By way of illustration, if prices are on average 5% above those in well-functioning markets, this would in aggregate lead to UK customers paying around £430 million more per year for these services than they would in more competitive markets.[2] If quality or innovation were lower by the same degree, this would also have a material impact on customers.

Our proposed remedies

34. We propose making recommendations to the CMA Board to use its new digital markets powers to prioritise commencing SMS investigations to consider

[2] Calculations are as follows: 9bn - (9bn/1.05) = 9bn - 8.57bn = 429m. £9bn was the value of IaaS and PaaS UK revenue in 2023. See revenue figures in chapter 2

designating the two largest providers AWS and Microsoft with strategic market status (SMS) in relation to their respective digital activities in cloud services.

35. These powers have been specifically designed to be effective in digital markets, in recognition of the fact that some digital markets share a combination of characteristics that can cause them to ‘tip’ in favour of one, or a few firms. The new regime will allow the CMA, if it designates one or both of AWS and Microsoft with SMS, to take a targeted and iterative approach to address these concerns.

36. For features such as technical barriers, egress fees and Microsoft's licensing practices, we have provisionally found that, while in principle, there are actions we could take using the remedy-making powers available in this market investigation to address these features, there would be material risks in doing so. We consider that the new digital markets powers are better suited to addressing the concerns we have identified, particularly as a result of their greater flexibility and better provisions for ongoing monitoring and oversight. Should AWS and Microsoft be designated, we recommend that the CMA consider imposing appropriate interventions such as those identified in this report.

37. We consider that measures aimed at AWS and Microsoft would address market-wide concerns by directly benefitting the majority of UK customers and producing wider indirect effects by altering the competitive conditions for other providers.

Next steps

38. We are now consulting on these provisional findings and will consider further evidence and submissions received before reaching our final decisions later this year.

Findings

1. Our task

Introduction

1.1 On 5 October 2023, the Office of Communications (Ofcom), in exercise of its powers[3] 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).[4]

1.3 This document sets out our provisional findings of our competition assessment and our proposed remedies.

Statutory duty

1.4 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 United Kingdom or a part of the United Kingdom'.[5] If we decide that there are such features or combination of features, then there is an adverse effect on competition (AEC).[6] A 'feature' of the market refers to:

(a) the structure of the market concerned or any aspect of that structure;

(b) 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

(c) any conduct relating to the market concerned of customers of any person who supplies or acquires goods or services.[7]

1.5 If we find that there is an AEC, we are required to decide the following additional questions:

[3] Sections 131 and 133 of the Enterprise Act 2002 (EA02), as provided for by the Communications Act 2003.
[4] Details on the members of the Group are on our case page: Cloud services market investigation.
[5] Section 134(1) EA02, for present purposes, 'relevant market' means a market in the United Kingdom for goods or services of a description specified in the reference (section 134(3)(b) EA02).
[6] Section 134(2) EA02.
[7] Section 131(2) EA02.

(a) whether action should be taken by us, 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; [8]

(b) and, if so, what action should be taken and what is to be remedied, mitigated or prevented.[9]

1.6 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'[10] and we may, in particular, have regard to the effect of any action on any relevant customer benefits of the feature or features of the market(s) concerned.[11]

This market investigation reference

1.7 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.[12] Ofcom focused its study on the market for public cloud infrastructure services.

1.8 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.[13]

1.9 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.[14]

---

[8] According to 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.

[9] Section 134(4) EA02.

[10] Section 134(6) EA02.

[11] Section 134 (7) EA02.

[12] Ofcom, Cloud services market study final report, paragraph 2.4.

[13] 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.

[14] Ofcom, Terms of Reference.

The role of public cloud infrastructure services in the UK today

1.10 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.

1.11 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 digital businesses providing services to consumers would not be able to function in the way they do today.

1.12 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 than the alternatives.[15]

Scope and focus of the investigation

1.13 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.

1.14 In this market investigation we are considering 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).[16]

1.15 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.[17]

1.16 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).[18]

[15] See Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), page 25.
[16] Ofcom, Cloud services market study, final report, Terms of Reference, 5 October 2023 (Terms of Reference).
[17] 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 1.16 (c) below.
[18] In practice, this vertical stack is not strictly applied. For example, SaaS may be built and deployed using IaaS only.

(a) The IaaS layer provides access to raw computing resources (compute, storage and network) for processing workloads and storing data.[19] The hardware associated with these computing resources take 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.[20] 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. IaaS includes accelerated compute used by customers like foundation models (FM) developers.

(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.

1.17 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

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

1.18 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.

1.19 Our market investigation focuses on public cloud infrastructure services, however other cloud deployment models and traditional IT have been considered where relevant.

[19] A workload is a general term which may mean a customer's application, service, capability or other task or activity.
[20] 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.

1.20 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 dynamics in cloud services and the barriers suppliers may face when looking to enter or expanding in cloud services, as well as on assessing both technical and commercial barriers to customers switching cloud provider or using multiple clouds and on assessing potential foreclosure of cloud providers.

1.21 We have also considered the potential impact of AI on how competition works in cloud services markets and have observed increasing importance of access to foundation models (FMs) and FM-enabled services to cloud customers. Our investigation therefore included consideration of the potential impact on cloud services of the increasing demand for both the compute resources for FM development and customers' access to FMs themselves.

1.22 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 and updated issues statement. 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

1.23 Since the launch of our investigation, we have consulted a large number of parties and gathered a broad range of evidence. This involved submissions from a range of parties, in response to our issues statement,[21] working papers and updated issues statement and our numerous information requests.[22] We held over 70 calls with stakeholders, including private sector and public sector customers, intermediaries and industry bodies. We held hearings with AWS, Microsoft and Google,[23] and conducted site visits at their premises. We have also shared parts of our working paper analyses with the external legal and economic advisors of these parties in a confidentiality ring to solicit more detailed feedback on our work. We are grateful for all those who have helped us progress our work.

1.24 In addition to the evidence we gathered ourselves, we commissioned qualitative customer research from Jigsaw Research (Jigsaw).[24] This research has captured a wider range and 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

[21] Responses to our issues statement are published on our case page Cloud services market investigation.
[22] Responses to our working papers and updated issues statement are published on our case page Cloud services market investigation
[23] Summaries of these hearings are published on our case page Cloud services market investigation.
[24] 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.

report.[25] We draw on the evidence set out in this report where relevant throughout this provisional decision report.

Figure 1.1: Sources of evidence - Overview

Source: CMA

Approach to our assessment and the structure of this document

1.25 In our issues statement we set out four theories of harm based on both the structure of the market(s) we investigate 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 provisional findings on each of these theories of harm separately[26] and we have considered them in the round.

1.26 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

[25] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024).
[26] See chapters 5, 6 and 7.

customers' demand for cloud services and how customers purchase these services.

- **Chapter 3** describes the competitive landscape by setting out our approach to market definition, the structure of the cloud services markets and the concentration therein as well as the market outcomes we observe. This chapter also sets out our considerations on the impact of AI on competition in cloud services.

- **Chapter 4** provides our assessment of the barriers to entry and expansion that we have identified in cloud services.

- **Chapter 5** provides our assessment of the barriers to customers' switching between public clouds and their use of multiple clouds. In this chapter we consider how competition is impacted by technical barriers and commercial barriers, namely egress fees.

- **Chapter 6** provides our assessment whether Microsoft has partially foreclosed its rivals;

- **Chapter 7** considers the impact of committed spend agreements; and

- **Chapter 8** brings together our findings into our provisional decision on competition, followed by our proposed remedies in **chapter 9.**

Next steps

1.27 This document together with its appendices constitutes our provisional decision on any AECs and on proposed remedies. The next steps in our investigation are outlined below.

Consultation

1.28 We now invite any interested parties to submit reasons, including supporting evidence for these, in writing as to why these provisional findings and proposed remedies package should not become final (or, as the case may be, should be varied). These reasons should be received by **23.59hrs on 18 February 2025**. Any party wishing to submit reasons should do so by emailing CloudMI@cma.gov.uk.

1.29 We aim to publish submissions we receive in response to this provisional decision report. In providing their response, parties should therefore also provide us with a non-confidential version of their submission. Respondents are also asked to describe any related interests or organisation that they represent when providing a submission.

Response hearings

1.30 We will hold further formal hearings with AWS, Microsoft and Google to discuss our provisional findings and the comments received in response. Summaries of those hearings will be published on our case page. We may additionally hold calls or hearings with other parties.

Final decision report

1.31 Following consideration of the submissions receive in response to this provisional decision, the further hearings as well as any additional evidence that may be submitted, we will publish our final decision report.

1.32 Following extension of our statutory timeline,[27] we are required by law to publish our final report by 4 August 2025. Our administrative timetable is published on our case page.[28]

[27] On 19 September the inquiry group decided to extend the original reference period by 4 months under section 137(2A) of the Act to 4 August 2025 see Extension notice.
[28] Administrative timetable (updated).

2. Industry background

- 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.
- Cloud services are now a vital input to businesses and organisations across the UK economy, with £9 billion spent on them in 2023. 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 revenue in the UK has grown significantly in the recent years: over the period 2020-2023, providers' revenues in IaaS and PaaS have grown at an annual average rate of 33%. 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 revenue.
- 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 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.[29]

2.2 In order to understand the main characteristics of the reference market(s), this chapter sets out our assessment and provisional 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 customer journey, including how products are purchased;

(c) the main providers focusing on the vertically integrated suppliers of cloud services[30] (which we refer to as 'providers') including their business models and strategies;

(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.

2.3 More details, including a detailed discussion of the evidence underlying our provisional conclusions are included in Appendices A: Demand for public cloud, and B: Parameters of competition.

The demand for cloud services

Customers of cloud services

2.4 Evidence from cloud providers shows that there is a large number of customers of cloud services in the UK.[31] These are both 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 – or which have started new workloads directly on cloud.[32]

2.5 These 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).[33] Our analysis of cloud

[29] See CC3 (Revised), Guidelines for market investigations: Their role, procedures, assessment and remedies (CC3) paragraph 97.
[30] That is, vertically integrated suppliers of cloud services that operate their own cloud infrastructure, ie they own the underlying raw computing resources.
[31] Responses to the CMA's information requests [✂].
[32] 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 22 October 2024.
[33] See chapter 4.

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

2.6 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 revenue and a large number of low-spend customers are responsible for a small proportion of their revenue. 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.[34]

Growth in demand for cloud services

2.7 In order to understand better how cloud services may develop in the future, we have looked at trends in customer spending, cloud provider revenue and data centre capacity.

(a) An industry report from Gartner indicates that worldwide customer spending for cloud services will grow significantly between 2023 and 2025.[35]

(b) The 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.[36]

2.8 We have estimated UK revenue from IaaS and PaaS in the UK from 2020 to 2023.[37] Over the period 2020-2023, the IaaS and PaaS markets have grown at an annual average rate of 33%.

Table 2.1: UK IaaS and PaaS revenues, 2020-2023 (£bn)

				£bn	*%*
	2020	2021	2022	2023	Annual growth
IaaS	2.4	3.0	4.3	5.1	29.3
PaaS	1.4	2.0	3.0	3.8	39.0
IaaS and PaaS	**3.8**	**4.9**	**7.3**	**9.0**	**33.1**

Source: CMA analysis of data from IaaS and PaaS providers, IDC and Synergy

2.9 We have looked at data on data centre capacity because it may indicate providers' expectations related to growth in customers' demand: if providers forecast increases in capacity, then this is consistent with an expectation among them that demand for cloud services will continue to increase.

[34] Data refers to customers in the UK with a spend greater than $1,000 in 2022. Responses to the CMA's information requests [✂].
[35] Gartner Forecasts Worldwide Public Cloud End-User Spending to Surpass $675 Billion in 2024, accessed 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)).
[36] [✂] response to the CMA's information request [✂].
[37] CMA analysis of data from IDC, Synergy and various IaaS and PaaS suppliers. See section on market structure and concentration in Chapter 3 for more details on methodology.

2.10 Data from AWS, Microsoft, Google, IBM and Oracle shows that data centre capacity has increased significantly since 2020 and is forecasted to increase substantially in the near future, partly driven by the increase in demand for AI.

Providers' market positions

2.11 We have gathered evidence from customers, providers, industry reports and providers' internal documents on the position of the top UK providers of cloud services. Below 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 customers' views on each provider.[38] This shows a sector subject to both challenges and opportunities driven by technological developments (eg in the area of AI) and evolving customer demand.

2.12 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

2.13 AWS, a subsidiary of Amazon, started providing cloud services in 2006.[39] AWS was the first provider of cloud services. It offers a wide variety of cloud services and has a global coverage.

2.14 Industry reports consistently identify AWS as one of the leading providers of cloud services. AWS was the main provider for many large customers we contacted. Most of the other large customers who responded to our information request identified AWS as an effective or fully effective alternative to their main provider.

2.15 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

[38] More details can be found in Appendix A
[39] About AWS, accessed 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'.[40]

2.16 The customers who participated in Jigsaw's research identified the following 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.[41]

2.17 At its earnings conference for Q1 FY24, AWS said that the 'combination of companies renewing their infrastructure modernization efforts and the appeal of AWS' AI capabilities is reaccelerating AWS' [global] growth rate (now at a $100 billion annual revenue run rate)'.[42]

2.18 AWS forecasts that its global revenue for 2024 to 2026 will [✂].[43] AWS expects its capex for its cloud business [✂] and it attributes this [✂].[44]

2.19 AWS identified the following strategic priorities as 'the most important things for [its own] business to get delivered or done in 2024': [✂].[45]

2.20 Other cloud providers described AWS as being strong across customers sizes, and particularly focused on startups and digital natives (ie customers without a pre-existing on-premises IT setup).[46]

Microsoft

2.21 Microsoft started providing cloud services in 2008 through Windows Azure,[47] but made its cloud services offering more widely available in 2010.[48] It offers a large number of IaaS, PaaS and SaaS services, and has a global reach.

2.22 Microsoft is identified by analyst reports as being the second leading provider overall behind AWS. One analyst said that Microsoft is closing the gap on AWS globally and particularly in Europe.[49] Microsoft was the main provider for many

[40] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), page 28.
[41] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), page 28.
[42] AWS' earnings conference for Q1 FY24, Amazon.com announces first quarter results, first quarter ended 31 March 2024, page 2.
[43] [✂]. AWS' response to the CMA's information request [✂].
[44] AWS' response to the CMA's information request [✂]. [✂].
[45] AWS' response to the CMA's information request [✂].
[46] Responses to the CMA's information requests [✂].
[47] About Microsoft - Stories accessed 11 February 2024.
[48] Windows Azure Platform Now Generally Available in 21 Countries, accessed 11 February 2024.
[49] [✂] response to the CMA's information request [✂].

large customers we contacted. Other large customers we contacted identified Microsoft as an effective or fully effective alternative to their main provider.

2.23 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'.[50]

2.24 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.[51]

2.25 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.[52] We cover Microsoft's licensing practices in more detail in Chapter 6.

Google

2.26 Google started providing cloud services in 2008 through Google App Engine, but made its offering in cloud services more widely available in 2011.[53] Since then, it expanded to offer a large number of IaaS, PaaS and SaaS services and has a global coverage.[54]

2.27 Google is described as being a smaller cloud provider than AWS and Microsoft by analysts, although they also recognise that it is expanding its capabilities across IaaS and PaaS and that in some areas it has been influencing the rest of the industry (eg Kubernetes technology[55]).

2.28 Google was the main provider for a small number of customers we contacted. Some other customers identified Google as an effective alternative to their main

[50] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), page 28.
[51] Microsoft's responses to the CMA's information requests [✂].
[52] Responses to the CMA's information requests [✂].
[53] The History of Google Cloud Platform, accessed 11 February 2024.
[54] Cloud Computing Services, accessed 11 February 2024; Global Locations - Regions & Zones, accessed 11 February 2024.
[55] Kubernetes, also known as K8s, is an open-source system for automating deployment, scaling and management of containerized applications. It groups containers that make up an application into logical units for easy management and discovery.

provider, but most identified it as 'neither an effective nor ineffective alternative'. The most common reasons for customers not seeing Google as having an effective offering were Google not having as advanced or as broad a range of functionality / features / services as the customers' main providers (AWS or Microsoft),[56] or having no direct experience of Google's offering or a lack of internal skills to use Google's cloud services.[57]

2.29 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'.[58]

2.30 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 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.[59]

2.31 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 [✂] for Google Cloud, Google stated that [✂].[60]

(b) winning startups, public sector, FTSE 100 & SI customers; and

(c) expanding Google's partners network.[61]

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.[62] IBM offers a variety of products[63] and is active in 10 global regions.[64]

2.33 IBM is described in analyst reports as being a smaller provider when compared to both the largest, AWS and Microsoft, but also Google. IBM was not the main

---

[56] Responses to the CMA's information requests [✂].
[57] Responses to the CMA's information requests [✂].
[58] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), page 27.
[59] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), page 30.
[60] Google's response to the CMA's information request [✂].
[61] Google's responses to the CMA's information requests [✂].
[62] What is IBM Cloud? Services Offered, Features & Pricing, accessed 11 February 2024.
[63] Cloud Products - IBM, accessed 11 February 2024.
[64] Locations for resource deployment - IBM, accessed 11 February 2024.

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,[65] and a lack of any experience or knowledge of IBM's offering on the part of the customer.[66]

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'.[67]

2.35 IBM identified several competitive factors. [✂].[68]

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.[69] Oracle offers a smaller set of services than other providers and is active in 48 commercial and government regions.[70]

2.37 Analysts say that Oracle is a smaller provider when compared to both the largest, AWS and Microsoft, but also Google.[71] Oracle was not the main provider for any of the customers we contacted. Amongst other customers none identified Oracle as an effective alternative to their main provider, with most identifying it as ineffective or very ineffective alternative.

2.38 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'.[72]

2.39 Oracle also identified AI as an area of emerging focus.[73] On this point, one provider submitted that 'similar to other providers, Oracle has worked aggressively

[65] Responses to the CMA's information requests [✂].
[66] Responses to the CMA's information requests [✂].
[67] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), page 31.
[68] IBM's response to the CMA's information request [✂].
[69] Oracle Cloud Infrastructure Platform Overview, accessed 11 February 2024.
[70] Public Cloud Regions and Data Centers - Oracle United Kingdom accessed 11 February 2024.
[71] [✂] responses to the CMA's information requests [✂].
[72] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), page 32.
[73] Oracle's response to the CMA's information request [✂].

to expand its range of AI and machine learning products and services with management continuing to highlight its technical differentiation for AI workloads'.[74]

2.40 Cloud providers submitted that Oracle has fuelled its cloud business by leveraging its legacy incumbency in on-premises business, especially in data bases.[75]

Other providers

2.41 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.

2.42 There are other cloud providers who may have a large presence globally such as Alibaba, Huawei and Tencent and all of these have been identified as a competitor by at least one cloud provider.[76] However, these providers have not been identified as suitable alternatives to the main providers by UK customers we have contacted.[77]

2.43 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.[78] 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.

2.44 Examples of ISVs include VMware, MongoDB, Snowflake, Yugabyte.

2.45 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.

2.46 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 2023.[79]

---

[74] [✂] response to the CMA's information request [✂].
[75] Responses to the CMA's information requests [✂].
[76] Responses to the CMA's information requests [✂]; [✂] response to Ofcom's information request [✂].
[77] For example, Alibaba received a very low rate 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 (Responses to the CMA's information requests [✂].), 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 [✂].
[78] There are some exceptions to this – for example, Salesforce is an ISV, but operates its own infrastructure.
[79] CMA analysis of data from IDC, Synergy and various PaaS suppliers. See section on market structure and concentration in Chapter 3 for more details on methodology.

2.47 ISVs compete in specific product categories rather than across the entire range of PaaS products.

Our assessment of the main providers

2.48 Microsoft and AWS are considered the leading providers of cloud services by customers and other providers.

2.49 Analysts see Google as a smaller cloud provider than AWS and Microsoft, although recognise that it is expanding its business across IaaS and PaaS.

2.50 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 an 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.

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

How customers purchase cloud services

2.52 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.

2.53 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 portals[80] and/or their

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

marketplaces[81].[82] These customers generally pay for cloud services on a pay-as-you-go basis, paying the providers' listed prices.[83]

(b) Some customers purchase cloud services through resellers who are a network of partners authorised to resell cloud services from providers.[84]

(c) Large enterprise customers – for example, classified by one provider as those with an estimated spend of over £1 million per year[85] – generally procure cloud services through bilateral negotiations with providers.[86] 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.[87]

(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.[88] Public sector procurement is discussed in Chapter Barriers to entry and expansion. Tenders are less frequent in the private sector.

(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,[89] and another offers discounts for resellers purchasing its cloud services.[90]

2.54 Most customers have standard contracts that have been agreed without negotiation. Some larger customers either engage in bilateral negotiations or tenders and are able to negotiate terms that depart from standard contracts.

2.55 In particular, there are two main types of contracts that customers can enter when purchasing cloud services: customer agreements and enterprise agreements.

(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

[81] 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 the CMA's information request [✂]; Responses to Ofcom's information requests [✂].
[82] [✂] response to the CMA's information request [✂].
[83] [✂] response to the CMA's information request [✂].
[84] Responses to Ofcom's information requests [✂].
[85] [✂] response to the CMA's information request [✂].
[86] Responses to the CMA's information requests [✂]; [✂] response to Ofcom's information request [✂].
[87] [✂] response to the CMA's information request [✂].
[88] Responses to Ofcom's information requests [✂].
[89] [✂] response to Ofcom's information request [✂].
[90] [✂] response to Ofcom's information request [✂].

prices. Three providers said that the vast majority of customers are on this type of contract and do not negotiate additional terms.[91]

(b) Enterprise agreements, which are individually negotiated, are generally reserved for larger customers with higher spending.[92] One provider submitted that they contain commonly requested terms such as invoicing and regulatory compliance commitments.[93]

Parameters of competition

2.56 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.[94]

Evidence from providers

2.57 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.[95] Providers compete on price through discounts,[96] including discounts under committed spend agreements,[97] cloud credits,[98] and free tiers.[99]

(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,[100] security and data protection,[101] innovation,[102] availability of advanced cloud features,[103] reliability[104] of a platform,[105] flexibility to deploy cloud services in combination with their traditional IT infrastructure,[106]

[91] Responses to the CMA's information requests [✂]; [✂] response to Ofcom's information request [✂].
[92] [✂] response to the CMA's information request [✂].
[93] [✂] response to the CMA's information request [✂].
[94] 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.
[95] Responses to Ofcom's information requests [✂]; Responses to the CMA's information requests [✂].
[96] [✂] response to the CMA's information request [✂].
[97] See Chapter 7.
[98] Responses to Ofcom's information requests [✂]; [✂] response to the CMA's information request [✂].
[99] Responses to Ofcom's information requests [✂].
[100] Responses to Ofcom's information requests [✂]; Responses to the CMA's information requests [✂].
[101] Responses to the CMA's information requests [✂; [✂] response to Ofcom's information request [✂].
[102] [✂] response to Ofcom's information request [✂]; [✂] response to the CMA's information request [✂].
[103] [✂] response to the CMA's information request [✂].
[104] 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.
[105] Responses to the CMA's information requests [✂]; [✂] response to Ofcom's information request [✂].
[106] Responses to the CMA's information requests [✂].

elasticity of its services,[107] customer's loyalty,[108] brand trust,[109] and technical support.[110]

2.58 Microsoft and IBM said that different customer groups prioritise different factors when choosing a cloud provider.[111] 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.[112]

2.59 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'.[113]

Evidence from customers

2.60 We asked large customers to rate the importance of a list of factors their organisation considers when choosing their main public cloud.

(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.[114]

(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.[115]

(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.[116]

Switching and multi-cloud

2.61 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

---

107 [✂] response to the CMA's information request [✂].
108 [✂] response to Ofcom's information request [✂].
109 [✂] response to the CMA's information request [✂].
110 Responses to the CMA's information requests [✂].
111 Responses to the CMA's information requests [✂].
112 Responses to the CMA's information requests [✂].
113 Microsoft's response to the CMA's information request [✂]. [✂].
114 Responses to the CMA's information requests [✂].
115 Responses to the CMA's information requests [✂].
116 Responses to the CMA's information requests [✂].

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.

2.62 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.[117]

2.63 Both of these concepts are important as they can significantly influence the competitive dynamics in the supply of cloud services.

2.64 Customers' ability to switch or use multiple providers is key to the competitive process. In Chapter 5 (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:

(a) how the switching process works for customers

(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

2.65 In this section, we describe the steps required to switch cloud provider and the factors that customers account for when deciding to switch.

2.66 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.[118]

2.67 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.[119]

2.68 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

[117] It is more rare to split a single workload across different providers.
[118] Responses to the CMA's information requests [✂].
[119] Responses to the CMA's information requests [✂].

size of the workloads being migrated, customers' priorities and chosen deployment method.[120]

2.69 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 providers[121] and in some cases customers have said that the effort involved outweighs the benefits.[122]

(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.[123]

2.70 Some customers noted that internal 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.[124]

2.71 Customers who switched tended to do so either for performance reasons, to improve cost-efficiency, or to reduce the complexity of their cloud architectures.[125] The Jigsaw report found that, in some cases, switching was a requirement for a customer following a merger or acquisition to rationalise providers.[126]

2.72 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.[127] 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.[128]

[120] Responses to the CMA's information requests [✂].
[121] Responses to the CMA's information requests [✂]; Notes of meetings with [✂].
[122] Responses to the CMA's information requests [✂]; Note of meeting with [✂].
[123] [✂] response to the CMA's information request [✂]; [✂] response to Ofcom's information request [✂].
[124] Responses to the CMA's information requests [✂].
[125] Responses to the CMA's information requests [✂]; Note of meeting with [✂]; Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 3.6.7.
[126] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 3.6.7.
[127] Note of meeting with [✂].
[128] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 3.6.8.

Types of multi-cloud

2.73 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.

2.74 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.[129]

(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.[130] 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.[131]

(b) Integrated multi-cloud, where customers can mix and match cloud services from different public cloud providers[132] and there is a degree of integration between these services (eg data is stored on one public cloud but analysed on a different one). Application components run on different clouds.[133] 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'.[134]

(c) Siloed multi-cloud, where customers use different cloud providers for different workloads with no or minimal integration between the different clouds.[135] 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.[136]

2.75 We consider 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).

[129] Responses to the CMA's information requests [✂].
[130] Responses to the CMA's information requests [✂].
[131] Google's response to the CMA's information request [✂].
[132] [✂] response to the CMA's information request [✂].
[133] [✂] response to the CMA's information request [✂].
[134] Google's response to the CMA's information request [✂].
[135] [✂] response to the CMA's information request [✂].
[136] Google's response to the CMA's information request [✂].

2.76 We have seen 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).[137]

(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.[138]

(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.[139]

2.77 We consider the technical costs of using multiple cloud architectures in greater depth in the section on technical barriers.[140]

2.78 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.[141]

2.79 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.[142]

2.80 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:[143]

(a) Digital native customers are better positioned to adopt multi-cloud strategies.[144]

---

[137] For example, for integration of IAM services between public clouds. Responses to the CMA's information requests [✂]; [✂] response to Ofcom's information request [✂]; Note of meeting with [✂].
[138] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 3.5.5(b).
[139] [✂] response to the CMA's information request [✂].
[140] See Chapter 5.
[141] Responses to the CMA's information requests [✂]; Notes of meetings with [✂].
[142] 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 [✂]; Notes of meetings with [✂].
[143] Responses to the CMA's information requests [✂].
[144] [✂]. [✂] response to the CMA's information request [✂].

(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).[145] Google said that large companies tend to have distinct departments with their own independent workloads, facilitating the adoption of a siloed multi-cloud approach.[146]

(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.[147]

(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.[148]

Benefits and disadvantages of using multi-cloud for customers

2.81 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.

2.82 The benefits 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.[149]

(b) Resilience: multi-clouding minimises the risk of downtime – cloud duplication gives the highest degree of resiliency and downtime protection.[150] However, integrated multi-cloud also creates multiple points of failure across different clouds, increasing the risks customers are exposed to, including security ones.[151]

(c) Compliance: multi-clouding allows customers to comply with specific legal or regulatory requirements; [152]

---

[145] Microsoft's response to the CMA's information request [✂]; Responses to the CMA's information requests [✂].
[146] Google's response to the CMA's information request [✂].
[147] Responses to the CMA's information requests [✂].
[148] [✂] response to the CMA's information request [✂].
[149] [✂] response to the CMA's information request [✂].
[150] [✂] response to the CMA's information request [✂].
[151] Responses to the CMA's information requests [✂].
[152] [✂] response to the CMA's information request [✂].

(d) Geographic reach: multi-clouding can help customers achieve global coverage.[153]

(e) Flexibility: a more integrated approach gives customers the flexibility to use their preferred services from different public cloud providers.[154]

2.83 The disadvantages include:

(a) Higher costs: because of the greater complexity and increased staffing requirements, customers may incur higher costs when using multiple clouds.[155] 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.[156]

(b) Complexity: multi-clouding can increase management complexity.[157] Relatedly, integrating workloads across multiple clouds can lead to challenges around data latency and data governance.[158]

2.84 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.[159] Reasons for this included that there are challenges with adopting this approach.[160] However, other responses from these customers suggest that they do have some form of integration or are considering some form of additional integration.

(b) The responses from a few customers suggest that they may be more open to adopting an integrated multi-cloud approach in the future. 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.[161] 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.[162]

(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

[153] [✂] response to the CMA's information request [✂].
[154] Responses to the CMA's information requests [✂].
[155] [✂] response to the CMA's information request [✂].
[156] [✂] response to the CMA's information request [✂].
[157] Responses to the CMA's information requests [✂].
[158] [✂] response to the CMA's information request [✂].
[159] Responses to the CMA's information requests [✂]; Note of meeting with [✂].
[160] [✂] response to the CMA's information request [✂]; Note of meeting with [✂].
[161] [✂] response to the CMA's information request [✂].
[162] Note of meeting with [✂].

process traversing multiple cloud providers, to prevent the impact of a single provider failure from introducing increased operational resilience risk to the business process.[163] 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.[164]

Provisional conclusions

2.85 Cloud services are now a vital input to businesses and organisations across the UK economy, with £9 billion spent on them in 2023. Cloud services revenue in the UK has grown at an annual average rate of 33% over 2020-2023 and that growth is expected to continue although there is some indication that it could be at a slightly slower rate.

2.86 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 revenue.

2.87 Microsoft and AWS are considered the leading providers of cloud services by customers and other providers. They are 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.

2.88 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.

2.89 Switching and multi-cloud require customers to take a series of steps and incur costs which may affect customers' ability and/or incentive to do so. We have found that decisions on whether to multi-cloud or not as well as the level of integration to adopt are driven by how customers weigh the different benefits and disadvantages. Both customers and providers recognise this.

2.90 The specific benefits and disadvantages of switching and multi-cloud depend on the 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.

[163] Note of meeting with [✂].
[164] [✂] response to the CMA's information request [✂].

3. Competitive landscape

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

- We have provisionally found a product market for IaaS based on standard compute infrastructure which excludes 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 also provisionally found a product market for platforms based on this infrastructure (known as platform as a service or PaaS). PaaS enables customers to develop and run applications in the cloud. We refer to the IaaS and PaaS product markets together as the 'cloud services markets'

- Neither software as a service (SaaS) nor traditional 'on-premise' IT and private cloud form part of the same markets as IaaS and PaaS, because most customers do not see them as substitutes to those cloud services.

- Cloud services markets are highly concentrated. The two largest providers, AWS and Microsoft, each has a high share of supply, particularly in IaaS where, in 2023, they had [40-50]% and [30-40]% shares of supply respectively. Across IaaS and PaaS markets together, their individual shares were [30-40]% and [30-40]% respectively. The third largest provider, Google, has much lower shares of supply in cloud services markets. There are also a few smaller 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 AWS, Microsoft and Google.

- We have provisionally found that AWS and Microsoft have been generating sustained returns from their cloud services substantially above their cost of capital in cloud services for a number of years.

- While customers recognise that cloud services offer both quality and innovation to them, we consider that a more competitive market would have sustained better market outcomes. These include more consistently competitive prices, as well as improvements in quality and innovation.

- We sought to understand the potential impact of AI on how competition works in cloud services and looked at: the extent to which competitive conditions in IaaS based on accelerated compute including the partnerships formed by cloud providers and FM developers, could affect the market for IaaS based on standard compute; and the rising importance of FMs for cloud customers.

- We have provisionally found that partnerships between larger cloud providers and FM developers are widespread and that they may play an important role in shaping the competitive conditions in the supply of accelerated compute to FM developers and in the supply of access to FMs to other customers.

- Although we note that AWS, Microsoft and Google each has a strong position in the supply of IaaS based on accelerated compute, we have provisionally found that there is currently no significant direct impact from IaaS based on accelerated compute on competition in IaaS based on standard compute.
- Providing access to FMs has emerged as a potential future driver of customers' choice of cloud service provider and cloud providers have differentiated strengths in this regard, but customers do not think that the current competitive landscape in respect of the supply of access of FMs is fixed. Based on the evidence we have considered, we have provisionally found that access to FMs by customers of cloud services is not currently a strong driver of customer choice and that the extent to which it will become a driver of choice in the future is uncertain.

Market definition

3.1 Our guidelines for market investigations set out that a part of our investigation is the collection and analysis of information about the main characteristics of the market and the outcomes of the competitive process within that market. In this chapter, we set out our analysis of market definition, before then considering market structure and outcomes.

3.2 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.[165]

3.3 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.[166]

3.4 There are normally two dimensions to the definition of a market:[167]

(a) a product dimension where the relevant product market comprises a set of substitute products; and

[165] CC3 (Revised), Guidelines for market investigations: Their role, procedures, assessment and remedies, paragraphs 94 and 132.
[166] CC3 (Revised), paragraph 133.
[167] CC3 (Revised), paragraph 142 also identifies customer groups or temporal factors.

(b) a geographic dimension, where the relevant geographic market may be national (or wider), regional or local.

3.5 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.[168]

3.6 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.[169]

3.7 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)').[170] The CMA may conclude that the market definition goes wider or narrower than those goods and services.[171]

3.8 We consider the product and geographic market definition below.

Product market

3.9 The willingness and ability of customers to switch to other products is a driving force of competition. This means in forming its views on market definition, the CMA will consider the degree of demand-side substitutability for a focal product. The CMA will, where relevant, include 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.[172] In determining whether there is supply-side substitutability, the CMA may consider factors such as whether: [173]

(a) suppliers supply a range of different products in the same broad category, using the same set of assets and capabilities; or

[168] CC3 (Revised), paragraph 152.
[169] CC3 (Revised), paragraph 138.
[170] CC3 (Revised), paragraph 26.
[171] CC3 (Revised), paragraph 131.
[172] CC3 (Revised), paragraphs 134.
[173] CC3 (Revised), paragraph 134 and footnote 75.

(b) suppliers regularly introduce new products or reposition existing ones within the category.

3.10 In approaching the definition of the product market, we start from the Terms of Reference, public cloud infrastructure services in the UK.[174] As part of our analysis, we have then considered whether:

(a) IaaS should be segmented into separate, narrower markets;

(b) IaaS should be segmented into separate markets for standard infrastructure and accelerated compute infrastructure;

(c) IaaS can be widened to include PaaS:

(d) PaaS should be segmented into separate, narrower markets;

(e) PaaS can be widened to include SaaS; and

(f) Alternative IT models are substitutable for IaaS and/or PaaS.

IaaS

Whether IaaS should be segmented into separate, narrower markets

3.11 IaaS includes three broad types of services that provide access to raw computing resources for processing workloads and storing data: compute, storage and networking.[175] This section considers whether IaaS should be segmented into separate, narrower markets, one for each of these types of services. We discuss accelerated compute separately.

3.12 We set out below the cloud providers' submissions in relation to market definition for IaaS before presenting our assessment. We have not received customer evidence on the segmentation of IaaS.

Cloud providers' views

3.13 On the relationship between compute, storage and networking, AWS said that customers typically look to solve a specific IT problem, which may involve one or more different services, such as compute, storage and networking, working together in a specific way.[176]

---

[174] Terms of Reference (ofcom.org.uk), page 2.
[175] 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.
[176] AWS' response to the CMA's information request [✂].

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

3.15 The cloud providers submitted that the market is wider than IaaS, and therefore IaaS should at least not be subdivided:

(a) AWS, Microsoft 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.[177] These providers said that they do not operate their businesses according to a strict IaaS/PaaS/SaaS segmentation.[178]

(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.[179] 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.[180] 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. [181]

(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.[182]

(d) Oracle said that IaaS and PaaS should not be distinguished for the purposes of market definition.[183]

(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.[184]

[177] Responses to the CMA's information requests [✂].
[178] Responses to the CMA's information requests [✂].
[179] AWS' response to the CMA's information request [✂].
[180] AWS' response to the CMA's information request [✂].
[181] AWS response to the CMA's working papers and updated issues statement, 25 June 2024, paragraph 7.
[182] Google's response to the CMA's information request [✂].
[183] Oracle's response to the CMA's information request [✂].
[184] AWS' response to the CMA's information request [✂]; Microsoft's response to the CMA's information request [✂].

Our assessment of whether IaaS should be segmented

3.16 We consider that the different elements of IaaS, namely standard compute (accelerated compute is assessed separately), storage and networking, serve different functions. We therefore consider that it is not likely that each element of IaaS are demand-side substitutes.

3.17 On the supply side, we understand that all IaaS providers supply standard compute, storage and networking. We have not received any evidence that suggests that competitive conditions are different for each element. Accordingly, we consider that there is likely to be supply-side substitution between each of these IaaS elements.

3.18 In any case, we will take into account the extent to which there is variation in the competitive strength of providers in our wider assessment.

3.19 The features that we are investigating are relevant across each of the IaaS services, and as such consider that it is reasonable to consider IaaS types of services in aggregate for the purposes of this market investigation.

3.20 In light of the potential for supply-side substitution, and the scope for the features under investigation to be considered in aggregate across the different types of IaaS service, we consider that it is reasonable for us not to subdivide the market for IaaS into its constituent services for the purposes of our competitive assessment.

Whether IaaS should be segmented into standard infrastructure and accelerated compute infrastructure

3.21 This section considers whether IaaS based on accelerated compute infrastructure should be segmented into a separate market from IaaS based on standard compute infrastructure. In our assessment, we have considered the extent to which accelerated compute infrastructure is a substitute for standard compute infrastructure.

3.22 We have examined demand for accelerated compute infrastructure in the context of the growing development and deployment of Foundation Models (FMs). FMs are a type of AI technology trained on large amounts of data that can be adapted to a wide range of tasks and operations.[185] Due to the substantial amount of computing power required to develop FMs, FM developers generally use accelerated compute provided by cloud providers. [186]

[185] Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., ... & Liang, P. (2021). On the opportunities and risks of foundation models.
[186] CMA's initial review of FMs, Full report, page 34.

3.23 Accelerated compute infrastructure involves specialised hardware – often accelerator chips – and accompanying software. This specialised hardware and associated software can run computing operations in parallel,[187] speeding up the computation of the large number of mathematical operations which are needed by FMs. The architecture of accelerator chips is suited to processing operations in parallel, making them more efficient (at least two times faster)[188] to use than Central Processing Units (CPUs).[189] Multiple accelerator chips can be used in parallel to increase the efficiency of computations, whereas CPUs do not scale well for this type of operations.[190]

3.24 There are two types of accelerator chips commonly used for FM development and deployment: 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 operations necessary in FM training and inference.[191] The most popular GPUs in the market are Nvidia's.[192]

(b) Application-Specific Integrated Circuits (ASICs), 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.[193]

Cloud providers' views

3.25 A cloud provider submitted that CPUs could be used for 'AI/ML training, development and inference'.[194] This provider also noted that [✂] is specifically optimised to be deployable using 'cost efficient infrastructure across [✂]', although we note that this only refers to inference of a relatively small FM.[195]

---

[187] What Is Accelerated Computing? | NVIDIA Blog and CMA's initial review of FMs, Full report, page 13.
[188] CMA's initial review of FMs, Full report. page 13.
[189] The Central Processing Unit (CPU) is the primary component of a computer that acts as its 'control center.' The CPU, also referred to as the 'central' or 'main' processor, is a complex set of electronic circuitry that runs the machine's operating system and apps. The CPU interprets, processes and executes instructions, most often from the hardware and software programs running on the device. Glossary Central Processing Unit (CPU).
[190] CMA's initial review of FMs, Full report, page 13.
[191] 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.
[192] Analyzing NVIDIA's growth strategy: How the semiconductor leader is powering generative AI and the future of computing - CB Insights Research, accessed 4th November 2024, Nvidia: The chip maker that became an AI superpower - BBC News, accessed 4th November 2024; [✂] response to the CMA's information request [✂].
[193] Responses to the CMA's information requests [✂].
[194] [✂] response to the CMA's information request [✂].
[195] [✂] response to the CMA's information request [✂].

3.26 This cloud provider submitted that ‘most of [its] AI customers use GPUs’ and that to ‘compete hard for AI workloads’ it must be ‘continually improving upon [✂]’ and ‘offering ample top-of-the-range GPUs’ – in this regard, there was no mention of CPUs.[196] Generally, internal documents received from this cloud provider discussed only investment in AI accelerator chips [✂] when discussing its AI infrastructure strategy and investment.[197]

3.27 Another cloud provider said that although ‘FM developers prefer GPUs for AI model training, GPUs can be used for other purposes (eg gaming infrastructure, video editing, fluid dynamics and non-large language model AI usage) or other demanding workloads requiring high performance computers’. This provider submitted that, ‘in [its]view, the requirements of FM developers are substantially similar, if not identical, to high performance computing customers’.[198]

3.28 Two other cloud providers said that accelerated compute is an essential input in the development and deployment of FMs, suggesting that it cannot be easily substituted by standard compute:

(a) [✂] 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’.[199]

(b) [✂] 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’.[200]

Customer views

3.29 Some evidence from FM developers indicates that they predominantly rely on accelerator chips for the development of FMs, and that there is limited demand-side substitution between the latter and standard compute infrastructure.[201]

(a) One FM developer submitted that it had previously used CPUs; however, as FMs have increased in complexity, it saw a need to switch to GPUs to be able to scale (for both training and inference).[202]

(b) One chip designer said that most FMs are designed to work using GPUs.[203]

---

[196] [✂] response to the CMA’s information request [✂].
[197] [✂] response to the CMA’s information request [✂].
[198] [✂] response to the CMA’s information request [✂].
[199] [✂] response to the CMA’s information request [✂].
[200] [✂] response to the CMA’s information request [✂].
[201] Notes of meetings with [✂].
[202] Note of meeting with [✂].
[203] Note of meeting with [✂].

(c) When asked about chips used for FM training and inference, an FM developer listed only GPUs.[204]

(d) One FM developer submitted that when users 'submit a prompt [to its GenAI tool for image developing], the [AI] image generation piece must run on GPUs, which are far more expensive to buy or rent than traditional CPUs'.[205]

Our assessment

3.30 We have considered the demand and supply-side substitutability between IaaS based on accelerated compute infrastructure and IaaS based on standard compute infrastructure.

3.31 On the demand side, we consider that accelerated compute infrastructure is technically differentiated from standard compute infrastructure: standard compute infrastructure relies on general-purpose CPU processors, whereas accelerated compute infrastructure depends on accelerator chips. Accelerator chips can 'accelerate' the computational process leading to higher overall computing performance compared to using CPUs.

3.32 This leads standard compute and accelerated compute infrastructure to have different use cases. Customers typically use IaaS based on standard compute infrastructure for tasks such as running software, analysing data, managing network traffic and fetching data from memory. On the other hand, IaaS based on accelerated compute infrastructure is used for FM training and inferencing, as well as other high-performance computing tasks.

3.33 There is some evidence suggesting that IaaS based on accelerated compute infrastructure is significantly more expensive than IaaS based on 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.[206] Cloud providers also offer larger discounts for IaaS based on accelerated compute infrastructure than they do for IaaS based on standard compute.[207]

3.34 We have seen evidence of customers that use IaaS based on accelerated compute infrastructure being unwilling or unable to substitute them with services based on standard compute infrastructure. This was because of the differences in performance, efficiency and latency.[208] Additionally, researchers have estimated

[204] Note of meeting with [✂].
[205] Note of meeting with [✂].
[206] Note of meeting with [✂]; 'GPU vs. CPU Cost Analysis in Shared Hosting Environments' by Medium', accessed 25 September 2024.
[207] [✂]. [✂] response to the CMA's information request [✂].
[208] Notes of meetings with [✂]; [✂] response to the CMA's information request [✂].

that building an AI model on non-state-of-the-art accelerator chips would be at least 33 times more expensive than using state-of-the-art accelerator chips.[209]

3.35 We have also seen evidence of customers of IaaS based on standard compute infrastructure being unwilling or unable to substitute them with services based on accelerated compute infrastructure because accelerated compute infrastructure is significantly more expensive to procure, and more technically complex to deploy.[210]

3.36 In light of the evidence set out above, we consider that there is a very limited degree of demand-side substitution between IaaS based on accelerated compute infrastructure and IaaS based on standard compute infrastructure.

3.37 In assessing the supply-side substitutability of IaaS based on accelerated compute infrastructure and IaaS based on standard compute infrastructure, we have considered how quickly and easily manufacturers can switch production between these two types of compute infrastructure.

3.38 Capacity shortages in the supply of accelerated compute infrastructure have prevailed despite excess capacity in standard compute infrastructure.[211] Enduring shortages suggest that supply lines for these two products are not fungible, 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.

3.39 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. Providers' internal documents categorise investments separately from other cloud or data centre capex. Cloud providers also track both their own and competitors existing and forecast AI-related compute capacity or GPU numbers (ie IaaS based on accelerated compute) separately from IaaS based on standard compute.[212] For example:

(a) In a [✂].[213] AWS said that it had invested a total of approximately [✂].[214]

(b) In a finance review submitted to Microsoft's Board in July 2023, Microsoft's [✂] between FY23 [✂] and FY24 [✂], whereas [✂] in FY23 and [✂] in

[209] Computational Power and AI (AI Now Institute), accessed 25 September 2024; Saif M. Khan and Alexander Mann, 'AI Chips: What They Are and Why They Matter' (Center for Security and Emerging Technology, April 2020).
[210] Note of meeting with [✂].
[210] CMA's initial review of FMs, Full report.
[211] 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 20 May. Note of meeting with [✂]; Responses to the CMA's information requests [✂].
[212] For example, [✂]. [✂] response to the CMA's information request [✂].
[213] AWS' response to the CMA's information request [✂].
[214] AWS' response to the CMA's information request [✂].

FY24).[215] Microsoft said that between July 2021 and June 2024 (including forecast spend), it expects to have spent in total [✂].[216]

(c) In a [✂] [internal document] Google outlined how [✂].[217]

3.40 We have seen evidence of differences in the competitive conditions for these two products. Although most of the large cloud providers compete in both IaaS based on accelerated and standard compute, their market positions differ between them. Furthermore, there are specialised providers that are heavily focused on supplying IaaS based on accelerated compute infrastructure – [✂] and [✂] – and offer IaaS based on standard compute infrastructure only to a limited extent.[218]

3.41 Despite this, and as discussed below, we acknowledge that the competitive conditions in IaaS based on accelerated compute are evolving and that the relationship between the supply of IaaS based on accelerated compute and the supply of IaaS based on standard compute may change in the future.

3.42 On the basis of the above, we consider that, at present, there is a very limited degree of both supply and demand-side substitution between IaaS based on accelerated compute infrastructure and IaaS based on standard compute infrastructure. Therefore, we consider it appropriate to define IaaS based on standard compute as a separate market from IaaS based on accelerated compute. However, we acknowledge that this might change in the future as the competitive conditions and customer preferences evolve.

Whether IaaS can be widened to include PaaS

3.43 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.

---

[215] Microsoft's response to the CMA's information request [✂].
[216] Microsoft's response to the CMA's information request [✂]. [✂].
[217] Google's response to the CMA's information request [✂].
[218] Responses to the CMA's information requests [✂].

Cloud providers' views

3.44 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.[219]

Customer views

3.45 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 service. Most customers we spoke to said they use a mix of IaaS and PaaS for their cloud workloads.[220] In general, customers expressed their preference for either IaaS or PaaS but they may use the other in specific cases.

3.46 We asked customers about the extent to which they regard IaaS and PaaS as substitutes for each other:

(a) Most 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.[221]

(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.[222]

(ii) Similarly, another customer explained that PaaS is not a perfect alternative because it did not allow it to control its own services.[223]

(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 more constraining in terms of flexibility, propriety and specific ways of working as well as being more subject to lock-in.[224]

---

[219] 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 customer switch from IaaS to PaaS.
[220] Responses to the CMA's information requests [✂].
[221] Responses to the CMA's information requests [✂].
[222] [✂] response to the CMA's information request [✂].
[223] [✂] response to the CMA's information request [✂].
[224] [✂] response to the CMA's information request [✂].

(b) A handful of customers said that IaaS and PaaS have the potential to be substitutes for specific workloads.[225]

   (i) For example, one customer said that substitutability is largely dependent on the service in context, but overall, there is minimal substitutability.[226]

   (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.[227]

   (iii) Another customer said that PaaS is a good substitute for IaaS when there is no customisation required at the operating level.[228]

(c) A few customers said that IaaS and PaaS are easily substitutable for most workloads.[229]

(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.[230]

(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.[231]

Our assessment

3.47 Although 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, we consider that this does not mean that the underlying products and services are in fact substitutes.

3.48 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:

(a) An ISV purchasing infrastructure services from a hypothetical monopolist provider of IaaS would not be able to avoid a price increase imposed by that monopolist on the cost of infrastructure by switching to a PaaS service. This

[225] Responses to the CMA's information requests [✂].
[226] [✂] response to the CMA's information request [✂].
[227] [✂] response to the CMA's information request [✂].
[228] [✂] response to the CMA's information request [✂].
[229] Responses to the CMA's information requests [✂].
[230] [✂] response to the CMA's information request [✂].
[231] [✂] response to the CMA's information request [✂].

is because the relevant PaaS service would use the same underlying infrastructure as an input to their PaaS products.[232] While some ISVs may be relatively large customers of cloud providers and thus able to negotiate better terms from specific cloud providers than smaller ISVs, they would not be able to switch away from the price rise of a hypothetical monopolist IaaS provider. In this regard, MariaDB said that there is not a level playing field to run its product in public cloud because the cloud provider controls the resources.[233] Most ISVs we received evidence from indicated that switching to self-supply for the underlying infrastructure is not feasible due to the significant costs involved.[234]

(b) Other IaaS customers (which are not ISVs) would also face the choice of absorbing the price increase or switching to PaaS from an alternative provider like an ISV. However, as discussed, ISVs also face an increase to their input costs. As IaaS is a variable cost to ISVs, it is likely that this will be passed through to the price they charge for supplying PaaS. In addition, any profits lost by the hypothetical monopolist would be ameliorated by margins on infrastructure services supplied to the ISVs. As such, both options for IaaS customers result in increased prices and any losses by the hypothetical monopolist may be mitigated.

3.49 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.

3.50 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, we consider that evidence we have seen 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.

3.51 In light of the above, our view is that it is unlikely that there would be a sufficient degree of demand-side substitutability to warrant widening the market to include PaaS.

3.52 We also consider that supply-side substitution is unlikely to be sufficient to warrant aggregating IaaS and PaaS together. As set out in paragraph 3.9 above, to aggregate markets on the basis of supply-side substitution, we may consider factors such as whether (i) suppliers supply a range of different products in the

[232] See paragraph 2.43 on the nature of ISVs' business models.
[233] Note of meeting with MariaDB [✂].
[234] Responses to Ofcom's information requests [✂].

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 larger cloud providers offer both IaaS and PaaS, we consider it 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 on Barriers to entry and expansion). We also consider that the competitive conditions in IaaS and PaaS are different, for example see our assessment of shares of supply later in this chapter.

3.53 As such, we consider that 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. However, 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.

3.54 In addition, we note that the evidence reviewed above is relevant also for the assessment of whether IaaS is substitutable for PaaS (ie responding to the question whether PaaS should be widened to include IaaS). We consider that the evidence indicates that there is limited demand-side substitutability and customers of PaaS would not easily switch to IaaS if a hypothetical monopolist of PaaS increased price by 5-10%. We consider the competitive conditions in the supply of PaaS to be significantly different from that in the supply of IaaS, as demonstrated by the large number of ISVs that do not supply IaaS services. For these reasons, we consider PaaS should not be widened to include IaaS.

PaaS

Whether PaaS should be segmented into separate, narrower markets

3.55 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 have not received customer evidence on the segmentation of IaaS.

Cloud provider submissions

3.56 As set out above, cloud providers generally submitted that different elements of PaaS were substitutes and therefore should be in the same market.

Our assessment

3.57 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.

3.58 Although PaaS could theoretically be subdivided into more granular segments, we consider that for the purposes of this market investigation, it is not reasonable to do so. In this context, taking a rigid approach to applying demand-side substitution considerations to a large number of subcategory focal products which are both arbitrary and do not have clear boundaries may not be helpful to the overall assessment of the wider supply of cloud services. For example, on the demand-side: (i) within PaaS categories there are products that are unlikely to be substitutable with each other and (ii) there are products in different PaaS categories that may be substitutable with each other.

3.59 We considered supply-side substitution between categories of PaaS. As set out above, to aggregate markets on the basis of supply-side substitution, we may consider factors including 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 larger cloud providers supply a range of products across categories using largely the same assets and capabilities, it is unlikely that the many ISVs that operate in PaaS have the capabilities or assets to quickly and easily redirect production from their current area of focus within PaaS to another area of PaaS.[235] In considering supply-side substitutability, we also consider whether competitive conditions are similar across markets. In relation to this, ISVs are present in particular niches of PaaS to varying degrees,[236] suggesting that competitive conditions are not similar across PaaS categories.

3.60 Our guidelines state that we may treat a group of product markets together for the purposes of assessing competitive effects; for example, where a feature manifests itself in a similar way across several different markets.

3.61 Therefore, we consider it is reasonable for us not to subdivide the market into narrower PaaS markets and instead consider PaaS services in aggregate for the purposes of this market investigation.

[235] For example, Salesforce said IaaS providers could use their control of IaaS to entrench and/or grow in other markets that are not at the infrastructure layer. Note of meeting with Salesforce [✂].
[236] See above.

Whether SaaS is substitutable for PaaS

Cloud provider submissions

3.62 As set out above, cloud providers generally submitted that IaaS, PaaS and SaaS were in the same market.

3.63 In addition, 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.[237]

3.64 Google said that the SaaS layer is even more fragmented than the IaaS or PaaS layers, [238] and IBM said that SaaS is a very crowded field with many large to small companies, some of which also offer PaaS.[239]

Customer views

3.65 We asked customers about the mix of PaaS and SaaS they use. Most customers submitted that they use a mix of services.[240] Some of these customers suggested that they prefer a particular service model (ie for IaaS, PaaS or SaaS).[241] For example, one customer said that its principle is to operate as high up the stack as 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.[242]

3.66 We asked customers about the extent to which they regard PaaS-based and SaaS-based solutions to be substitutes for each other. Almost all customers said that they do not consider PaaS-based and SaaS-based solutions as substitutes.[243] 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.[244] 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

[237] Microsoft's response to the CMA's information request [✂].
[238] Google's response to the CMA's information request [✂].
[239] IBM's response to the CMA's information request [✂].
[240] Responses to the CMA's information requests [✂].
[241] Responses to the CMA's information requests [✂].
[242] [✂] response to the CMA's information request [✂].
[243] Responses to the CMA's information requests [✂].
[244] Responses to the CMA's information requests [✂].

ease of use. The customer explained that the two solutions complement each other but address different needs.[245]

(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.[246]

(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.[247]

3.67 Some customers also submitted that the layers of cloud services require different skills and experience. 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.[248]

3.68 A handful of customers indicated that SaaS may be a substitute for PaaS in certain circumstances.[249] 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.[250]

(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.[251]

Our assessment

3.69 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.

3.70 In relation to supply-side substitution, based on the evidence we have seen, we consider that competitive conditions for PaaS and SaaS are significantly different

[245] [✂] response to the CMA's information request [✂].
[246] [✂] response to the CMA's information request [✂].
[247] [✂] response to the CMA's information request [✂].
[248] [✂] response to the CMA's information request [✂].
[249] Responses to the CMA's information requests [✂].
[250] [✂] response to the CMA's information request [✂].
[251] [✂] response to the CMA's information request [✂].

and we do not consider there to be supply-side substitution between the two layers.

3.71 On this basis, we do not consider that PaaS should be widened to include SaaS.

Alternative IT models

3.72 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 or PaaS. Together, we refer to these as alternative IT models or environments.

3.73 In assessing the constraint from traditional on-premises IT and private cloud services, we have sought to understand the extent to which customers would switch to those environments following a small but significant increase in price or deterioration in quality of public cloud. 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.

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

3.75 We are undertaking this market investigation during a period of significant migration to public cloud, which may be characterised by a process of customers learning which environment is most suitable for their requirements and moving their workloads accordingly. Therefore, we have considered the evidence in light of the possibility that switching may arise from competitive diversion and/or migration trends and the associated learning process.

3.76 In our assessment of the evidence, we have been mindful of the potential for asymmetric constraints and migration effects to create patterns of shifting demand that may be not related to ongoing competition. We have 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

3.77 AWS and Google said that from the customer's perspective, on-premises and public cloud are substitutable for most use cases.[252] 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,[253] 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.[254]

3.78 Some 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 will 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.[255]

(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, which it implies indicates that both on-premises and cloud services can often meet the same customer needs. 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.[256]

(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.[257]

---

252 AWS' response to the CMA's information request [✂]. Google's response to the CMA's information request [✂].
253 AWS response to the CMA's working papers and updated issues statement, 25 June 2024 paragraph 7.
254 AWS' response to the CMA's information request [✂].
255 AWS response to the CMA's working papers and updated issues statement, 25 June 2024 paragraph 9.
256 AWS response to the CMA's working papers and updated issues statement, 25 June 2024 paragraph 10.
257 Note of meeting with Okta [✂].

(d) Google said that the key difference between the two solutions are flexibility and scalability, rather than technical differentiation. It also said that traditional IT, unlike cloud services, is costly and can be inefficient.[258]

(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.[259] It also said that traditional IT can be more resource-intensive to manage and maintain. [260] 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 their IT solutions or companies that are looking for access to the most cutting-edge computing capabilities.[261]

3.79 AWS submitted that traditional IT solutions are exerting competitive pressure on cloud services providers,[262] and customers switch regularly between cloud and traditional IT.[263] 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'.[264]

(b) AWS told us about a news article which identified a customer that said it boosted its profits from leaving the cloud.[265] 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.[266]

(c) AWS submitted survey evidence on switching. We set out this evidence, and our assessment of the evidence below.

3.80 AWS, Google and Oracle said that there are very few technical differences between on-premises and cloud services.[267] Oracle added that the difference is in method of delivery, not the underlying technology.[268]

[258] Google's response to the CMA's information request [✂].
[259] Microsoft's response to the CMA's information request [✂].
[260] Microsoft's response to the CMA's information request [✂].
[261] Microsoft's response to the CMA's information request [✂].
[262] AWS' Economic response to the CMA's Competitive landscape working paper, paragraph 5.
[263] AWS' response to the CMA's information request [✂].
[264] AWS' response to the CMA's information request [✂].
[265] Transcript of hearing with AWS [✂].
[266] Are rainy days ahead for cloud computing? - BBC News
[267] AWS' response to the CMA's information request [✂]; Google's response to the CMA's information request [✂]; Oracle's response to the CMA's information request [✂].
[268] Oracle's response to the CMA's information request [✂].

3.81 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.[269] In addition, AWS cited research and said that Dell, VMware and HPE are pivoting towards cloud-like 'as a service' operating models, which AWS said shows strong supply side substitutability.

Private cloud

3.82 AWS said that, from a customer's perspective, IT services are substitutable for most use cases, regardless of their delivery method.[270]

3.83 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.[271]

3.84 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.[272]

3.85 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.[273] It provided an example where a customer chose instead to invest in a private cloud platform.[274]

3.86 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.[275]

Our assessment

3.87 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. To the extent that customers do not view alternative IT as a good

[269] AWS' Economic response to the CMA's Competitive landscape working paper paragraph 21 and 22.
[270] AWS' response to the CMA's information request [✂].
[271] Google's response to the CMA's information request [✂].
[272] Google's response to the CMA's information request [✂]; Microsoft's response to the CMA's information request [✂]; Oracle's response to the CMA's information request [✂]; IBM's response to the CMA's information request [✂].

[273] Google's response to the CMA's information request [✂].
[274] Google's response to the CMA's information request [✂].
[275] AWS' response to the CMA's information request [✂].

substitute (because alternative IT does not meet those requirements well), we could still find that public cloud is a separate market.

3.88 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. We therefore consider it informative to consider the use cases for traditional IT, and evidence on substitutability between it and public cloud.

3.89 With respect to anecdotal examples of customers switching, we note that this shows only that such switching occurs. The question we are concerned with is whether such switching is likely in response to a small but significant increase in price or deterioration in quality of public cloud, and to such a sufficient extent that it would prevent a hypothetical monopolist of all IaaS services from raising its prices profitably.

3.90 With respect to the cloud providers' qualitative submissions that cloud and alternative IT models are substitutable, we assess these submissions alongside other evidence on substitutability which includes views from various stakeholders, responses to requests for information sent to cloud customers, and qualitative market research commissioned by the CMA.

Customer views

3.91 We asked large customers about the extent to which they consider IT environments other than the public cloud (namely private cloud and traditional IT) as 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.

3.92 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, or that doing so would involve significant resources.[276] Most of these customers said that public cloud had different characteristics, such as advanced functionality, cost, elasticity, scalability, flexibility and resiliency.[277]

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

(a) For example, one 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.[278]

(b) 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.[279]

(c) Another 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.[280]

(d) A further 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.[281]

(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’.[282]

3.93 A handful of customers said that non-public cloud IT environments was only a substitute for public cloud for certain workloads or had significant qualifications about how good of a substitute it was. For example:[283]

(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;[284]

(b) Similarly, another customer said that substitutability depends on the application stack development;[285]

278 [✂] response to the CMA’s information request [✂].
279 [✂] response to the CMA’s information request [✂].
280 [✂] response to the CMA’s information request [✂].
281 [✂] response to the CMA’s information request [✂].
282 [✂] response to the CMA’s information request [✂].
283 Responses to the CMA’s information requests [✂].
284 [✂] response to the CMA’s information request [✂].
285 [✂] response to the CMA’s information request [✂].

(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;[286] 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 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.[287]

3.94 A couple of customers said that non-public cloud IT environments are a good substitute for public cloud:[288]

(a) One customer said that most of its workloads can be relatively easily migrated to non-public cloud environments.[289]

(b) Another customer said that most environments are a viable substitute.[290]

3.95 We also asked customers whether they had ever switched any workloads hosted on a public cloud to be hosted on non-public cloud IT environments instead, such as private cloud or traditional IT, what they switched and the reasons why. This question can be expected to identify both switching due to competitive diversion, and switching as customers learn that their requirements are not met by public cloud. Where possible, we have set out customers' reasons for switching to help consider whether non-public cloud IT environments pose a strong constraint on the remaining workloads. Further, the question asked for the reasons customers switched – but not 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. We set out the evidence below:

(a) Most customers said that they had not switched to on-premises,[291] although two of these were exploring switching.[292]

(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

[286] [✂] response to the CMA's information request [✂].
[287] [✂] response to the CMA's information request [✂].
[288] Responses to the CMA's information requests [✂].
[289] [✂] response to the CMA's information request [✂].
[290] [✂] response to the CMA's information request [✂].
[291] Responses to the CMA's information requests [✂].
[292] Responses to the CMA's information requests [✂].

requirements for those workloads, rather than as a response to a price rise or degradation of quality.[293] For example:

(i) one customer switched because of application latency constraints and poor client application architecture;[294] and

(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.[295]

(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.[296] No other customers we spoke to identified moving to the private cloud.

3.96 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.[297]

(a) Of these customers, most qualified their answer, saying that there would be cost and/or effort involved in switching.[298]

(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.[299]

(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.[300]

(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.[301]

---

[293] Responses to the CMA's information requests [✂]. [✂].
[294] [✂] response to the CMA's information requests [✂].
[295] [✂] response to the CMA's information requests [✂].
[296] [✂] response to the CMA's information requests [✂].
[297] Responses to the CMA's information requests [✂].
[298] Responses to the CMA's information requests [✂].
[299] Responses to the CMA's information requests [✂].
[300] [✂] response to the CMA's information requests [✂].
[301] [✂] response to the CMA's information requests [✂].

3.97 A supplier of professional services also commented on customer needs.[302] 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 that customers might choose public cloud for other reasons such as technical strategy or commitments with cloud providers.[303]

Jigsaw report

3.98 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, the customers spoken to said that public cloud was reliable and came with less risk to continuity than on-premises or private cloud.[304] In addition, the customers spoken to said 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. [305]

3.99 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 consider GPUs in the section on accelerated compute above) and data transfer (we consider the costs of data transfer in the section on egress fees in the Barriers to switching and multi-cloud chapter). [306]

Evidence from surveys

3.100 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.[307] 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

[302] Suppliers of professional services assist customers with their cloud strategy (see, for example, Cloud Offerings - KPMG UK), which may include negotiating with cloud providers on customers behalf.
[303] Note of meeting with [✂].
[304] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), page 25.
[305] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), page 25.
[306] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), page 26.
[307] AWS' economic response to the Competitive landscape working paper, paragraph 5.

organisations that switched workstreams away from 'hyperscalers', 65% switched to their own on-prem services.[308]

(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 private cloud customers).[309] Overall, 7.5% of IaaS/PaaS users that responded submitted that they switched to an on-premises solution.

(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%.[310] 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.

3.101 We set out our observations on this evidence below:

(a) We place limited evidential weight on the surveys that AWS submitted on switching to traditional IT. We consider that there are significant weaknesses to using quantitative surveys in this market for our purposes due to lack of validity and vulnerability to the quality and coverage of the customers spoken to.[311] In fact, these weaknesses are recognised by the research that AWS submitted:

(i) In relation to lack of validity, the research recognises that the repatriation terminology is 'imprecise' and 'somewhat of a loaded term'.[312] We consider that such ambiguity makes it difficult to interpret the results in a meaningful way.

(ii) The research indicated that the results are sensitive to the survey used, stating that analysing the results of 'a different survey gives us a different story'.[313]

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

[308] AWS' economic response to the Competitive landscape working paper, paragraph 13.
[309] AWS' submission to the CMA [✂].
[310] AWS' economic response to the Competitive landscape working paper, paragraph 24.
[311] As set out in Appendix C.
[312] Cloud repatriation: What it is, what it isn't, and why it's not going away, pages 1 and 3.
[313] Cloud repatriation: What it is, what it isn't, and why it's not going away, page 3.

(i) The statistics calculated are more akin to a ‘switching ratio’ than a diversion ratio because they capture all customers switching, irrespective of changes in price or quality. As we set out above, switching may arise from competitive diversion and/or migration trends and the associated learning process. As such, we need to consider the reasons for switching and the extent to which customers view traditional IT as a good alternative to public cloud. While we consider it appropriate to place limited weight on this evidence, we note that the research itself 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 of the customers spoken to 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 represent competitive diversion. In any case, above we have set out the evidence we have gathered on the extent to which customers consider traditional IT a good substitute starting at paragraph 3.92.

(ii) Switching of workloads from public cloud to traditional IT is not definitive evidence of a constraint from traditional IT. Customers may choose to deploy workloads on public cloud to test the suitability of public cloud for certain applications, or to develop an understanding of which applications are best deployed on public cloud versus on traditional IT (as discussed above). In such cases, public cloud and traditional IT are complements, rather than substitutes. Without sufficient clarity on the motivation for customers’ switching from public cloud to traditional IT, survey evidence indicating that customers have switched a proportion of workloads from public cloud to traditional IT is ambiguous. 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%.

(iii) The switching ratios overstate the strength of traditional IT:

(1) The research identified in the paragraph 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. 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 a customer that switched all its workloads to another cloud provider. These customers would be equally weighted in this survey. In addition, 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 using this question.

(2) The research identified in paragraph 3.100 considers whether customers moved any workloads to traditional IT. Similar to the previous survey, customers that switched one or a limited number of workloads to traditional IT would be included in this proportion, which we do not consider to be reflective of the constraint that traditional IT poses on public cloud customers.

Our assessment of alternative IT models

3.102 Most customers do not consider alternative IT models like private clouds and traditional IT to be a good substitute for public cloud. 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.

3.103 While there were a few instances where customers identified that alternative IT models are a good substitute for public cloud, they are few in number and, where they arise, they 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.

3.104 The evidence we have seen indicates that public cloud workloads could technically be hosted on private cloud. However, evidence from customers shows that there would be significant costs and time associated with switching to private cloud and only two customers indicated that they had previously switched.

3.105 The Jigsaw report also found that customers consider that their requirements are best matched by public cloud, and the characteristics of public cloud are not matched well by alternative IT models.

3.106 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. We consider it appropriate to attach weight to the customer evidence we have gathered using random sampling and for which we have gathered information on the context and background to the decisions being made.

3.107 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'.[314] This indicates that there is a learning process associated with migration to the public cloud (as we set out above). 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. 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.

3.108 AWS also submitted several surveys, which it said show traditional on-premises IT should be included in the market. However, we do not consider it appropriate to attach 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, and the survey suggests 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.

3.109 With respect to the submission that people use traditional on-premises IT and private cloud in hybrid solutions with public cloud, we consider that this provides little evidence on substitutability. In fact, customers may derive distinct benefits from each alternative IT model, rather than viewing them as interchangeable. In this sense, the alternative IT models may be better considered as complementary than substitutable. Indeed, all customers that said that they would not substitute their public cloud workloads to non-public cloud environments also had hybrid environments.

3.110 Based on the evidence set out above, we consider that traditional on-premises IT is not 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. Alternative IT models do not meet those requirements well. We recognise that alternative IT models pose some level of constraint on public cloud and some customers may be

[314] 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

able to switch their workloads to such environments. However, we do not consider that such switching would be sufficient to constrain a hypothetical monopolist from raising prices by a small but significant amount. As such, we consider that:

(a) IaaS and PaaS should be treated as separate from traditional on-premises IT for the purposes of this investigation; and

(b) IaaS and PaaS should be treated as separate from private cloud for the purposes of this investigation.

Provisional conclusions

3.111 Our provisional conclusions on the product market definition are that:

(a) we have defined a product market for IaaS based on standard compute which excludes IaaS based on accelerated compute infrastructure. For the purposes of this investigation, 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) we have defined the 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 consider it appropriate to consider PaaS services in aggregate for the purposes of this market investigation and therefore have not considered whether the market should be subdivided further into PaaS types of services;

(e) the market for PaaS should not be widened to include SaaS;

(f) IaaS and PaaS should be treated as separate from alternative IT models, including traditional on-premises IT and private cloud for the purposes of this investigation; and

(g) for the purpose of this investigation, we refer to the IaaS and PaaS markets together as the 'cloud services markets'.

Geographic market

3.112 Our guidelines state that geographic markets can be defined based on the location of either suppliers or customers,[315] by considering the degree of substitutability, ie

[315] CC3 (Revised),paragraph 145.

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.[316]

3.113 The Terms of Reference in this case state the supply of public cloud infrastructure services in the UK.[317] As noted above, 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.[318]

Geographic market for IaaS and PaaS

3.114 In this section we consider 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 consider them together. The geographic scope of cloud infrastructure services based on accelerated compute infrastructure is considered separately.

Cloud Providers' views

3.115 Cloud providers said that it is not necessarily a requirement to have physical infrastructure based in the UK to compete effectively.[319] Microsoft added that a cloud provider can rely on physical infrastructure outside but nearby the UK and still compete effectively[320] and AWS said that many of its UK customers use infrastructure located outside of the UK.[321] This is consistent with the capacity shares relative to the revenue shares, which are both set out later in this chapter.

3.116 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.[322]

(b) Microsoft and IBM said that UK-based infrastructure may be important for certain customers such as those with specialised security needs.[323]

[316] CC3 (Revised), paragraph 147.
[317] Terms of Reference, page 2.
[318] CC3 (Revised), paragraph 26 and 131.
[319] Responses to the CMA's information requests [✂].
[320] Microsoft's response to the CMA's information request [✂].
[321] AWS' response to the CMA's information request [✂].
[322] Oracle's response to the CMA's information request [✂].
[323] Microsoft's response to the CMA's information request [✂]; IBM's response to the CMA's information request [✂].

(c) AWS said that, while infrastructure does not need to be located near customers, providers may choose to do this to decrease latency or to address customer preferences for data location.[324]

3.117 AWS and Google said that significant UK-specific investments are not necessary to competitively provide cloud infrastructure services to UK customers.[325] However, Microsoft, Oracle and [✂] 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 networking equipment;[326]

(b) Oracle said that costs would include sales, marketing (though noting marketing could rely on corporate resources) and access to data centre infrastructure;[327]

(c) [✂] said that there is a multitude of UK-specific investments required to compete effectively, and that it has invested in [✂].[328]

3.118 While three providers said that, where appropriate, they tailor certain aspects on narrower geographical basis, most cloud providers set their pricing, advertising and marketing strategies globally.[329]

3.119 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. [330] Our evidence from cloud providers and customers, however, does not totally align with this statement.

Evidence from customers

3.120 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

[324] AWS' response to the CMA's information request [✂].
[325] AWS' response to the CMA's information request [✂]; Google's response to the CMA's information request [✂].
[326] Microsoft's response to the CMA's information request [✂].
[327] Oracle's response to the CMA's information request [✂].
[328] [✂] response to the CMA's information request [✂].
[329] Responses to the CMA's information requests [✂].
[330] 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

these factors from one to five, with one being not important at all, and five being very important.

3.121 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 attaining an average rating that round four out of five.

3.122 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.[331]

(b) A few said that a European presence was important,[332] 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.[333]

(c) A few said that they require data centres in the UK.[334] One said that this was due to GDPR implications.[335]

3.123 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.[336]

(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.[337]

3.124 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 not listed by the CMA.

3.125 Alibaba is heavily concentrated in China,[338] and it was not identified as an effective alternative to UK customers' main providers by customers we spoke to. In particular, Alibaba received a low average rating and was not identified by any respondent as a suitable alternative.

[331] Responses to the CMA's information requests [✂].
[332] Responses to the CMA's information requests [✂].
[333] [✂] response to the CMA's information request [✂].
[334] Responses to the CMA's information requests [✂].
[335] [✂] response to the CMA's information request [✂].
[336] Responses to the CMA's information requests [✂].
[337] [✂] response to the CMA's information request [✂].
[338] Alibaba Cloud Global Locations - Deploy Around the World Including Mainland China

3.126 Reasons that customers gave for their response included a lack of experience or knowledge of Alibaba's offering,[339] that using Alibaba would be a supply chain risk or not appropriate (eg for data sovereignty or security reasons),[340] that Alibaba has more limited services or capabilities compared to larger providers,[341] and that they are only considered for demand relating to China.[342]

Our assessment

3.127 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.

3.128 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 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.

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

(a) certain customers require UK, or sometimes EEA, data centres for regulatory or security purposes;

(b) customers may prefer data centres that are located relatively close to reduce latency;

(c) customers named the number and location of data centres and data sovereignty requirements as important factors in selecting a cloud provider; and

(d) Alibaba, a Chinese provider, was universally seen as an ineffective alternative by customers.

339 Responses to the CMA's information requests [✂].
340 Responses to the CMA's information requests [✂].
341 Responses to the CMA's information requests [✂].
342 Responses to the CMA's information requests [✂].

Provisional conclusion

3.130 Our provisional conclusion is that the geographic scope of the markets for both IaaS and PaaS is Europe-wide (ie UK and EEA).

Market structure and concentration

3.131 In this section we consider 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

3.132 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.[343] 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.[344]

(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 incentive to refuse access to the input or to raise its prices, and consequently increase the costs of competing downstream firms.[345]

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

3.133 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.

[343] CC3 (Revised), paragraph 100 and 101.
[344] CC3 (Revised), paragraph 187.
[345] CC3 (Revised), paragraph 268-270.

3.134 We present shares of supply for IaaS and PaaS separately, which aligns with our provisional views on the relevant product markets (see above). We also present shares of supply for IaaS and PaaS combined (cloud services).

3.135 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 large 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 large 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.

3.136 Our calculation of market shares in IaaS and PaaS required us to categorise different cloud services between IaaS, PaaS and SaaS.[346] Categorising cloud services in this way is not always simple: some services will be on the boundary between two categories and categorisation required some judgement.[347] Therefore, while we have presented market shares in IaaS and market shares in PaaS individually, we have also presented shares in IaaS and PaaS combined. These combined shares are not sensitive to the allocation of revenues between IaaS and PaaS. However, they are also less reflective of the two separate product markets for IaaS and PaaS. We attach greater weight to the shares of supply in IaaS and in PaaS separately.

3.137 We have largely gathered data on a UK basis, which is informative of providers' success in competing for UK customers and therefore useful in assessing their strength in the market. As set out above, we consider the geographic market for these markets to be Europe-wide, therefore, where available, we present some share measures considering a wider geographic area than the UK.

Shares of supply by revenue

3.138 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.[348] High and stable or increasing shares of supply can be a strong indicator of market power, although they should be considered alongside

[346] See detail in Appendix D.
[347] 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.
[348] CC3 (Revised), Annex A, paragraph 2.

other indicators, such as high barriers to entry and expansion, high profitability and high barriers to switching.[349]

Our approach

3.139 The analysis set out here 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.

3.140 We have calculated shares of supply by revenue using the annual UK revenue[350] data from AWS, Microsoft, Google, Oracle, IBM, UpCloud, Civo, [✄], Hyve, Wasabi and Centerprise.[351] We manually categorised individual services for some cloud providers using the methodology set out in Appendix D: Market structure and concentration methodology. Some of these 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 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.

3.141 For the total revenue across all providers (ie the denominator for our shares) we used data from IDC and Synergy and adjusted these using the revenue data provided by the 11 IaaS (plus PaaS) providers (referred to as first party revenue). We placed equal weight on these data sources. Based on the firms present in the data sets, we considered that the definitions of IaaS and PaaS used by IDC and Synergy were not in line with our definitions. We used revenue submitted by cloud suppliers when available, as we considered this would be the best estimate of supplier's UK revenue. We used data from IDC and Synergy to supplement this revenue, combining the two data sets using a different methodology for IaaS and for PaaS, which are set out in Appendix D.[352]

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

(a) The IDC data set includes a relatively sizeable aggregated 'other' segment, over which we do not have visibility over the revenue breakdown by provider. However, this revenue is included as part of the market size as reported. Based on the firms included in the data set, we consider that this aggregated

---

[349] We consider barriers to switching as part of prevalence of switching and multi-cloud in Market outcomes below.
[350] For the purposes of this analysis, we defined UK revenues as revenues generated from UK customers in the UK and we defined UK customers as cloud services customers that are operating or trading in the UK. We defined annual revenues as revenues generated within a calendar year.
[351] In the following analysis we refer to UpCloud, Civo, [✄], Hyve, Wasabi and Centerprise as small UK IaaS providers.
[352] For the denominator of the market shares calculations in the Competitive landscape working paper we did not have access to the IDC and Synergy data sources, instead we had rounded values from Ofcom calculations. The figures here reflect our methodology which is outlined in Appendix D based on IDC and Synergy data to calculate the total revenue across all providers, and therefore, historical figures are updated from those published in the working paper.

'other' category may include firms providing a broader range of services that would not be captured by our definition of IaaS.[353]

(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 providers in the calculation even if they provide only one of compute, storage or networking. However, we consider that these providers are likely to be less close competitors to the large cloud providers than would be cloud providers that offer a comprehensive offering consisting of all three services.

UK shares of supply for IaaS

3.143 UK shares of supply for IaaS for 2020 – 2023 are presented in the table below. These shares include some revenues from IaaS based on accelerated compute which could not always be distinguished in the Parties' data, but we understand the contribution of accelerated compute to the total revenues to be small.

Table 3.1: UK shares of supply for IaaS, 2020 – 2023

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

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 3.140).

3.144 The evidence shows that the IaaS market is highly concentrated. Between 2020 and 2022 our analysis based on shares by revenue shows that concentration in IaaS was increasing. The revenue shares from 2023 show a stabilisation in the level of concentration.

(a) AWS' and Microsoft's combined share of supply for IaaS was increasing between 2020 and 2022, and remained high in 2023 at [70-90]%. 2023 figures illustrate that this combined share is stable.

[353] 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 from both Synergy and IDC data sets. Therefore, we consider that the 'other' category may also contain revenue from providers which do not align with our definition of IaaS.

(b) The next largest provider is Google, but in 2023 AWS remained at [least four] times its size. Google is growing and gaining share, but despite that the combined market share of Microsoft and AWS has grown since 2020.

3.145 Our analysis also shows that:

(a) AWS is the largest provider of IaaS, and between 2020 and 2022, its share remained largely the same [✄]. From 2022 to 2023 its share decreased slightly from 2022 [40-50]% to 2023 [40-50]%;

(b) Microsoft is the second largest provider of IaaS. Its share increased from [30-40]% in 2020 to [30-40]% in 2022. Its share has decreased slightly from [30-40]% in 2022 to [30-40]% in 2023;

(c) Google is the third largest IaaS provider and has a significantly lower share of IaaS revenues than AWS and Microsoft: it held just a [5-10]% share in 2023 which was higher than its share of [5-10]% in 2022, and [5-10]% in 2020.

(d) For IBM and Oracle, shares have remained in the 0-5% range from 2020 to 2023.[354]

3.146 [✄], the largest of the small UK IaaS providers for which we have collected revenue, had a declining share of supply in UK IaaS between 2021 and 2023, and held [0-5]% share in 2023. The combined share of all other IaaS suppliers has decreased from [10-20]% in 2020 to [5-10]% in 2023.

3.147 Our analysis shows that concentration in IaaS is stable in the sense that the combined share of AWS and Microsoft has remained persistently high from 2021 – 2023 at around [70-90]%. Over this period, AWS' share has decreased slightly while Microsoft's has increased slightly.

3.148 While the IaaS market is growing overall and this may provide opportunities for smaller providers as there is still new business to be won, the small (0-5%) shares of other providers and lack of growth in these shares over the period 2020 to 2023, indicates that these providers are likely to remain considerably smaller than their rivals. Google, with a share increasing on average by [0-5] percentage points from 2020 to 2023 also remains small in comparison to AWS and Microsoft.

UK shares of supply for PaaS

3.149 UK shares of supply for PaaS for 2020 – 2023 are presented in the table below.

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

Table 3.2: UK shares of supply for PaaS, 2020 – 2023

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

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 3.140).

3.150 The PaaS market is less concentrated than IaaS:

(a) AWS and Microsoft account for a smaller share of PaaS than IaaS at [40-60]% combined, while Google accounts for a larger share.

(b) Google, nevertheless, has a much smaller share in PaaS than AWS or Microsoft.

(c) The combined share of all other smaller cloud providers and ISVs in PaaS has increased slightly in 2023.

3.151 Our analysis also shows that:

(a) AWS is the largest provider of PaaS. Its share increased from [20-30]% in 2020 to [20-30]% in 2022. In 2023, its share has decreased slightly to [20-30]%;

(b) Microsoft is the second largest provider of PaaS and its share has increased from [20-30]% in 2020 to [20-30]% in 2022. In 2023, its share decreased slightly to [20-30]%;

(c) Google is larger in PaaS than it is in IaaS and its share in PaaS has also been growing: its share has increased from [5-10]% in 2020 to [5-10]% in 2023;

(d) For IBM and Oracle, shares, have remained in the 0-5% range from 2020 to 2023.[355]

3.152 The combined share of the other smaller cloud providers and ISVs has declined slightly from [40-50]% in 2020 to [30-40]% in 2023 but remains much higher than

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

in IaaS. Shares in PaaS do not capture the distribution of infrastructure (IaaS) which ISVs use.[356]

3.153 Our analysis also shows that concentration in PaaS is stable with the combined share of AWS and Microsoft remaining between [40-60]% and [40-60]% between 2020 and 2023.

UK shares of supply for IaaS and PaaS combined

3.154 UK shares of supply for IaaS and PaaS combined for 2020 – 2023 are presented in the table below.[357]

Table 3.3: UK shares of supply for IaaS and PaaS combined, 2020 – 2023

				%
	2020	*2021*	*2022*	*2023*
AWS	[30-40]%	[30-40]%	[30-40]%	[30-40]%
Microsoft	[20-30]%	[20-30]%	[30-40]%	[30-40]%
Google	[5-10]%	[5-10]%	[5-10]%	[5-10]%
IBM	[0-5]%	[0-5]%	[0-5]%	[0-5]%
Oracle	[0-5]%	[0-5]%	[0-5]%	[0-5]%
Other	[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 3.140).

3.155 The overall cloud services sector (IaaS and PaaS combined) is concentrated, and concentration is stable over time. Google is much smaller than AWS or Microsoft, however it is growing, while AWS' and Microsoft's shares have declined slightly in 2023 – although Microsoft's market shares has grown over the last four years. The combined shares of all other smaller cloud providers and ISVs have increased marginally between 2022 and 2023 but declined over the last four years.

3.156 Our analysis shows the shares of supply in cloud services by revenue – that is IaaS and PaaS in combination. Our analysis shows that:

(a) AWS is the largest provider of IaaS and PaaS combined and its share has remained broadly stable, albeit declining slightly from 2020-2023: its share was [30-40]% in 2020, decreasing to [30-40]% in 2022 and [30-40]% in 2023;

[356] 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.
[357] Similar to shares of supply for IaaS, these shares include some revenues from IaaS based on accelerated compute which could not always be distinguished in the Parties' data, but we understand the contribution of accelerated compute to the total revenues to be small.

(b) Microsoft is the second largest provider of IaaS and PaaS combined and its share has remained broadly stable, increasing from [20-30]% in 2020 to [30-40]% in 2022, albeit declining slightly to [30-40]% in 2023;

(c) Google is the third largest provider of IaaS and PaaS combined and has a significantly lower share than AWS and Microsoft: Google is significantly smaller than AWS and Microsoft, and in 2023 has reported less than a quarter of AWS' revenue. It has increased its share from [5-10]% in 2020, to [5-10]% in 2022 and to [5-10]% in 2023; and

(d) for each of IBM and Oracle, shares have remained in the 0-5% range from 2020 to 2023.[358]

3.157 The 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]% between 2020 and 2023.

3.158 The combined share of the other smaller cloud providers and ISVs has remained largely stable, between [20-30]% and [20-30] % between 2020 and 2023.

Providers' and stakeholder views

3.159 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.[359] It 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.[360]

3.160 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.[361] We note we conducted the shares of supply analysis in line with the scope of the market investigation and our market definition.[362]

3.161 AWS reported UK concentration index analysis (Herfindahl–Hirschman index (HHI))[363], which it said showed concentration had decreased in segments

[358] CMA analysis of first-party revenue data and IDC and Synergy data. Responses to the CMA's information requests [✂].
[359] Google's submission to the CMA Google Cloud's response to the CMA's Competitive landscape working paper dated 23 May 2024, paragraph 8.
[360] Google's submission to the CMA Google Cloud's response to the CMA's Competitive landscape working paper dated 23 May 2024, paragraph 8a.
[361] AWS' submission to the CMA AWS' economic response to the CMA's Competitive landscape working paper dated 02 August 2024, paragraph 28.
[362] See Market definition section above.
[363] 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.

considered by the CMA's Competitive landscape working paper, ie IaaS, PaaS and IaaS and PaaS combined, between 2019 and 2023.[364]

3.162 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.[365] We note that the HHI for cloud services provided by AWS indicates it is likely to be at least concentrated.[366]

3.163 Our calculations of UK HHIs gives 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.[367]

3.164 The decline shown in the HHI calculated by AWS was only for 2023 and should not necessarily be interpreted as a trend, particularly in light of five previous years of overall stability.

3.165 AWS' reports HHI for cloud infrastructure services, as defined in the Competitive landscape working paper (IaaS and PaaS combined).[368] We consider that HHI calculated across IaaS and PaaS combined is less useful than HHIs for IaaS and PaaS separately, as a combined measure does not accurately reflect the boundaries of the relevant markets. Appendix D describes in more detail some reasons for the differences in the CMA's and AWS' HHI estimates. In particular, we consider the CMA's estimate of HHI for IaaS to have the fewest limitations therefore place most weight on this measure.

3.166 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 the CMA and AWS calculations. Therefore, as we present above, we consider shares of supply to be our primary indicator of market concentration.

3.167 AWS also cited data showing that AWS' share in IaaS, PaaS and IaaS with PaaS has decreased from 2019 to 2023.[369] 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

[364] AWS' submission to the CMA AWS' economic response to the CMA's Competitive landscape working paper dated 02 August 2024, paragraph 31.
[365] CC3 (Revised), Annex A, paragraph 7.
[366] AWS' submission to the CMA [✂].
[367] CMA analysis of first-party revenue data and IDC and Synergy data. Refer Appendix D.
[368] The Competitive landscape working paper, paragraph 1.3.
[369] AWS' submission to the CMA AWS' economic response to the CMA's Competitive landscape working paper dated 02 August 2024, paragraph 32.

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 two sources of data, in particular from IDC and from Synergy, which we refined based on the definitions of IaaS and PaaS we have used for the purposes of our assessment, as well as first party revenue data.[370]

3.168 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.[371] 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.[372]

3.169 An industry report states that cloud services are highly consolidated, with just two primary vendors.[373]

Shares of supply by capacity

3.170 Shares by capacity look at the market structure from the supply side. At a basic level they show us the relative strength of each provider in terms of their production capability.[374] 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 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.

3.171 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

---

[370] See Appendix D.
[371] [✂] submission to the CMA [✂].
[372] 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).
[373] [✂] response to the CMA's information request [✂].
[374] CC3 (Revised), Annex A, paragraph 2.

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.

3.172 We present shares for the calendar years 2020-2026. The shares for 2020-2023 are based on realised capacity while the 2024-2026 shares are based on the 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

3.173 We calculated the shares of supply by capacity using data from AWS, Microsoft, Google, IBM, Oracle and some smaller IaaS providers that serve UK customers,[375] on their data centre capacity in megawatts (MW) within UK+EEA,[376] globally and in the UK. These shares therefore do not include the capacity of other small IaaS providers and as such each provider's share is likely a small overestimate across all providers and should be interpreted as an indicator of relative share between the cloud providers included in the calculation.

3.174 We consider that European shares give the best reflection of the market structure with respect to UK customers. These shares include data on IaaS based on accelerated compute which could not always be distinguished in the Parties' data, but we understand the contribution of accelerated compute to the total capacity to be relatively small.

UK and European shares of supply by data centre capacity

3.175 UK and European shares of supply by data centre capacity for 2020–2026 are presented in the table below.

Table 3.4: UK and European shares of supply by data centre capacity, 2020 – 2026

%

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

375 Centerprise, CoreWeave, [✂], Hyve and Wasabi.
376 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.

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. Data from 2024 – 2026 are forecasts. 'Other' category includes Centerprise, CoreWeave, [✂], Hyve and Wasabi.

3.176 UK and 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 providers and Microsoft's share has been increasing relative to AWS' share. In contrast to revenue shares, Microsoft has been the largest provider in the UK and EEA based on capacity since 2021. Smaller cloud providers have a small share of capacity, which is not forecast to change substantially. Our analysis shows that:

(a) Microsoft has become and is forecast to remain the largest provider by capacity in Europe: Microsoft's share was [40-50]% in 2020 and has increased to [40-50]% in 2023. It is the only one of the main cloud providers whose share based on forecasts will increase in future, to [✂]% by 2026.

(b) AWS is the second largest provider by capacity in Europe: AWS' share was [40-50]% in 2020 and has decreased to [30-40]% in 2023. It is forecast to decrease further to by 2026.

(c) Google is the third largest 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 2023. It is forecast to decrease slightly to [✂]% by 2026.

(d) Smaller cloud providers do not have a similar scale. Oracle and IBM's share of capacity has been [0-5]% each.

(e) Microsoft and AWS' combined UK and EEA shares of supply by capacity was [70-90]% in 2023, and it is forecast to increase to [✂]% by 2026.

Global shares of supply by data centre capacity

3.177 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*	*2024F*	*2025F*	*2026F*
AWS	[40-50]%	[40-50]%	[40-50]%	[40-50]%	[✂]	[✂]	[✂]
Microsoft	[30-40]%	[30-40]%	[30-40]%	[30-40]%	[✂]	[✂]	[✂]
Google	[10-20]%	[10-20]%	[10-20]%	[10-20]%	[✂]	[✂]	[✂]
Oracle	[0-5]%	[0-5]%	[0-5]%	[5-10]%	[✂]	[✂]	[✂]
IBM	[0-5]%	[0-5]%	[0-5]%	[0-5]%	[✂]	[✂]	[✂]
Other	[0-5]%	[0-5]%	[0-5]%	[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. Data from 2024 – 2026 are forecasts. 'Other' category includes CoreWeave, [✂], Hyve and Wasabi.

3.178 Global shares of supply by capacity indicate that AWS and Microsoft have the largest capacity, followed by Google and then by Oracle and IBM. The positions of AWS, Microsoft and Google appear unlikely to change in the next few years; however, the positions of Oracle, IBM and the 'other' small providers (CoreWeave, [✂], Hyve and Wasabi) may alter.[377]

(a) AWS is the largest global provider in terms of capacity, although its share has declined: AWS' share was [40-50]% in 2020 and has decreased to [40-50]% in 2023. It is forecast to decrease slightly to [✂]% by 2026.

(b) Microsoft is the second largest in terms of capacity with an increasing share: Microsoft's share was [30-40]% in 2020 and has decreased slightly to [30-40]% in 2023. It is forecast to increase slightly to [✂]% by 2026.

(c) Google is the third largest in terms of capacity with a declining share: Google's share was [10-20]% in 2020 and increased slightly to [10-20]% in 2023. It is forecast to decrease slightly to [✂]% by 2026.

(d) Oracle's share went from [0-5]% in 2020 to [5-10]% in 2023. IBM's share went from [0-5]% in 2020 to [0-5]% in 2023.

UK shares of supply by data centre capacity

3.179 UK shares of supply by capacity suggest that Microsoft is by far the largest cloud provider in the UK, followed by AWS.[378] Evidence suggests the relevant markets are wider than the UK, so we place less weight on these shares.[379]

Cloud providers' views

3.180 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.[380]

3.181 The denominators 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

[377] CMA analysis of data centre capacity in megawatts submitted by the cloud providers.
[378] 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. CMA analysis of data centre capacity in megawatts submitted by the cloud providers.
[379] See Market definition section above.
[380] AWS' submission to the CMA [✂].

slight overestimates. As described above, we place less weight on UK shares by capacity.

Shares of supply by flows of new business

3.182 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. 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).

3.183 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 provider, and expressing that growth as a proportion of the total revenue growth of all of the providers combined (calculated for both IaaS, and IaaS and PaaS combined).

(b) Shares by new customers acquired. This involves calculating the number of customers that each provider acquired as a proportion of the total customers acquired by all of the providers in each year.

(c) Shares by revenue from newly acquired customers. This involves calculating the revenues that each provider earned from customers acquired in each year as a proportion of the total revenues from new customers earned by all of the providers combined.

3.184 We consider that the analysis gives a useful indication of the relative importance of revenue growth from existing and new customers and which providers are gaining new business at a faster pace. However, the granularity of the data we have gathered means we cannot distinguish between the following:

(a) If the new customers a 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 the that provider and switching an existing workload (representing competition for existing workloads).

(b) If changes in a 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.

Our approach

3.185 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,[381] 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 market size 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 2023.

3.186 There are some caveats to this analysis:

(a) The definition of a new customer is binary and is also based on an arbitrary threshold. A new customer was defined as a customer that spent more than $100 for the first time in a year (in that 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. [382]

(b) In relation to the new customer data, each cloud provider submitted data on a slightly different basis.[383]

(c) The number of new customers analysis doesn't account for the size of customers, so all are given the same weighting.

3.187 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 supply by year-on-year UK IaaS revenue growth

3.188 For each firm we calculated its annual growth in IaaS revenues, then divided this by the sum of annual revenue growth across all providers to give the share.[384]

3.189 UK shares of supply by year-on-year IaaS revenue growth for 2021 to 2023 are presented in the table below.

[381] Centerprise, Civo, [✂], Hyve, Wasabi and UpCloud. Shares of supply by newly acquired customers and by revenue from newly acquired customers are calculated for AWS, Microsoft, Google, Oracle, IBM and these selected smaller IaaS providers only and do not include all cloud providers active in the UK.
[382] 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.
[383] See Appendix D.
[384] Similar to the shares of supply for IaaS, the underlying data includes some revenues from IaaS based on accelerated compute which could not always be distinguished in the Parties' data, but we understand the contribution of accelerated compute to the total revenues to be small.

Table 3.6: UK shares of supply by year-on-year IaaS revenue growth, 2021 – 2023

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

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.

3.190 Shares of supply by 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.

(a) AWS' share of revenue growth fell slightly over 2021 to 2023: it won [40-50]% and [40-50]% of overall new revenues in 2021 and 2023, respectively;

(b) Microsoft's share of revenue growth fell slightly from [30-40]% in 2021 to [30-40]% in 2023;

(c) Google's share of overall growth grew over 2021 to 2023 from [5-10]% to [10-20]%;

(d) For IBM and Oracle, shares of growth have generally remained in the 0-5% range.[385]

Shares of supply by year-on-year UK IaaS and PaaS revenue growth

3.191 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.

3.192 UK shares of supply by year-on-year IaaS and PaaS revenue growth for 2021–2023 are presented in the table below.

[385] CMA analysis of first-party revenue data and IDC and Synergy data.

Table 3.7: UK shares of supply by year-on-year IaaS and PaaS revenue growth, 2021 – 2023

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

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.

3.193 Shares of supply by year-on-year revenue growth in IaaS and PaaS combined show that AWS and Microsoft have the largest shares of overall revenue growth, though both declined in 2023. Google has the third largest share of overall revenue growth, consistent with its position in the market.

(a) Microsoft's share of revenue growth fell slightly from [30-40]% in 2021 to [20-30]% in 2023;

(b) AWS' share of overall revenue growth fell slightly over 2021 to 2023: it won [30-40]% and [30-40]% of overall new revenues in 2021 and 2023, respectively;

(c) Google's share of overall growth grew over 2021 to 2023 from [10-20]% to [10-20]%;

(d) For IBM and Oracle, shares of growth have remained in the 0-5% range; and

(e) Share of revenue growth from all other providers increased from [10-20]% in 2021 to [20-30]% in 2023. This is across an estimate of other providers based on more than 300 other providers.[386]

Shares of supply by newly acquired UK customers

3.194 For each firm, we divided its total number of new customers it acquired in each year by the sum of all providers' new customers to give the share. This is less useful than shares by revenues from new customers (see next section) as it doesn't consider the relative size of each customer.

3.195 UK shares of supply by newly acquired customers for 2021–2023 are presented in the table below.

[386] CMA analysis of first-party revenue data and IDC and Synergy data.

Table 3.8: UK shares of supply by newly acquired customers, 2021 – 2023

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

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, [✂], Hyve, Wasabi and UpCloud.

3.196 Shares of supply by newly acquired customers show that Microsoft is winning customers at a significantly higher rate than other cloud providers. It won more than twice as many customers in each of 2021 to 2023 as AWS, and at least three times as many as Google.

(a) Microsoft won more than [30-40]% of new customers in each year from 2021 to 2023.

(b) AWS' share of new customers was approximately [10-20]% in 2021 and 2022, decreasing to [0-10]% in 2023.

(c) Google's share of new customers is consistent with its other shares of supply presented above: it won [5-10]% of new customers in 2021 and [0-5]% in 2023.

(d) Oracle and IBM's share of new customers remained in the 0-5% range in 2021 to 2023.[387]

3.197 Our sensitivity analysis showed that these findings were robust to changes in the definition of a new customer.[388] Microsoft continued to have the largest share of number of new customers in the first two scenarios ($500 and $1,000), though its share drops as the threshold increases. In the final scenario ($10,000), AWS' share of new customer is the largest. This could suggest AWS has a larger proportion of existing customers who are increasing their spend.

[387] CMA analysis of revenue and customer acquisition data from Microsoft, AWS, Google, Oracle and IBM. Responses to the CMA's information requests [✂].
[388] Refer to Appendix D.

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

3.198 For each cloud provider, we calculated the annual growth in IaaS plus PaaS revenues from new customers, then divided this by the sum of annual revenue growth across revenues from all providers' new customers to give the share.

3.199 UK shares of supply by revenue from newly acquired customers for 2021 – 2023 are presented in the table below.

Table 3.9: UK shares of supply by revenue from newly acquired customers, 2021 – 2023

			%
	2021	*2022*	*2023*
Microsoft	[50-60]%	[60-70]%	[60-70]%
Google	[5-10]%	[5-10]%	[10-20]%
AWS	[10-20]%	[10-20]%	[10-20]%
Oracle	[0-5]%	[10-20]%	[5-10]%
IBM	[10-20]%	[0-5]%	[0-5]%
Other	[5-10]%	[0-5]%	[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, [✂], Hyve, Wasabi and UpCloud.

3.200 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.

(a) Microsoft's share of revenues from new customers has increased, from [50-60]% in 2021 to [60-70]% in 2023;

(b) AWS' position here is relatively weaker than in the other shares presented above: its share of revenues from new customers has decreased slightly from [10-20]% in 2021 to [10-20]% in 2023;

(c) Google appears to be relatively strong in winning new business: its share of revenues from new customers has increased slightly from [5-10]% in 2021 to [10-20]% in 2023;

(d) Oracle and IBM's shares ranged between [0-5]% to [10-20]% during this period.[389]

3.201 Our sensitivity analysis showed that these findings were robust to changes in the definition of a new customer.[390] The results showed that as the threshold

[389] CMA analysis of revenue and customer acquisition data from Microsoft, AWS, Google, Oracle and IBM. Responses to the CMA's information requests [✂].
[390] Refer to Appendix D.

increases, Microsoft's share of revenues from new customers declines, although Microsoft continued to have the largest share in each scenario.

Cloud providers' views

3.202 In response to our Competitive landscape working paper, Microsoft provided an analysis of cloud provider year-on-year revenue growth, based on quarterly S&P Capital IQ and 10-K financials data.[391] Microsoft said this showed Oracle has experienced double-digit revenue growth since Q4 2021, and IBM had double-digital growth between Q1 2020 and Q3 2022 (after which the revenue data was no longer available).

3.203 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 providers' year-on-year growth rate. We put more weight on the shares of supply by year-on-year UK IaaS, and IaaS and PaaS combined, revenue growth we have calculated because this is based on first party revenue data.

3.204 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, the 2022 shares are: Amazon 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).[392]

3.205 We consider the weight we can place on this analysis may be limited for the following reasons: this data relates to a global market definition rather than European; 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.

Provisional conclusions

3.206 We have provisionally found that the cloud services sector is highly concentrated, particularly in IaaS and also, although to a lesser extent, in PaaS. Forward looking metrics suggest this position is likely to endure.

[391] Microsoft's response to the Competitive landscape, Committed spend agreements and Egress fees working papers dated 23 May 2024, paragraph 20.
[392] Microsoft's response to the Competitive landscape, Committed spend agreements and Egress fees working papers dated 23 May 2024, paragraphs 34-37.

3.207 We place most weight on shares which map to the product markets we have defined: these are separate product markets for IaaS and PaaS, rather than the shares which consider IaaS and PaaS combined. If we changed the boundaries we have defined for IaaS and PaaS slightly, this would be unlikely to change our overall findings of concentration because we have found that the markets for IaaS and PaaS individually are concentrated.

3.208 In particular, the IaaS market remained highly concentrated in 2023. While the IaaS market is growing overall and this may provide opportunities for smaller providers as there is still new business to be won, the small ([0-5]%) shares of other providers and lack of growth in these shares over the period 2020 to 2023, indicates that these providers are likely to remain considerably smaller than their rivals.

3.209 The PaaS market remains concentrated but 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 is also an IaaS sale for one of AWS, Microsoft 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.

3.210 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 and EEA based on capacity since 2021.

(b) Shares of supply by year-on-year revenue growth in (i) IaaS and (ii) IaaS and PaaS combined both show that AWS and Microsoft have the largest shares of overall revenue growth; and shares of supply by revenue from new customers show that Microsoft is winning significantly more new customers than other providers.

(c) While Google's market share has increased in recent years, it remains significantly smaller than that of Microsoft or AWS whose combined market share remains 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.

Market outcomes

Introduction

3.211 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.[393] 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'.[394]

3.212 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

3.213 AWS and Microsoft 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.[395] 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'.[396]

3.214 Our guidelines state that outcomes of the competitive process in their different forms in a market can provide evidence about its functioning.[397] Such outcomes include profitability, prices, quality and innovation. Of these, profitability and prices are among the more observable and measurable outcomes, 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.[398] Our guidelines note that evidence on quality and innovation, on the other hand, tends to be qualitative.[399]

3.215 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 more measurable outcomes, our guidelines are clear that they should be treated

[393] CC3 (Revised), paragraph 103.
[394] CC3 (Revised), paragraph 103.
[395] 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.
[396] AWS' Economic response to the CMA's Competitive landscape working paper, paragraphs 108-109.
[397] CC3 (Revised), paragraph 103.
[398] CC3 (Revised), paragraphs 104 and 126.
[399] CC3 (Revised), paragraph 127.

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,[400] depending on the broader market context.[401]

3.216 Notwithstanding our view that it would not be appropriate to specify quantitative benchmarks, a number of providers also suggested comparative benchmarks, none of which we think are appropriate for this market:[402]

(a) **Comparison with on-premises:** Some providers have submitted that the cloud market produces better outcomes than traditional IT services.[403] 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:** [✂] and Microsoft have submitted evidence to show continuing improvement in relation to particular market outcomes, for example that prices have fallen over time or that numbers of patents have increased.[404] However we think using the past as a benchmark in this way is only instructive to a degree because of the risk of concurrent events influencing market outcomes

3.217 In the rest of this section, we consider:

(a) profitability;

(b) pricing trends;

(c) quality and innovation;

(d) multi-cloud; and

(e) switching.

[400] 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.)
[401] 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.
[402] 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).
[403] See Market definition section.
[404] [✂] submission to the CMA [✂]; Microsoft's submission to the CMA [✂].

Profitability

3.218 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.[405] 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.

3.219 In this section we set out the role of profitability analysis and assess the profitability of cloud providers. We have sought to consider whether profits among certain providers reflect a 'normal' rate of return based on the nature of competition in the supply of public cloud infrastructure services.[406]

Role of profitability analysis

3.220 The aim of profitability analysis is to understand competitive conditions within a market, by examining the outcomes of that market in terms of the financial performance of the participating firms.

3.221 We consider 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.[407] The profitability of firms representing a substantial part of the market can therefore be a useful indicator of competitive conditions in a market.[408]

3.222 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.

3.223 We do not regard 'excess' profitability 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.[409]

[405] CC3 (Revised) paragraphs 104, 114 and 116.
[406] We provide additional details on our analysis in Appendix E.
[407] CC3 (Revised), paragraph 116.
[408] CC3 (Revised), paragraphs 116, 118 and 119.
[409] CC3 (Revised), paragraphs 117, 118 and 126.

Stakeholder views on profitability analysis

3.224 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'.[410]

3.225 Two cloud providers submitted that gross margins are not meaningful or sufficiently informative metrics for comparing profitability between different providers.[411] 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.[412]

3.226 An academic submitted that there were issues with use of global profitability to assess the competitive landscape in the UK.[413]

3.227 We received the following submissions on the relevance of profitability to assessing competition:

(a) AWS submitted that 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.[414]

(b) AWS submitted that there is no empirical link between cloud service providers' gross margins or operating margins and their size/growth, thus there is no systematic evidence that larger cloud service providers are more profitable.[415]

(c) AWS also submitted that our Competitive landscape 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

[410] Microsoft's submission to the CMA Microsoft's response to the Competitive landscape, Committed spend agreements and Egress fees working papers dated 1 July 2024, paragraph 32. .
[411] Submissions to the CMA [✂].
[412] [✂] submission to the CMA [✂].
[413] R. Parisi, The Cloud Services Markets' Competitive Landscape: A contribution to the Competition and Markets Authority, page 13.
[414] AWS' submission to the CMA [✂].
[415] AWS' submission to the CMA [✂].

levels exceeding the WACC do not represent evidence against the market being workably competitive.[416]

(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.[417]

Margins and ROCE

3.228 AWS submitted that its gross margins are in line with those of other cloud service providers for the period 2021-2023.[418]

3.229 AWS submitted that ROCE estimates for AWS show a downward trend, and that with the inclusion of estimates for brand value, goodwill and IPR&D intangible assets, AWS' ROCE ranges between [✂] during 2013-2023. [✂].[419]

Innovation

3.230 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.[420]

(b) Two academics 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.[421] One of these characterised the nature of the cloud market as new, risky, uncertain,

---

[416] AWS' submission to the CMA [✂].
[417] Microsoft's response to the Competitive landscape, Committed spend agreements and Egress fees working papers, paragraph 32.
[418] AWS' submission to the CMA [✂].
[419] AWS' submission to the CMA [✂].
[420] AWS' submission to the CMA [✂].
[421] 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. 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). R. Parisi's submission did not include a declaration of conflicts.

innovative technology or way of doing business, and noted that in the long run, the profitability or rate of return falls in markets.[422]

Investment

3.231 In our Competitive landscape working paper,[423] 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 follow submissions:

(a) AWS submitted that the presence of significant investments is consistent with effective competition.[424]

(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'.[425]

Our analysis of profitability

Our approach

3.232 We consider that we can use 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 time period.

3.233 Also, a lack of 'steady state' in the market might be relevant if we saw significant fluctuation in ROCE or profit margins, however our analysis finds that margins for AWS and Microsoft's cloud businesses have been broadly stable (or growing in the case of Microsoft's Azure) and even with the recent increase in capex spending (discussed further below) AWS' and Microsoft's cloud businesses have consistently achieved ROCE in excess of WACC.

[422] 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.
[423] Competitive landscape working paper, paragraphs 6.41 and 6.42
[424] AWS' submission to the CMA [✂].
[425] Microsoft's response to the Competitive landscape, Committed spend agreements and Egress fees working papers, paragraph 33.

3.234 We note that two cloud providers consider gross margins to be uninformative metrics for comparing profitability between different providers.[426] 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, but note that this often has limitations when seeking to determine if profitability exceeds a 'normal' level which makes other measures, such as ROCE, preferable where possible.

3.235 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.[427] We examine the profitability of cloud services for the largest cloud providers in the UK (AWS, Microsoft 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.[428]

3.236 We examined the global profitability of providers in our analysis. While one stakeholder submitted that the use of global data to assess profitability for the UK cloud market was incorrect,[429] we consider it more suitable to assess the global profitability of providers due to (i) the global nature of 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 level which limits our ability to do meaningful UK-level profitability assessment.

3.237 We have analysed gross margins and earnings before interest and tax (EBIT) margins for cloud providers as indicators of financial performance.

3.238 We have also analysed and compared the return on capital employed (ROCE) for AWS and Microsoft's Azure and Cloud & Enterprise business segments, as these are the two largest providers in the UK markets and represent [60%-70]% of the market together,[430] to our estimate for the weighted average cost of capital (WACC) to assess their profitability.[431]

---

[426] Submissions to the CMA [✂].
[427] Our profitability analysis includes revenues, costs and capital base from IaaS based on accelerated compute which is not separately identifiable in the Parties' data.
[428] 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.
[429] R. Parisi, The Cloud Services Markets' Competitive Landscape: A contribution to the Competition and Markets Authority, page 15.
[430] This is the share of the combined UK IaaS and PaaS markets, of which AWS and Microsoft each have [30-40]% market share based on revenue. See Market structure and concentration above.
[431] 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.

3.239 We do not include Google in our ROCE analysis as it has only recently (since 2023) reported profits,[432] and it has a significantly lower share of cloud services revenues compared to AWS and Microsoft.[433] Similarly we do not include other smaller cloud providers which have lower market shares in our ROCE. We consider the profitability trends of Google and other cloud providers in our margin benchmarking for comparative purposes, but as we consider the profitability of firms representing a substantial part of the markets to be a useful indicator of competitive conditions in the cloud markets, we are primarily interested in assessing the profitability of the largest incumbent providers in the markets.[434]

3.240 We compare the ROCE for AWS and Microsoft's cloud businesses to the WACC, in order to assess the extent to which these providers earn a 'normal' rate of profit. Where firms persistently earn in excess of a normal return, this signals that there may be limitations in the competitive process. 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.[435]

3.241 We have examined profitability over the last five years but have 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).[436]

3.242 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.[437] We consider that AWS and Microsoft each entered the cloud markets over ten years ago, and that much of their initial investment 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 Amazon and Microsoft's other businesses). We consider that 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 reflect their steady state profitability in the markets.

3.243 We cross-check 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

[432] Alphabet 2023 Form 10-K, page 87.
[433] 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 [60-70]% market share.
[434] CC3 (Revised), paragraphs 114 and 116.
[435] CC3 (Revised), paragraphs 118 and 119.
[436] [✂] response to the CMA's information requests [✂].
[437] CC3 (Revised), paragraph 121 and Annex A, paragraph 10.

some point, there is a clear downward trend towards the cost of capital that is apparent.

Our analysis

3.244 Below we summarise the findings of our profitability analysis.[438] We have found that, with regard to gross margins:

(a) AWS, Microsoft, Google and Oracle have generated gross margins for their cloud businesses which are higher than those of other providers.

(b) AWS and Microsoft's Cloud and Cloud & Enterprise segments have consistently had the highest gross margins, in excess of [✂]% for AWS and [✂]% for both Microsoft segments over the last five years.

3.245 We have found with regard to EBIT margins that:

(a) The EBIT margins for AWS have consistently been between 24% and 30% for the last nine financial years (2015 to 2023) and are increasing in 2024.[439]

(b) The EBIT margins for Microsoft's Azure segment have consistently been between [20-30]% and [30-40]% for the last four years, and for Microsoft's Cloud and Enterprise segment have consistently been between [40-50]% to [50-60]% for the last eight years.[440] 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 nine percentage point improvement in margin in the last 24 months).[441] The EBIT margins for Google Cloud are currently significantly lower than AWS [✂].

3.246 With regard to ROCE analysis, we have found:

(a) AWS ROCE has consistently been substantially above our estimated WACC of [10-20]% to [10%-20]% for the last nine 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 of [10-20]% to [10-20]% and in excess of [✂]% under all sensitivities used in our analysis for the last seven years.

[438] Our full profitability analysis is set out in Appendix E.
[439] CMA analysis of Amazon Form 10-Ks and 10-Qs.
[440] CMA analysis of Microsoft Form 10-Ks and Microsoft's responses to the CMA's information requests [✂].
[441] Period to 30 September 2024. CMA analysis of Alphabet Form 10-Qs.

(c) Microsoft Azure ROCE has been substantially above our estimated WACC for the last four years and is trending upwards. Microsoft Azure ROCE has been in excess of [✂]% under all sensitivities used in our analysis for the last three years.

3.247 Our ROCE analysis shows that the ROCE for AWS and Microsoft Cloud & Enterprise ROCE has been falling slightly since financial year 2021 ([✂] and [✂] respectively for our baseline ROCE calculations), but still substantially above the WACC.

3.248 Our estimate of forecast AWS ROCE for financial year 2024 using AWS' internal forecasts for EBIT and capex also indicates AWS ROCE will [✂].[442]

Interpretation of our analysis

3.249 We recognise that there can be many factors that influence profitability in any given year but our analysis focuses on whether there have been returns in excess of the costs 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.

3.250 We do not consider it 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.[443]

3.251 We compare cloud providers' margins in our analysis but do not put forward any views on whether there is a causal relationship between cloud provider margins and size/growth.

Margins and ROCE

3.252 As set out in our analysis, we find AWS and Microsoft's Cloud and Cloud & Enterprise segments have consistently had the highest gross margins. AWS' submission that its gross margins are in line with other 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).

---

[442] CMA analysis of Amazon Form 10-Ks and AWS' response to the CMA's information request [✂]. We have not extended our ROCE forecast for AWS beyond 2024 as [✂]. AWS response to the CMA's information request [✂].
[443] CC3 (Revised), paragraphs 120, 122 and 124.

3.253 In the analysis AWS submitted [✂].[444]

3.254 We do not agree with AWS' submission that its ROCE ranges between [✂] during 2013-2023, as these figures are based on estimates of capital employed that include intangible assets that we have not received supporting evidence for and have determined do not merit inclusion in capital employed.[445]

3.255 AWS submitted that there is a clear downward trend in AWS' ROCE over the last three years (2021-2023). AWS submitted that this implies a convergence between ROCE and WACC which suggests the presence of fierce competition despite a temporal competitive advantage granted to AWS 'through its constant innovation'.[446]

3.256 However, as set out in the following section, our analysis indicates that the downward trend for AWS ROCE (in the last two years) is largely driven by a new wave of investment in AI which has impacted AWS' capital employed for its overall cloud business. We also note that AWS ROCE remains in excess of WACC, at [20-30]% [✂].[447]

3.257 We discuss our interpretation of the recent decrease 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.

3.258 We consider innovation and investment in the following section.

Considerations when assessing ROCE

3.259 As stated above, we will consider reasons why ROCE may exceed WACC for a firm. We have not seen any indications, and none have been suggested to us, that AWS and Microsoft Azure and Cloud & Enterprise 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.

3.260 As explained above, we consider the last five years to be relevant to assess the profitability of AWS' and Microsoft's cloud businesses using our ROCE analysis. We also recognise that 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 need to consider profitability over a relatively long timescale.[448] We consider

---

444 AWS' submission to the CMA [✂].
445 See criteria in CC3 (Revised), annex A paragraph 14. We set out our consideration of intangible assets in more detail in Appendix E.
446 AWS' submission to the CMA [✂].
447 For our baseline ROCE estimate. CMA analysis of Amazon Form 10-Ks; AWS' response to the CMA's information request [✂].
448 CC3 (Revised), paragraph 121.

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 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),[449] we consider seven or more years to be a 'sustained' period rather than point in time for the cloud market.

3.261 In addition, AWS' and Microsoft's investments in their cloud businesses will be taken into account in our measure of capital employed used in calculating ROCE. As AWS and Microsoft have characterised their investments in the cloud market as continuous,[450] (rather than an initial one-off investment), these investments will continually be adding to capital employed and hence reflected in our ROCE figures.

3.262 We have also considered whether the profitability of AWS' and Microsoft's cloud businesses may be explained by a high degree of innovation in the cloud market.

3.263 While 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,[451] we consider the ROCE substantially in excess of WACC for AWS and Microsoft's cloud businesses over a sustained period to be an indication of limitations in the competitive process in the cloud market.

3.264 We also note 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, particularly due to the ongoing rather than one-off nature of investment that cloud providers have described. For example, neither AWS nor Microsoft 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.

3.265 We therefore do not consider the sustained ROCE in excess of WACC for AWS' and Microsoft's cloud businesses to be explained by the level of innovation in the cloud market.

[449] CC3 (Revised), paragraph 117.
[450] 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.
[451] CC3 (Revised), paragraphs 117 and 118.

Impact of investments in AI

3.266 Additionally, we note that AWS and Microsoft appear to have been recently substantially increasing their investments in their cloud infrastructure, in particular in investments to support the development of their AI services.

3.267 Whilst both AWS and Microsoft have 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.

3.268 AWS and Microsoft have been making substantial capex investments in supplying accelerated compute (see our assessment of this later in this chapter). For example:

(a) [✂];[452] and

(b) Microsoft's cloud infrastructure and server capex in financial years 2022 and 2023 was [✂].[453]

3.269 This increased investment has contributed to the recent downward trend in ROCE for AWS and Microsoft Cloud & Enterprise, as it increases capital employed at a greater rate than the growth in EBIT. This may indicate that a new investment cycle is commencing for these firms' cloud businesses in relation to AI services.

3.270 For a new investment cycle, we might expect to see initial upfront investment, that then generates returns in the future. Our review of AWS and Microsoft's public announcements and internal documents indicates that both companies are expecting large, accelerating revenue growth from their AI services.[454]

3.271 AWS and Microsoft have noted the [✂] margins resulting from the significant capex associated with AI. Microsoft's internal documents also indicate that:

(a) [✂].[455] [✂].[456]

(b) [✂].[457]

(c) [✂].[458]

---

[452] CMA analysis of AWS' response to the CMA's information request [✂].
[453] CMA analysis of Microsoft's response to the CMA's information request [✂]; Microsoft's response to Ofcom's information request [✂].
[454] See The impact of AI on cloud services section below.
[455] Microsoft's response to the CMA's information request [✂].
[456] Microsoft's response to the CMA's information request [✂].
[457] Microsoft's response to the CMA's information request [✂].
[458] Microsoft's response to the CMA's information request [✂].

3.272 Based on the comments in relation to AI growth in Microsoft's public announcements[459] 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.

3.273 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.

3.274 Similarly, our analysis of [✂].[460] As noted above, our estimate of forecast AWS ROCE for financial year 2024 indicates AWS ROCE will [✂].

3.275 We recognise that AWS and Microsoft 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 AWS and Microsoft includes their AI cloud investment and the comparatively [✂] profit contribution associated with this currently, and ROCE levels have consistently remained above WACC even with this increase in investment.

3.276 Overall, we consider the forecast data and commentary reviewed for AWS and Microsoft to be consistent with ROCE for each of their cloud businesses stabilising at around current levels in the future, rather than indicating future decline.

Pricing trends

3.277 In markets subject to effective competition, prices are likely to respond to changing supply and demand conditions and 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 the nature of competition. However, there may be several factors affecting prices and we will take this into account when considering inferences from this type of analysis.[461]

3.278 In particular, we have sought to understand the trends in prices over time among AWS, Microsoft and Google and whether that can tell us anything about competition in cloud services. In doing this, 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 only indicative and needs to be

[459] Microsoft Q2 FY24 Earnings Conference Call transcript; Microsoft Q3 FY24 Earnings Conference Call transcript.
[460] CMA analysis AWS' response to the CMA's information request [✂].
[461] CC3 (Revised), paragraph 108

considered alongside other evidence such as costs – this is why we also considered profitability above.

3.279 Pricing trends can be informative in that a trend of increasing prices over time in a sector may be consistent with a competition concern subject to other supporting evidence. For example, it may not be profitable for a firm to increase or even hold prices stable over time when competition is working effectively, as we might expect rivals to undercut incumbent firms in a bid to win customers. If a firm is increasing prices (or holding prices constant while costs are falling, for example, due to economies of scale) while maintaining market share over time then this could suggest that that firm has market power.

3.280 Google submitted that the cloud market is characterised by a downward pricing trend, and the availability of numerous deals and pricing offers.[462] Dr Barker stated that prices have fallen, not risen.[463] Mr Parisi stated that cloud computing's prices continue to be deflationary despite economic inflation, which highlights the competition among advanced companies.[464]

3.281 During Ofcom's market study, AWS submitted to Ofcom an analysis of global net price trend data and UK list price trend data for [✂] cloud infrastructure services (S3, EC2 and data transfer-out).[465] AWS said its analysis shows that:

(a) global average net prices decreased significantly in both nominal and real terms for all three services;

(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.[466]

3.282 AWS also noted that it has innovated and improved the quality of these three services, implying a further reduction in quality-adjusted real prices.[467]

[462] Google, response to the CMA's Issues Statement dated 17 October 2023, paragraph 9.
[463] 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.
[464] R. Parisi, The Cloud Services Markets' Competitive Landscape: A contribution to the Competition and Markets Authority, page 11.
[465] AWS' submission to Ofcom [✂].
[466] AWS' submission to Ofcom [✂].
[467] AWS' submission to Ofcom [✂].

3.283 Microsoft provided us with pricing analysis that it had conducted.[468] Microsoft said its analysis shows that quality-adjusted real prices have fallen.[469] 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.[470]

3.284 We have reviewed both price analyses submitted by AWS and Microsoft. We have some reservations about some of the methodological choices taken by both Microsoft and AWS, which potentially overstate prices decreases:

(a) For the AWS analysis, we note there are differences in scope between the UK list price and global net price series (beyond differences in geographic scope) which mean these are not directly comparable.

(b) Microsoft and AWS deflated prices using the consumer price indexes (CPIs). These are a general inflation indexes that might not reflect (they could either underestimate or overestimate) the specific inflation of costs related to the cloud industry.

(c) Microsoft's categorisation of prices as decreasing, not changing, or increasing omits any information related to the magnitude of such changes.

(d) Microsoft and AWS both analysed unweighted prices, ie it considered each product equally regardless of its importance (eg in terms of customer spend on them).

(e) Microsoft and AWS 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.

3.285 In our own analysis, we improved on all of these aspects.[471] We applied a consistent methodology to data provided by Microsoft and Google; however, AWS was not able to provide us with data to perform a similar analysis.

3.286 Our analysis shows that the picture is more mixed than that presented by Microsoft. When looking at real (ie inflation-adjusted) prices net of discounts, we find that:

---

[468] Microsoft's submission to the CMA [✂]. 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.
[469] Microsoft's Response to the Competitive landscape, Committed spend agreements and Egress fees working papers, paragraph 12 (a) and (b).
[470] Microsoft's submission to the CMA [✂].
[471] See Appendix F for a detailed description of this.

(a) Some services show declining prices. However, some services increased in price over the last five years.

(b) Within each service, there is a variety of price trends at the product level. In particular, even for those services whose overall price decreases, a significant share of underlying products (which make up services) increased in prices; and the other way round.

(c) Price changes are driven by old products being discontinued and replaced by new ones, rather than price changes to single products.[472]

3.287 In interpreting these results, we account for the fact that this analysis does not:

(a) account for changes in costs. Therefore, looking at profitability is important;

(b) account for changes in quality of the products; and

(c) cover all providers to the same extent: AWS did not provide comparable data, while Google's data was more limited in scope than Microsoft's.

3.288 Therefore, we assign more evidential weight to our assessment of the 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) and has a clear and established method for assessing the market outcomes against a benchmark (the cost of capital), which can be used to compare across providers.

3.289 The provisional conclusion of our assessment is that there is no clear trend in prices: some services and products are increasing in prices while others are decreasing.

Quality and innovation

3.290 Outcomes related to quality, innovation and range may also be useful indicators of competition.[473] The Jigsaw report indicates that customers value innovation and find value in sharing the burden of architecture design, upgrades and maintenance. Many businesses valued providers innovating on infrastructure and the user being able to focus on their core business function.[474]

3.291 In order to assess these market outcomes, we have assessed a range of metrics including:

[472] CMA analysis of responses to the CMA's information requests [✂].
[473] CC3 (Revised), paragraph 127.
[474] Jigsaw Research, page 25.

(a) product quality assurance;

(b) customer satisfaction;

(c) software updates;

(d) the number of new services and features of services launched;

(e) the number of patents and patent citations received; and

(f) levels of investment in R&D.[475]

3.292 AWS, Microsoft and Google submitted that levels of quality and innovation in the market are high.[476] For example, Google said that the cloud market is characterised by continuous innovation.[477]

3.293 AWS submitted that:

(a) innovation is omnipresent in the cloud segment;[478]

(b) [✂];[479]

(c) the quality of AWS' innovations is evidenced by the fact that the innovations are often directly responsive to customer requests and are aimed at ensuring the quality of a service for customers;[480]

(d) [✂];[481]

(e) it regularly introduced additional services and features each year between 2011-2023, and new EC2 instance families each year between 2008-2022;[482] and

(f) providers are characterised by a high degree of innovation consistent with a dynamically competitive market.[483]

3.294 These results are sensitive to the time period considered. For example, for AWS, if we only looked at the period 2019-2023, the [✂] would have been of less than

[475] We present this evidence in detail in Appendix G.
[476] AWS' submissions to the CMA [✂]; AWS Economic response to the CMA's Competitive landscape working paper paragraph 110 and 112; Microsoft's submission to the CMA [✂]; Google's response to the CMA's Issues Statement dated 17 October 2023, paragraph 9.
[477] Google's response to the CMA's Issues Statement dated 17 October 2023, paragraph 9.
[478] AWS' Economic response to the CMA's Competitive landscape working paper, paragraph 110.
[479] AWS' submission to the CMA [✂].
[480] AWS' Economic response to the CMA's Competitive landscape working paper, paragraph 112.
[481] AWS' submission to the CMA [✂].
[482] AWS' submission to the CMA [✂].
[483] AWS' submission to the CMA [✂].

[✂] and the [✂] would be considerably less significant than during the period 2018-2023.[484]

3.295 Similarly, Microsoft submitted that its cumulative number of feature updates for the top 20 UK cloud services was large in 2023.[485] In addition, Microsoft presented a submission on innovation in the cloud market.[486] In this submission, Microsoft:

(a) presented the advantages of cloud relative to other forms of IT models as well as significant technological developments in the industry;[487]

(b) said that R&D spend and capital investment has increased significantly since 2010;[488]

(c) said that it advanced first party and third party innovation across a range of services and features, throughout the whole stack and with significant customer uptake;[489] and

(d) said that innovation occurred also at the infrastructure level.[490]

3.296 We have gathered both quantitative and qualitative information that relate to quality and innovation, but different business practices across providers as well as the different availability of data limits our ability to assess trends over time or compare metrics across providers. Information on quality and innovation is also less readily quantifiable than prices and costs.

3.297 As discussed above, we also consider that specific thresholds for assessing levels of quality and innovation would not be necessary or helpful, as it is not feasible to determine exactly at what level and based on what measures the quality or extent of innovation would be consistent with a market without the features that we have considered in this investigation.

3.298 Even if there was a way to usefully assess whether a given measure of innovation implies 'high' or 'low' levels of innovation, we consider that looking at trends in measures of quality and innovation over time or across providers would provide an incomplete picture at best of the extent to which competition is driving any changes and of the extent to which any feature or features of the market may be harming competition and leading to a reduced level of innovation.

3.299 In this respect, AWS said that innovation brought by the prospect of growing the market would be an indication of the value added that is created by innovation and

[484] CMA analysis of AWS' submission to the CMA [✂].
[485] Microsoft Response to The Competitive landscape, Committed spend agreements and Egress fees working papers, paragraph 12 (c).
[486] Microsoft's submission to the CMA [✂].
[487] Microsoft's submission to the CMA [✂].
[488] Microsoft's submission to the CMA [✂].
[489] Microsoft's submission to the CMA [✂].
[490] Microsoft's submission to the CMA [✂].

that it would also increase output which is another positive outcome.[491] We agree that innovation that expands the market is a positive outcome. However, the observation that there may be innovation to expand the market is not inconsistent with features of the market restricting competition. If, for example, in a hypothetical scenario, customers were not able to switch or multi-cloud, cloud providers may still have an incentive to innovate to the extent this would generate additional revenue and profits from their cloud services, particularly where the revenues from those innovations exceed the cost of introducing them.[492] However, their incentives to innovate may still be stronger if customers of cloud services were able to put competitive pressure on providers by being able to switch and/or multi-cloud because more of the cloud providers' revenues would depend on innovating.

3.300 We consider that, while it is difficult to evaluate evidence on quality and innovation, there has been some innovation in the market and we have seen that providers innovate and invest to win business. This is also reflected in the perception that customers have of some of the providers as 'great innovators'.[493] However, while we acknowledge the benefits this brings to customers, we cannot conclude purely on the basis that innovation is taking place that competition is working fully effectively.

3.301 We assign greater evidential weight to our assessment of the providers' profitability than the analysis of quality and innovation indicators described above. This is because the former accounts for broader financial indicators (eg revenue, which is determined by prices and other variables, and costs including costs to improve quality and innovate) and has a clear and established method for assessing the market outcomes against a benchmark (the cost of capital), which can be used to compare across providers.

Prevalence of switching and multi-cloud

3.302 We have considered how common it is for customers to switch or multi-cloud.

3.303 The level of switching or use of multiple clouds could have two different interpretations:

(a) On the one hand, 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.

[491] AWS, Economic response to the CMA's Competitive landscape working paper, paragraph 116.
[492] For example, such innovations may involve developing additional features that can then be charged at a premium or products and services that do not overlap with existing products and services such that they generate incremental revenue. In this respect there are different types of innovation and each type may have different implications for competition.
[493] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), page 25.

(b) On the other hand, a lack of switching or multi-cloud would 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.

3.304 This second interpretation, supported by some cloud providers,[494] may be more compelling where the level of switching and use of multiple cloud providers is not particularly low (as particularly low switching or multi-cloud may be inconsistent with customers being able to make a credible threat to switch or multi-cloud), and where other evidence on the ease of switching and on the broader market context (such as profitability) are consistent with that interpretation.

Prevalence of switching

3.305 We have assessed how common switching is in the market.

Cloud providers' views

3.306 AWS said that switching is far more prevalent than Ofcom's Final Report suggested and that, due to the way cloud services are designed and priced (pay-as-you-go) compared to previous IT environments, it has never been easier for customers to switch IT provider.[495] 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).[496]

3.307 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.[497] 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).[498]

3.308 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.[499] 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.[500]

[494] AWS' response to the CMA's updated issues statement and working papers, paragraph 18.
[495] AWS' response to the Issues Statement, 17 October 2023, paragraph 17.
[496] AWS' response to the Issues Statement, 17 October 2023, paragraph 14.
[497] Microsoft's response to the Issues Statement, 17 October 2023, paragraph 26.
[498] Google's response to the CMA's information request [✂].
[499] Microsoft's submission to Ofcom [✂].
[500] Microsoft's submission to Ofcom [✂].

(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.[501]

3.309 AWS and Google identified practices that restrict customer choice and make switching more difficult such as certain licensing practices.[502]

3.310 Microsoft said that, given the prevalence of multi-cloud (which we discuss above), 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).[503]

3.311 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.[504]

3.312 However, providers submitted 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. 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.[505]

(a) Google gave an estimate of the number of customers that were not billed in that 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 a number of 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. Google submitted that it lost around [✂]% of its UK Google Cloud Platform customers in 2022, altogether accounting for $[✂] in revenue in 2021.[506]

---

501 Microsoft's submission to Ofcom [✂].
502 AWS' response to the Issues Statement, 17 October 2023, paragraph 33; Google's submission to the CMA [✂].
503 Microsoft's response to the CMA's information request [✂]; Microsoft's submission to the CMA [✂].
504 Microsoft Response to the Competitive landscape, Committed spend agreements and Egress fees working papers, paragraph 17.
505 AWS' response to the CMA's information request [✂].
506 Google's responses to the CMA's information requests [✂].

(b) 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 [✂] Google, but there are limitations to what we can take from this given that [✂].[507]

Quantitative analysis from one provider

3.313 One provider submitted quantitative analysis of its customer data and said that the results show that customers can and do switch.[508] According to this analysis, approximately [✂]% of its customers churned on an annual basis, considering all the cloud services it offered. The provider said that [✂]% of its customers, accounting for [✂]% of its revenues, [✂] 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.

3.314 We consider that there are limitations to this quantitative analysis on switching.[509] 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 [✂]%), the much lower churning rates for customers with high spend, and limitations as to how customers are identified as having churned. We consider that the results from this quantitative analysis should be interpreted in light of these caveats.

Quantitative surveys

3.315 Below, we consider the evidential value we can attribute to quantitative surveys used to measure switching and multi-cloud in the cloud market.[510]

3.316 In general, quantitative surveys have some advantages: if the sample is representative, the results are generalisable to the whole population. 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.[511]

3.317 However, we have concerns about using surveys to assess the prevalence of multi-cloud and switching in the cloud market due to:

[507] [✂] submission to the CMA [✂].
[508] [✂] response to the CMA's information request [✂]; Ofcom's analysis in its Final Report, Annex 3, paragraph A3.37
[509] We include our full assessment of this analysis in Appendix H.
[510] A more detailed analysis of these surveys can be found in the Appendix C.
[511] 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).

(a) Uncertain validity: quantitative surveys require the customers spoken to, 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.[512]

(b) Vulnerability to the quality 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 of these factors in this market. 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.[513] [514]

3.318 For these reasons, we consider that there are significant weaknesses with using a quantitative research method to estimate the prevalence of multi-cloud in this market. We therefore place far more reliance on the in-depth qualitative research we commissioned ourselves than the quantitative survey evidence submitted to us by cloud providers. We have chosen to use qualitative research methods for our customer research in this investigation as have judged that these are better suited to researching our key research questions and navigating some of the complexities of customer behaviour in this technical market.

Surveys on prevalence of switching

3.319 AWS and Microsoft submitted a survey (Public First) that included an estimate of the prevalence of switching in the market.[515] 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.[516]

---

[512] For example, as identified in the 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.

[513] 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).

[514] 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.

[515] AWS' submission to the CMA [✂]; Microsoft's submission to the CMA [✂].

[516] Our views on quantitative research in cloud services are set out in Appendix C.

3.320 As discussed in Appendix C, we consider there are significant limitations to such estimates, in particular in relation to uncertain validity and vulnerability to the quality and coverage of the customers spoken to.

3.321 Notwithstanding these methodological limitations, we also consider that the survey evidence referenced by AWS and Microsoft offers only limited insight into the extent of switching by customers between public cloud providers. For example, the Public First survey asked: 'Have you ever switched one of your cloud infrastructure providers in the past?' [517] As phrased, the question is not informative on the materiality of customer switching, ie in terms of the proportion of workloads that were switched from one cloud provider to another. The customers spoken to could answer 'yes' even if they switched only a handful of workloads from one cloud provider to another. 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. The results of the survey are also at odds with our own analysis of full switching based on customer data – this is described below.

3.322 Therefore, we consider it appropriate to place limited evidential weight on the quantitative evidence we have received.

3.323 The research commissioned by Ofcom included both a qualitative research phase and a quantitative online survey. [518] The quantitative online survey has the same limitations outlined earlier about the available public survey evidence in this market, relating to issues of validity and sample quality. The qualitative research is not subject to the same methodological limitations.[519]

3.324 We note that these qualitative results also indicate that switching is lower than the quantitative results reported. 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 AWS, Microsoft or Google to another cloud provider;

---

[517] Public First Poll [✂].
[518] Cloud Services Market Research - Summary of Findings March 2023.
[519] 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.

(c) the switching that was described by customers tended to involve a relatively small portion of data and workloads, moving from one minority 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

(e) in most cases, firms were still on their way into, not out, of their IaaS/PaaS environments.

3.325 In light of these limitations, we consider it appropriate to place limited evidential weight on the quantitative results on the prevalence of switching from Ofcom's research.

Customer views

3.326 The Jigsaw report states that few of the customers interviewed had chosen to switch between public cloud providers.[520] 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.[521] 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.[522]

3.327 We also collected customer evidence on switching as part of our assessment of technical barriers and committed spend agreements.

3.328 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.[523] 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',[524] or in relative terms eg some customers described technical barriers as the main barrier to switching.[525] In particular:

[520] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 1.3.14.
[521] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 1.3.16.
[522] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 1.3.13.
[523] Responses to the CMA's information requests [✂]; Notes of meetings with [✂].
[524] Note of meeting with [✂].
[525] Responses to the CMA's information requests [✂]; Notes of meetings with [✂].

(a) Some of these customers indicated that these costs had stopped them from switching or considering switching.[526]

(b) Some customers also said that there are a range of operational concerns that arise when they consider switching between public clouds.[527] This is consistent with the findings of the Jigsaw report.[528]

(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.[529] This is consistent with the findings of the Jigsaw report.[530]

3.329 This last point was corroborated by other customer evidence. Some customers and other market players (eg professional services suppliers) viewed AWS, Microsoft and Google as having broadly equivalent offerings, in terms of products, features and prices.[531] 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.[532]

(b) Other customers said that any innovations in one cloud provider's offerings are matched quickly by the others.[533]

(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.[534]

3.330 Where customers did mention differences in the offering of public cloud providers, they said that these are currently relatively minor or cover edge cases.[535]

3.331 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

[526] Notes of meetings with [✂].
[527] Responses to the CMA's information requests [✂].
[528] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 4.2.1
[529] Responses to the CMA's information requests [✂]; Note of meeting with [✂].
[530] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024).
[531] Responses to the CMA's information requests [✂]; Notes of meetings with [✂].
[532] [✂] response to the CMA's information request [✂].
[533] Note of meeting with [✂]; [✂] response to the CMA's information request [✂].
[534] Note of meeting with [✂].
[535] Notes of meetings with [✂]; Responses to the CMA's information requests [✂].

cloud providers, including the proportion of spend that these services accounted for.[536]

(a) This evidence gives an indication of the share of existing demand for AWS' and Microsoft's 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.

(b) Based on our analysis, there was significant variation in customer responses to this question, with the proportion of AWS' and Microsoft's 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'.[537]

Our analysis

3.332 AWS, Microsoft and Google provided customer data sets that identified customer names and annual spend on their respective clouds from 2020 to 2023. We have matched these data sets to identify customers that are likely to have fully switched from one of these three providers to another one of these three providers.

3.333 In our baseline results, we identify a customer as having switched if:

(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.

3.334 The data, methodology, results and limitations of our analysis are set out 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. Partial switching would be impossible to distinguish from natural fluctuations of customers' spend on cloud.

(b) Our analysis is limited to the three largest providers, AWS, Microsoft, Google. Therefore, we are not able to identify customers who switched to or from other smaller providers. Nonetheless, these three providers cover the vast majority of the market, so we consider the amount of switching not captured due to this limitation to be limited.

---

[536] This evidence is described more in detail in the CSA chapter.
[537] Responses to the CMA's information requests [✂].

(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 have not. 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.[538]

(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.[539]

3.335 Overall, our analysis suggests that regardless of whether they are weighted by revenue or not, customers that do fully switch are relatively small, mostly with a spend around $1,000 and $5,000.

Prevalence of multi-cloud

3.336 This section sets out the evidence on the prevalence of multi-cloud use among customers. We consider the following:

(a) cloud providers' views,

(b) AWS' quantitative analysis;

(c) publicly available survey data and the results of a survey that Ofcom commissioned during its market study;[540] and

(d) our analysis of customer data provided by cloud providers.

Cloud providers' views

3.337 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 said that using multiple clouds is particularly prevalent among large customers.

3.338 However, one cloud provider's submission that the use of multi-cloud is widespread is not consistent with an internal document from that provider which

[538] For further detail, see Appendix I.
[539] For full details see Appendix H.
[540] For more information on the research that Ofcom commissioned, see Cloud Services Market Research - Summary of Findings March 2023.

states, in the context of launching a new product, that it was not at the beginning of a multi-cloud support strategy.[541] The document said that key decision-makers in private and public companies tended to choose one single provider. In particular, the document states that, in relation to [✂]. This suggests that even large customers tend to use one main provider.

3.339 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 provider and/or to run different workloads separately in different clouds.[542]

3.340 As explained in Chapter 2, we consider 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 customers weigh the different benefits and disadvantages of adopting multi-cloud setups. However, as mentioned above, low levels of multi-cloud could be consistent with a lack of ability or incentive for customers to switch or multi-cloud, and therefore with weaker competition, where other evidence also corroborates that view.

3.341 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.[543]

(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.[544]

3.342 We consider that both of these metrics can provide an indication of the prevalence of multi-cloud in the market. While neither can reveal anything about the level of integration a customer may be running, both can identify whether multi-cloud is occurring or not.

3.343 Overall, cloud providers generally submitted that there is a high prevalence of multi-cloud in the market, and Microsoft submitted that customers multi-cloud 'when it works for them'.[545] However, we note that in general these submissions provide limited evidence on the prevalence of multi-cloud.

541 [✂] response to the CMA's information request [✂].
542 Microsoft's response to the Issues Statement, paragraph 40.
543 AWS' submission to CMA [✂].
544 Google's response to the CMA's information request [✂].
545 Microsoft's response to the CMA's Competitive landscape, Committed spend agreements and Egress fees working papers, paragraph 3.

3.344 To this point, [✄] and [✄] highlighted surveys that sought to estimate prevalence of multi-cloud, and we consider 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 one provider

3.345 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 provider participated in from 2009 to 2021 were issued by existing customers. The provider also said that its win rate in tenders for customers with existing workloads in its cloud is [✄] its win rate for other customers. It said that this shows it does not enjoy a significant advantage as an incumbent cloud provider.[546]

3.346 The same 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. [547] This analysis showed that many of its customers [✄] used the provider after awarding tenders to at least one other cloud provider between 2018 and 2022. The provider submitted that this is an indication that these customers were using multiple clouds.[548]

3.347 The provider said that its analysis likely understates the prevalence of multi-cloud because:[549]

(a) many customers acquire IT services without a tender process. These customers would not be recorded in the opportunity data set;

(b) the 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 provider knows the identity of the other competitor that won the tender; and

(d) some customers may have awarded tenders before or after the sample period.

3.348 We consider that the analysis has several conceptual and technical limitations, including sample selection bias, customer inertia/lack of credible options to tender to, limited market coverage. Additionally, tenders are rarely used outside of the public sector, so this analysis is only concerned with the behaviour of a small subset of customers. Therefore, we consider this analysis needs to be interpreted

546 [✄] submission to the CMA [✄].
547 [✄]. [✄] response to Ofcom's information request [✄].
548 [✄] response to the CMA's information request [✄].
549 [✄] submission to the CMA [✄].

with care in light of these caveats.[550] For these reasons, we place less weight on this analysis than our own assessment, set out below.

Quantitative surveys on the prevalence of multi-cloud

3.349 AWS and Microsoft submitted that independent surveys and industry reports show that using multiple clouds is common.[551]

3.350 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.

3.351 Further, even if we did not have any reservations about the probative value of the survey evidence relied on by AWS and Microsoft, we do not consider this evidence to be fully informative of the extent of multi-cloud among customers. This is because the survey evidence provides no information on the split of workloads deployed on different public clouds by customers that multi-cloud. For example, it is possible that customers that multi-cloud do so for only a small proportion of workloads. As a result, the survey might overestimate the prevalence of multi-cloud in the market.

3.352 The results of the survey on the prevalence of multi-cloud at a customer level are broadly consistent with the results of our own analysis of multi-cloud (described below). However, our analysis goes one step further. Our analysis complements customer-level findings with a breakdown of the customers spend across providers. As explained below, this shows that when multi-cloud occurs, most of the spend tends to concentrate on one provider only.

Customer views

3.353 The Jigsaw report notes that, amongst those customers that participated, single-cloud and siloed multi-cloud are the main operating models. Other operating models were found to be less common. [552]

3.354 Only a few participants were described as having models of multi-cloud that were not siloed:[553]

(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 provider failure eg through use of ‘open-source’

[550] We explore these limitations further in Appendix I.
[551] AWS’ submission to the CMA [✂]; Microsoft's submission to the CMA [✂].
[552] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 3.5.2.
[553] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 3.5.5.

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 BiqQuery from Google or Microsoft IAM alongside infrastructure from a different public cloud provider as these services were superior to those offered by their main 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 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 provider availability and price for uses like machine learning).

3.355 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.[554]

3.356 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.[555]

3.357 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.[556]

Our analysis

3.358 We have built our own estimate of multi-cloud prevalence using customer data from cloud providers. We requested customer data sets from AWS, Microsoft and Google that identified customer names and annual spend on their respective clouds for 2020 - 2023. By analysing these data sets to identify customers using

[554] Responses to the CMA's information requests [✂]; Note of meeting with [✂].
[555] Responses to the CMA's information requests [✂].
[556] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 3.6.4.

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.

3.359 We highlight some key considerations that we account for when interpreting the results of our analysis:[557]

(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 AWS, Microsoft 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 providers too: we assumed that 50% of customers (by cloud spend) not on AWS, Microsoft or Google use multi-cloud. We then added these customers to our prevalence of multi-cloud statistics.

(d) Customers are matched across data sets using an approximation technique called fuzzy matching.[558] This can lead to false positives (two different customers on two different data sets being matched) or false negatives (the same customer on two different data sets not being matched).

3.360 In response to our working papers, cloud providers provided views on our analysis.[559] We have incorporated their submissions in our analysis, where appropriate. Below, we discuss some additional points:

(a) AWS and Microsoft 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, we understand, signal weak competition or bad outcomes. AWS added that not all customers wish to multi-cloud, as it would not be efficient for some of them to do so.[560] In line

[557] The data, methodology and limitations of our analysis are fully considered in Appendix I.
[558] See Appendix I for further details.
[559] AWS' response to the CMA's updated issues statement and working papers, paragraph 16; [✂] submission to the CMA [✂]; Microsoft response to the Competitive landscape, Committed spend agreements and Egress fees working papers, paragraph 40.
[560] AWS' submission to the CMA [✂]; Microsoft response to the Competitive landscape, Committed spend agreements and Egress fees working papers, paragraph 41.

with our approach to market outcomes (see paragraph 3.225 above), 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 barriers to multi-cloud. In our view, such an approach would be one of spurious precision. Instead, we consider the evidence set out above on the prevalence of multi-cloud within the broader context of the market, including evidence from customers expressing concerns with specific barriers to multi-cloud.[561]

(b) Microsoft submitted that industry studies confirm that multi-cloud is the new norm when appropriate. It also said that the working papers suggest only one kind of integrated multi-cloud shows a well-functioning market, even though customers disagree.[562]

(c) AWS submitted revised figures for prevalence of multi-cloud based on our analysis.[563] However, these estimates are based on a different market definition than the one adopted by us (see section market definition). This has the effect of leading AWS' revised analysis overestimating the prevalence of multi-cloud.

(d) 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.[564] 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 time (eg basing conclusions off two customers in 2019 might not reliably predict the behaviour of 20 customers in 2025).

3.361 Table 3.10 below 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 3.10: prevalence of multi-cloud, unweighted and weighted by spend, 2020-2023

Prevalence of multi-cloud	*2020*	*2021*	*2022*	*2023*
Unweighted (%)	7.04%	7.50%	7.48%	7.50%
Weighted by revenue (%)	34.80%	37.41%	38.07%	38.47%

Source: CMA analysis of customer data provided by AWS, Microsoft and Google

3.362 This data suggests that multi-cloud is more common among larger customers, as suggested by some cloud providers.

561 See, further, Market outcomes above and Chapter 8.
562 Microsoft Response to the Competitive landscape, Committed spend agreements and Egress fees working papers, paragraph 17.
563 AWS' submission to the CMA Economic response to the CMA's Competitive landscape working paper, paragraph 56.
564 AWS' submission to the CMA Economic response to the CMA's Competitive Landscape working paper, paragraph 54.

3.363 To explore this further, we looked at the prevalence of multi-cloud across customer spend bands: 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.

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

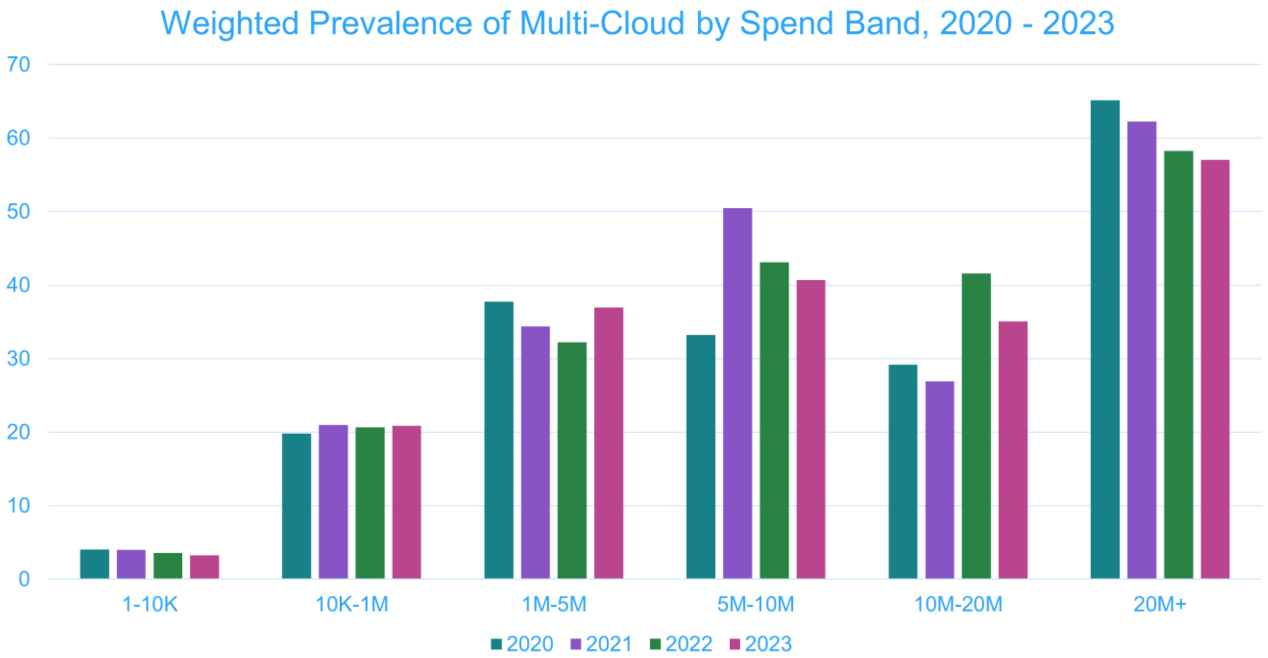

Source: CMA analysis of customer data provided by AWS, Microsoft and Google

3.364 Our results do not show the level of integration of customers' multi-clouds. In Appendix I Prevalence of multi-cloud, we looked at the customers' average spend split across providers. However, we are cautious at interpreting more even splits (eg closer to 50/50 spend across two providers) as suggesting a higher level of integration as this is highly dependent on the type of architecture adopted by customers.

3.365 Nonetheless, we observe how customers typically allocate their spend among providers when they use multi-cloud: this is shown in Table 3.11 below. The table indicates that for customers spending more than $10,000 a year on cloud, which is the group of customers that is most likely to multi-cloud, their spend is concentrated on one primary provider.

Table 3.11: 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	66.9	67.1	66.8	66.5
10K – 1M	82.3	82.6	82.7	83.0
1M – 5M	88.9	90.9	91.3	90.5
5M – 10M	86.7	89.9	91.4	92.4

10M – 20M	94.3	82.0	82.7	83.7
Over 20M	79.3	85.4	87.1	86.1

Source: our analysis of customer data provided by AWS, Microsoft and Google

3.366 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 overall benefits of multi-cloud being greater than the disadvantages.

(b) Multi-cloud is not uncommon among larger customers – this might be consistent with larger customers benefitting more from multi-cloud and/or the barriers they face from doing so to be 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) Among customers that multi-cloud, customers often have one primary cloud provider accounting for the significant majority of expenditure.

3.367 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, consideration of multi-cloud is likely to become a more important factor in customers' cloud strategies. If so, barriers to multi-cloud would become even more important to the competitive process.

Provisional conclusions

3.368 We have looked at market outcomes including profitability, prices, quality and innovation. These outcomes of the competitive process may provide evidence about the functioning of the market.[565]

3.369 In relation to profitability, we have provisionally found that the ROCE for AWS and for Microsoft Cloud & Enterprise has consistently been substantially above our estimated WACC for the last nine and seven years respectively. Microsoft Azure ROCE has been substantially above our estimated WACC for the last four years and is trending upwards.

3.370 We have provisionally found that the ROCE for AWS and for Microsoft Cloud & Enterprise has been falling slightly since 2021. However, we think that these recent trends are, in large part, a result of increased levels of investment in cloud infrastructure aimed at supporting the development of AI services that would represent an upfront investment ahead of developments, that could be expected to generate returns in the future.

[565] CC3 (Revised), paragraph 103.

3.371 Recent trends do not indicate that ROCE for AWS and Microsoft's cloud services business will decline to a 'normal' competitive level, that is the cost of capital.

3.372 Our provisional view is that AWS and Microsoft have been generating sustained returns from their cloud services above their cost of capital, and we consider that this is likely to continue in the future.

3.373 In relation to prices, our provisional view is that prices for different cloud services have moved in different directions, with some services and products increasing in price over time, while others are falling.

3.374 In relation to quality, our provisional view is that, while there is evidence that providers have invested in improvements in quality and innovation, it is not possible to determine the extent to which that is due to competition or other factors.

3.375 Our findings on profitability are consistent with a finding that market outcomes in terms of prices, quality and innovation could be better in a more competitive market. However, they do not, on their own, provide conclusive evidence that the market could be more competitive, and we consider them alongside other findings when we assess the features of the market that may be harming competition.[566]

3.376 The prevalence of switching and multi-cloud is helpful in understanding the balance between benefits and costs of switching and multi-cloud, and the current levels of interest in multi-cloud and switching.[567]

3.377 We have found that full switching is extremely rare in the market. 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 largest providers to be consistent with the presence of high barriers to switching. This is reflected by customers' views. This indicates that switching costs outweigh the benefits of changing provider for many customers.

3.378 The barriers to multi-cloud are not so high that it is prevented to the same degree as switching and we have found that multi-cloud is used by many larger customers. But its overall prevalence indicates that some barriers exist and that, in particular, customers' ability to integrate workloads on more than one cloud is subject to barriers. The barriers may also be greater for smaller customers.

[566] CC3 (Revised), paragraph 126
[567] 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 and discussed further in Chapter 8, we consider that such an approach is neither necessary nor helpful.

3.379 We consider the prevalence of switching and multi-cloud alongside evidence on customers' ability and incentive to pursue their preferred cloud strategy later in this report.

The impact of AI on cloud services

3.380 This section considers the potential impact of AI on how competition works in cloud services.[568] This is important because it is primarily cloud services that provide the computing resources and infrastructure needed to develop and use AI FMs at scale.[569] Furthermore, cloud services are a key route to market for organisations that develop FMs, which we refer to as 'FM developers' (or 'AI labs').

3.381 In particular, this section focuses on whether and how the competitive conditions in IaaS based on accelerated compute could affect the market for IaaS based on standard compute and whether the rising importance of access to FMs could change the way customers of cloud services adopt multi-cloud strategies.

3.382 As discussed earlier in this chapter, providers' estimated revenue from IaaS and PaaS in the UK has grown substantially over recent years.[570] Both the growing demand for accelerated compute from FM developers and demand for access to FMs from other cloud customers is contributing to this revenue growth in cloud services.

3.383 Cloud providers' internal documents show that they recognise the importance of AI to their recent and future global growth:

(a) In a November 2023 update to Amazon's board, AWS stated that 'given the current economic environment, [✂].[571] In a planning document for 2025, AWS forecast that its Machine Learning (ML) infrastructure and GenAI services would deliver revenue of [✂].[572]

(b) In a March 2024 memo to its Board of Directors, Microsoft highlighted that [✂].[573]

(c) In a 2023 Google internal document, Google stated that AI is a [✂]. It predicted [✂] growth of Google Cloud AI revenue between 2023 and 2024 [✂].[574]

568 Updated issues statement, paragraph 30.
569 AI has a wide variety of forms and applications. Here, our use of the term generally refers to foundation models, which typically require substantial specialist compute resources. Foundation models (including, among others, language models or LLMs) are a type of AI technology that are trained on vast amounts of data that can be adapted to a wide range of tasks and operations. See AI Foundation Models: Initial Report.
570 See Nature of competition section.
571 AWS' response to the CMA's information request [✂].
572 AWS' response to the CMA's information request [✂].
573 Microsoft's response to the CMA's information request [✂].
574 Google's response to the CMA's information request [✂].

(d) In June 2024, Oracle's CEO announced that in Q3 and Q4 of FY24 Oracle signed the largest sales contracts in its history 'driven by enormous demand for training AI large language models in the Oracle Cloud. These record level sales drove [Oracle's Remaining Performance Obligation] up 44% to $98 billion'. Throughout FY25, Oracle's CEO expected 'continued strong AI demand to push Oracle sales and [Remaining Performance Obligation] even higher—and result in double-digit revenue growth this fiscal year'.[575]

3.384 We have assessed two aspects of how AI could affect competition in cloud services:

(a) the supply of accelerated compute by cloud providers to FM developers; and

(b) the supply of access to FMs by cloud providers to their customers.

3.385 As background to our assessment, we have set out:

(a) what accelerated compute is and why cloud providers supply accelerated compute to FM developers;

(b) the factors involved in supplying accelerated compute, such as accessing AI accelerator chips and establishing supply agreements with FM developers; and

(c) how cloud providers offer access to FMs to customers of their public cloud infrastructure.

Accelerated compute and its importance to FM development

3.386 Generally, it is not feasible to develop FMs using conventional computer chips, such as Central Processing Units (CPUs), due to the large size of FMs and the amount of training data required. Instead, specialised computer chips known as accelerator chips are typically used for FM development. Once trained, FMs can 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 accelerator chips.[576] Figure 3.2 below illustrates different levels of the value chain and the role of compute.

3.387 Different types of accelerator chips include Graphical Processing Units (GPUs) and Application-Specific Integrated Circuits (ASICs).[577] We refer to compute that involves AI accelerator chips as accelerated compute.

[575] Oracle's earnings conference for Q4 of FY24, June 11th 2024, Oracle Announces Fiscal 2024 Fourth Quarter and Fiscal Full Year Financial Results, accessed 25 August 2024.
[576] AI Foundation Models: Initial Report, paragraph 2.18.
[577] AI Foundation Models: Initial Report, paragraph 2.14.

3.388 Most FM developers rent the accelerated compute that they need from cloud providers rather than build it themselves, due to the large up-front cost and complexity involved in building accelerate compute capacity and the pace at which the compute requirements of FMs change.[578] Other options for organisations that train FMs are to invest in building their own infrastructure (on-premises) or to access a public supercomputer.[579]

Figure 3.2**: Illustration of the value chain for FM development and deployment**

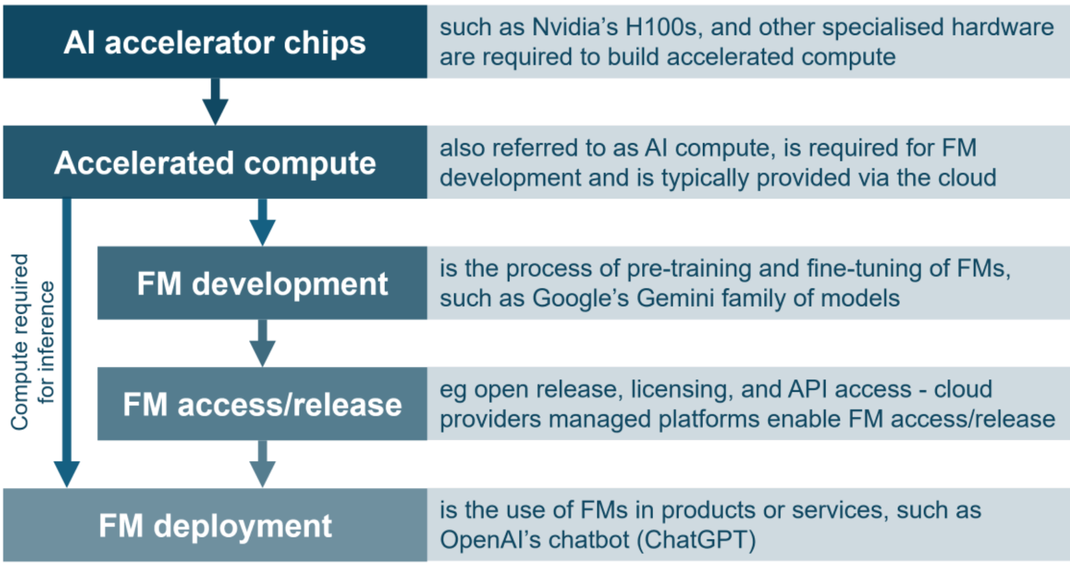

Source: CMA

3.389 We have identified at least 13 FM developers globally that procure accelerated compute from cloud providers.[580] Table 3.12 below maps the commercial

[578] Note of meeting with [✂].
[579] One provider gave two examples of Dell collaborations to offer on-premises solutions for FM deployment (with Hugging Face and Meta, respectively). [✂] response to the CMA's information request [✂]. Meta is an example of a large technology company that has opted to build its own in-house AI infrastructure and reduce its use of cloud computing for these purposes; Note of meeting with [✂].
[580] Unlocking Innovation: AWS and Anthropic push the boundaries of generative AI together - AWS Machine Learning Blog, accessed 30 September 2024; Meta/Facebook turns to AWS as 'long-term strategic cloud provider' for acquisitions, third-party collaborations and AI - DCD, accessed 30 September 2024.; AI21 Labs Accelerates Generative AI Model Adoption Using Amazon SageMaker, accessed 19 April 2024; Amazon will invest up to $4B in Anthropic to advance generative AI, accessed 19 April 2024; Runway News - Scaling our in-house research infrastructure with AWS , accessed 30 September 2024; Stability AI builds foundation models on Amazon SageMaker, accessed 30th September. Meta selects Azure as strategic cloud provider to advance AI innovation and deepen PyTorch collaboration, accessed 30 September 2024; Introducing Mistral-Large on Azure in partnership with Mistral AI, accessed 19 April 2024; [✂] response to the CMA's information request [✂]; OpenAI and Microsoft extend partnership. With a systems approach to chips, Microsoft aims to tailor everything 'from silicon to service' to meet AI demand, accessed 30 September 2024. What is Azure OpenAI Service? accessed 19 April 2024; Generative AI startups choose Google Cloud, accessed 19 April 2024; Anthropic Partners with Google Cloud, accessed 19 April 2024; Accelerating language model training with Cohere and Google Cloud TPUs, accessed 30 September 2024; AI21 Labs Case Study, accessed 30 September 2024; Character.AI and Google Cloud Partner to Build the Next Generation of Conversational AI - May 10, 2023 , accessed 19 April 2024; Midjourney Selects Google Cloud to Power AI-Generated Creative Platform, accessed 19 April 2024; Runway to Make Content Creation More Accessible with Google Cloud's Generative AI, accessed 3 December 2024; AI Startups Find an Unlikely Friend: Oracle, accessed 19 April 2024; Oracle to Deliver Powerful and Secure Generative AI Services for Business, accessed 19 April 2024; MosaicML Trains Generative AI Models Faster with Oracle, accessed 19 April 2024; CoreWeave Partners with EleutherAI & NovelAI to Make Open-Source AI More Accessible, accessed 19 April 2024; [✂] response to the CMA's information request [✂]; Microsoft inks deal with CoreWeave to meet OpenAI cloud demand; Amazon.com announces first quarter results, accessed 19 July 2024.

relationships between these and FM developers.[581] We note that several cloud providers are also FM developers and deployers, including AWS, Microsoft and Google: these firms are vertically integrated across the FM value chain, from the supply of accelerated compute to FM development and deployment.[582]

3.390 AWS and Microsoft internal documents show that the provision of accelerated compute to FM developers is becoming an increasingly sizeable source of revenue for their public cloud infrastructure businesses, although it remains overall small as a proportion of total revenue.[583]

Table 3.12: FM developers as customers of public cloud services.

	AWS	***Microsoft***	***Google***	***Oracle***	***CoreWeave***
AI21 Labs	**Yes**		**Yes**		
Adept AI		**Yes**		**Yes**	
Anthropic	**Yes**		**Yes**		
Character AI			**Yes**	**Yes**	
Cohere			**Yes**	**Yes**	
EleutherAI					**Yes**
Meta	**Yes**	**Yes**			
Midjourney			**Yes**		
Mistral		**Yes**			
Mosaic ML				**Yes**	
OpenAI		**Yes**			**Yes**
Runway ML	**Yes**		**Yes**		
Stability AI	**Yes**				

Source: Unlocking Innovation: AWS and Anthropic push the boundaries of generative AI together | AWS Machine Learning Blog (amazon.com), accessed 20 September 2024. Meta/Facebook turns to AWS as 'long-term strategic cloud provider' for acquisitions, third-party collaborations and AI - DCD (datacenterdynamics.com), accessed 20 September 2024; AI21 Labs Accelerates Generative AI Model Adoption Using Amazon SageMaker | Case Study | AWS, accessed 19 April 2024; Amazon will invest up to $4B in Anthropic to advance generative AI (aboutamazon.com), accessed 19 April 2024; Runway News | Scaling our in-house research infrastructure with AWS (runwayml.com), accessed 30 September 2024; Stability AI builds foundation models on Amazon SageMaker; | AWS Machine Learning Blog, accessed 30 September 2024. Meta selects Azure as strategic cloud provider to advance AI innovation and deepen PyTorch collaboration | Microsoft Azure Blog, accessed 30 September 2024; Introducing Mistral-Large on Azure in partnership with Mistral AI | Microsoft Azure Blog, accessed 19 April 2024; [✂] response to the CMA's information request [✂]. OpenAI and Microsoft extend partnership. With a systems approach to chips, Microsoft aims to tailor everything 'from silicon to service' to meet AI demand - Source, accessed 30 September 2024;. What is Azure OpenAI Service? - Azure AI services | Microsoft Learn, accessed 19 April 2024. Generative AI startups choose Google Cloud | Google Cloud Blog, accessed 19 April 2024;; Anthropic Partners with Google Cloud \ Anthropic, accessed 19 April 2024; Accelerating language model training with Cohere and Google Cloud TPUs | Google Cloud Blog, accessed 30 September 2024 , AI21 Labs Case Study | Google Cloud, accessed 30 September 2024; Character.AI and Google Cloud Partner to Build the Next Generation of Conversational AI - May 10, 2023 (googlecloudpresscorner.com) , accessed 19 April 2024;; Midjourney Selects Google Cloud to Power AI-Generated Creative Platform (prnewswire.com), accessed 19 April 2024; Runway to Make Content Creation More Accessible with Google Cloud's Generative AI, accessed 19 April 2024; AI Startups Find an Unlikely Friend: Oracle — The Information, accessed 19 April 2024; AI Startups Find an Unlikely Friend: Oracle — The Information, accessed 19 April 2024; Oracle to Deliver Powerful and Secure Generative AI Services for Business, accessed 3 December 2024; MosaicML Trains Generative AI Models Faster with Oracle, accessed 19 April 2024; CoreWeave Partners with EleutherAI & NovelAI to Make Open-Source AI More Accessible — CoreWeave, accessed 19 April 2024; [✂] response to the CMA's information request [✂]. Microsoft inks deal with CoreWeave to meet OpenAI cloud demand (cnbc.com); AMZN-2024.03.31-EX99.1 (q4cdn.com), accessed 19 July 2024.

[581] This list is illustrative rather than exhaustive. Where a customer is associated with more than one cloud provider, in some cases this is due to switching over time, in other cases it is due to a multi-cloud approach.

[582] AI Foundation Models: initial review. See the Update Paper published on 11 April 2024 and the Technical Update Report published on 16 April 2024.

[583] Microsoft's CFO delivered an update to the Board in September 2023, noting that they had guided the Azure [✂] to grow by [✂] year-on-year in constant currency and [✂]. Microsoft's response to the CMA's information request [✂]. In a November 2023 update to Amazon's Board of Directors, AWS noted that '[✂].' AWS anticipated [✂]. AWS' response to the CMA's information request [✂]. In comparison, another AWS internal document from 2023 outlines its expectation [✂] (This suggests that [✂] may account for approximately [✂] of revenue increase between 2024 and 2025). AWS' response to the CMA's information request [✂].

How cloud providers supply accelerated compute to FM developers

3.391 The provision of accelerated compute requires access to large numbers of AI accelerator chips. These chips are networked together in data centres or supercomputers such that they can be used in parallel to perform the large computations involved in FM development.[584]

3.392 This section outlines the sources of supply for AI accelerator chips and summarises how cloud providers supply FM developers with accelerated compute.

Sources of supply for the AI accelerator chips

3.393 AI accelerator chips are in high demand because they are the most efficient for FM developers to use.[585] A cloud provider must therefore be able to acquire AI accelerator chips in sufficient quantities.

3.394 The main sources of supply of AI accelerator chips are[586]:

(a) established chip providers, namely Nvidia (the market leader in data centre GPUs), AMD and Intel;[587] and

(b) self-supply (or 'custom silicon') ie cloud providers developing their own ASICs.[588]

3.395 In some cases, cloud providers access additional accelerated compute capacity by procuring it from other providers of accelerated compute. For example, as discussed below, Microsoft has previously procured accelerated compute from specialised providers CoreWeave or Lambda Labs.

3.396 Pursuing a custom silicon strategy enables cloud providers to optimise the chip architecture for performance of its own (or partners') FMs, thereby enhancing computational performance.

---

[584] While all the hardware and software components comprising AI compute – and the way they are configured – are important in influencing computational performance, AI accelerator chips are a particularly key component (this was noted, for example, see page 4 Amazon.com Q3 2023 Earnings Call). Developing (and in some cases, deploying) FMs typically requires thousands of accelerator chips working together – for example, it has been reported that OpenAI used a system of 25,000 GPUs to train its GPT models Analysts positive on Nvidia's ChatGPT, AI opportunity By Investing.com, accessed 20 May 2024.

[585] Note of meeting with [✂]; [✂] response to the CMA's information request [✂]; It has been estimated that building an AI model on non-SOTA AI chips would be at least 33 times more expensive than using SOTA AI chips. Computational Power and AI - AI Now Institute, accessed 30 September 2024.

[586] Computational Power and AI - AI Now Institute, accessed 30 September 2024.

[587] 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.

[588] 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.

3.397 Cloud providers can also custom design other hardware components necessary to build a data centre or supercomputer, as well as the software used to programme the chips. This custom approach enables systems-level optimisation and hence lowering of costs.[589]

3.398 Self-supply reduces reliance on third-party suppliers, and this may be especially important in the current context of data centre GPU supply constraints.

3.399 Examples of cloud providers developing their own AI accelerator chips for self-supply include:[590]

(a) Google's series of Tensor Processing Units (TPUs) designed for machine learning workloads including FM development and deployment.[591] 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,[592] Anthropic,[593] and Midjourney.[594]

(b) Amazon's 'Trainium' and 'Inferentia' series of AI accelerator chips, which are only available to AWS customers. Customers currently include Anthropic and Snap.[595]

(c) Microsoft announced its first custom AI accelerator chip, the 'Maia 100', in November 2023.[596] In 2024 the Maia 100 will power Microsoft's services such as Copilot and Azure OpenAI Service.[597] [✂].[598]

3.400 Cloud providers also require access to additional hardware and infrastructure to build data centres, including memory chips, networking technologies, cooling systems and power supplies in sufficient volume.[599] They also need software to

---

589 With a systems approach to chips, Microsoft aims to tailor everything 'from silicon to service' to meet AI demand - Source, accessed 30 September 2024; Google AI Infrastructure Supremacy: Systems Matter More Than Microarchitecture (semianalysis.com), accessed 30 September 2024; Announcing Cloud TPU v5e and A3 GPUs in GA | Google Cloud Blog, accessed 30 September 2024.

590 Oracle said that it has [✂]. Oracle's response to the CMA's information requests [✂]. IBM said that its research division has invested in several prototype chip projects, but that IBM Cloud does not currently offer access to IBM-build accelerators to customers. IBM's response to the CMA's information requests [✂].

591 Tensor Processing Units (TPUs), accessed 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 requests [✂].

592 Building the most open and innovative AI ecosystem, accessed 20 September 2024.

593 Cloud TPU v5e is generally available, accessed 20 September 2024.

594 Building the most open and innovative AI ecosystem, accessed 20 September 2024.

595 AWS AI chips powering Amazon's partnership with Anthropic, accessed 30 September 2024. AI Chip - AWS Inferentia , accessed 30 September 2024. In its [✂] report for 2024, AWS stated that: '[✂].' AWS's response to the CMA's information request [✂].

596 Microsoft's response to the CMA's information request [✂]; Microsoft Azure delivers purpose-built cloud infrastructure in the era of AI, accessed 30 September 2024.

597 Microsoft's response to the CMA's information request [✂]; Microsoft Azure delivers purpose-built cloud infrastructure in the era of AI, accessed 30 September 2024.

598 [✂] response to the CMA's information requests [✂].

599 Computing Power and the Governance of Artificial Intelligence , accessed 30 September 2024.

programme the AI accelerator chips and integrate them with other cloud services and applications.[600]

Cloud providers' provision of accelerated compute to FM developers

3.401 AWS, Microsoft, Google, Oracle and IBM all supply FM developers with accelerated compute via their cloud platforms, alongside their standard, non-accelerated compute. As with non-accelerated compute, these cloud providers offer a range of software and tools for managing accelerated compute workloads.

3.402 There are also several smaller cloud providers that specialise in the provision of accelerated compute for compute intense workloads such as FM development and deployment (examples include CoreWeave and Lambda Labs).[601] Specialised providers' offerings tend to be targeted at a narrower group of customers and therefore offer a more limited range of cluster management options, including more basic management software or none at all ('bare metal').[602]

3.403 Based on our review of internal documents, we consider that these bare metal offerings are particularly attractive to smaller, sophisticated customers such as AI startups, who may have the expertise to directly manage their workloads and also be keen to cut costs by avoiding paying for additional software.

3.404 Cloud providers typically supply accelerated compute to FM developers in one of three ways:

(a) supplying compute at commercial on-demand rates;

(b) entering into a supply agreement (usually one to three years) offering compute at a reduced rate;[603] or

(c) entering into a commercial 'compute partnership' that involves the FM developer gaining preferential or discounted access to accelerated compute in exchange for supplying the cloud provider with preferential or exclusive access to its FMs (either for the cloud provider to license and use within its own products or services or to distribute via its cloud platform).[604]

3.405 In addition to the supply of compute, FM distribution and product licensing agreements, compute partnerships may also involve profit-sharing agreements and cloud providers gaining voting or consultation rights. Partnerships between

[600] Computational Power and AI - AI Now Institute, accessed 30 September 2024.
[601] Responses to the CMA's information requests [✂].
[602] AI Infrastructure Explained, accessed 30 September 2024.
[603] For example, one cloud provider's internal document compares pricing of H100 instances across GCP, Azure and AWS, based on (i) on-demand pricing, (ii) a 1-year agreement, or (iii) a 3-year agreement. [✂] response to the CMA's information request [✂]. See also EC2 Reserved Instance Pricing, accessed 19 April 2024.
[604] CMA AI Foundation Models: Initial Report, page 35. Also see paragraph 3.465 and its references.

cloud providers and FM developers may also involve agreements to collaborate on the development of cloud providers' AI accelerator chips.[605]

3.406 We consider price competition and compute partnerships further below.

Cloud providers' supply of access to FMs to customers

3.407 FM developers can monetise their FMs by enabling customers to use (or 'deploy') them for business or personal purposes. Customers may access FMs in order to develop tools or services that they then use internally, or in order to integrate those FMs into products and services that they sell to their own customers.

3.408 Access to FMs is typically provided through cloud providers' managed platforms, which enable customers to access various first and third party open and proprietary FMs hosted on the provider's cloud.[606] Access to FMs via a managed platform is often on a metered basis.[607]

3.409 Table 3.13 below shows the managed platforms offered by each cloud provider. Table 3.14 below lists a selection of the FMs that are available on some cloud providers' managed platforms and services.

Table 3.13: Cloud providers' managed platforms

AWS	Microsoft	Google	Oracle	IBM
• Amazon Bedrock • AWS Marketplace • Amazon SageMaker Jumpstart	• Azure AI studio • Azure Machine Learning Studio • Azure Marketplace	• Vertex AI	• Generative AI	• Watson X

Source: Figure 3 of AI Foundation Models technical update report (publishing.service.gov.uk). Responses to the CMA's information requests [✂].

3.410 Cloud providers sometimes enter commercial 'compute partnerships' with FM developers. Examples of such partnerships include Microsoft and OpenAI, Amazon and Anthropic, and Google and Anthropic. Below, we discuss the importance of these partnerships in a cloud provider's strength in supplying customers with access to FMs.[608]

[605] With a systems approach to chips, Microsoft aims to tailor everything 'from silicon to service' to meet AI demand, accessed 30 September 2024; Summary of phase 1 decision.

[606] Responses to the CMA's information requests [✂].
[607] Responses to the CMA's information requests [✂].
[608] This is a non-exhaustive list of commercial compute partnerships. We note that each of these partnerships has been, or is being, examined by the CMA in separate investigations conducted under the UK merger control regime. (See Microsoft - OpenAI partnership merger inquiry; Amazon / Anthropic partnership merger inquiry; Alphabet Inc. (Google

Table 3.14: Non-exhaustive list of FM models available on managed platforms

AWS Bedrock	Microsoft Azure Machine Learning	Google Vertex AI
• Anthropic's Claude family of FMs • Amazon's Titan family of FMs • Stability AI's Stable Diffusion models • Meta's Llama 2 • Mistral's 7B and Mixtral 8x7B • Cohere's Command & Embed	• Exclusive access to OpenAI models including GPT-3.5, GPT-3.5 Turbo, GPT-4, GPT-4 Turbo • Curated models from Meta, Nvidia and Mistral, and open models from HuggingFace. • Microsoft's series of small language models eg Phi, Orca	• API access to Google models, including Gemini 1.0, PaLM 2, Codey, Imagen • The Model Garden provides access to 130+ FMs including Anthropic's Claude family of FMs and open source models Google's Gemma and Meta's Llama 2

Source: Figure 3 of AI Foundation Models technical update report (publishing.service.gov.uk)

3.411 When providing cloud customers with access to FMs, there are two main access routes:

(a) API access: the third party can send prompts to an FM via an application programming interface (API) and receive a response.

(b) Model access: the third party can obtain a license to use the FM on their own systems, preventing the need for data sharing with the FM provider.

3.412 Cloud providers have also explained how the billing relationship works between customers and cloud providers and between cloud providers and the FM developers whose models they list in their managed platforms:

(a) One cloud provider stated that when FM developers make their FMs available on its managed platform [✂], after receiving sales proceeds from the customer, the provider deducts the cost of the computing infrastructure and the operation of its managed service and remits the rest of the proceeds to the third- party FM developer in accordance with the cloud provider's commercial arrangements with FM providers.[609]

(b) Another cloud provider stated that when customers access FMs through its managed platform, they typically pay for the latter on a pay-as-you-go basis (ie per API call). It added that for open access FMs, customers are only charged for the compute resources required for their use.[610]

LLC) / Anthropic merger inquiry). Additionally, the CMA has examined the partnership between Microsoft and Mistral (See Microsoft / Mistral AI partnership merger inquiry) and the arrangement between Microsoft and Inflection (See Microsoft / Inflection inquiry).

[609] [✂] response to the CMA's information request [✂].

[610] [✂] response to the CMA's information request [✂].

The supply of accelerated compute by cloud providers to FM developers

3.413 We have examined the relative position and strength of AWS, Microsoft and Google, as well as other providers of accelerated compute, including Oracle, IBM, CoreWeave and Lambda Labs.

3.414 The cloud providers' accelerated compute capacity (measured in TFLOPS)[611] shows that AWS, Microsoft and Google have approximately [✂] the amount of accelerated compute capacity than CoreWeave which in turn has approximately [✂] the accelerated compute capacity of Oracle and Lambda Labs. IBM had approximately [✂] orders of magnitude less accelerated compute than the other providers. [612]

3.415 In order to understand how cloud providers' accelerated compute capacity was allocated to FM development, we asked cloud providers the proportion of their total global AI accelerator chip capacity that was allocated to (i) the development and provision of first-party FMs, (ii) the development or provision of any FMs by FM developers/providers they were in partnership with, (iii) the development or use provision of any FMs by previous partners and (iv) other use cases. The distribution of accelerated compute capacity for first-party versus partners' FM development and provision varies between AWS, Microsoft and Google.

3.416 According to cloud providers' estimates, in 2023:[613]

(a) Approximately [✂]% of AWS' Amazon EC2 'Accelerated Computing' instance capacity was procured by all Amazon-wide use of such instances (including the development and provision of its own FMs) and external customer usage of AWS and Amazon services, such as Bedrock. [✂]. The remaining [✂] was procured by all other customers.[614]

(b) Only [✂]% of Microsoft's accelerated compute capacity was used for the development and provision of its own FMs, whereas [✂]% was used for the development (training) of its partners' FMs (we understand [✂] to be included in this category). The remaining [✂]% was used for inferencing, [✂] being of Microsoft's partners' FMs.[615]

(c) Approximately [✂]% of Google's estimated accelerated compute capacity was used for development and provision of its own FMs (including DeepMind's workloads, such as the development of first party FMs, as well as inference of those FMs by Google Cloud customers via Vertex AI) and

[611] Teraflops (TFLOPS) are used to measure computational performance. One TFLOPS is equal to one trillion floating point operations per second. Teraflop Definition & Meaning.
[612] Responses to the CMA's information requests [✂].
[613] We note that the cloud providers may have varying approaches in how they made these estimates. Regardless of this potential inconsistency, the statistics are illustrative of a trend.
[614] [✂] response to the CMA's information request [✂]. [✂].
[615] [✂] response to the CMA's information request [✂].

that approximately another [✂]% was used for the development and provision of FMs developed by its partners. The remaining approximately [✂]% was used for other purposes, including other areas of the Alphabet Group.[616]

Specialised providers

3.417 There are several specialist cloud providers offering solely (or mainly) accelerated compute:[617]

(a) CoreWeave was established in 2017 as a company focused on cryptocurrency applications.[618] Today, CoreWeave describes itself as a cloud provider that specialises in GPU-accelerated workloads.[619] The startup was reportedly valued at $19bn in May 2024.[620]

(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.[621] It was reportedly valued at $1.5bn in February 2024.[622]

3.418 Microsoft submitted that the rapid increase in demand for accelerated compute had accelerated entry and expansion of these suppliers of specialised cloud infrastructure services focused on large-scale GPU-accelerated workloads.[623] Microsoft also submitted that, in order to access additional GPU capacity for AI workloads, it has agreements with smaller cloud providers to supply Microsoft with Nvidia GPU capacity – these include an agreement signed in February 2023 with CoreWeave,[624] [✂].[625]

3.419 Microsoft's internal documents demonstrate that Microsoft considers its principal competitors in the supply of accelerated compute to be [✂], rather than CoreWeave or Lambda Labs:

---

616 [✂] response to the CMA's information request [✂].
617 AWS said it is one of numerous options available to customers to obtain the compute capacity they need and identified CoreWeave, Lambda Labs, Denvr Dataworks, G42, Omniva, Crusoe, Cirrascale, TensorWave and Runpod, as newer entrants responding to the opportunities created by generative AI [specialist IT providers offering solely (or mainly) accelerated compute]. AWS' response to the CMA's information request [✂]. AWS' response to the CMA's information request [✂]; Microsoft identified CoreWeave, Lambda Labs, Scaleway, Hewlett Packard Enterprise, Together AI, Nvidia (its DGX Cloud platform), LaminiAI, Paperspace, Crusoe Cloud, Cirrascale, Clever Cloud, Denvr Dataworks, FluidStack, Runpod, Oblivus, Vast.ai, Tensordock, TensorWave, Modal and OctoAI. Microsoft's response to the CMA's information request [✂].
618 CoreWeave came 'out of nowhere.' Now it's poised to make billions off AI with its GPU cloud, accessed 25 August 2024.
619 CoreWeave's response to the CMA's information request [✂].
620 CoreWeave Raises $7.5B In Debt For AI Push, accessed 20 September 2024.
621 About Lambda, accessed 20 September 2024.
622 AI Compute Startup Lambda Hits $1.5B Valuation After Massive $320M Raised, accessed 20 September 2024.
623 Microsoft's responses to the CMA's information requests [✂].
624 Microsoft's response to the CMA's information request [✂].
625 Microsoft's response to the CMA's information request [✂].

(a) In instances when Microsoft [✄].[626]

(b) Oracle, CoreWeave and Lambda Labs are described in Microsoft internal documents as [✄].[627]

3.420 One cloud provider has 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.[628] It has been reported that CoreWeave and Lambda Labs were two of the first cloud providers to get general access to flagship Nvidia's H100 chip when it launched in 2023.[629] CoreWeave's and Lambda Labs' accelerated compute capacity lags substantially behind that of the main providers.

Cloud providers' supply of proprietary accelerator chips

3.421 AWS, Microsoft and Google already have or are developing proprietary AI accelerator chips for in-house use and for their cloud customers' use. These cloud providers have markedly differentiated offerings in the supply of first-party AI accelerator chips.

3.422 All three 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.[630]

(b) AWS considers that its custom silicon (Trainium and Inferentia chips) are beginning to 'differentiate [its] infrastructure offering in the GenAI space' [✄].[631]

(c) Microsoft's internal documents note [✄].[632]

Competition on price

3.423 Cloud providers' internal documents show that they also compete for FM developers' workloads on price, with the largest FM developers often benefiting from substantial usage discounts.

[626] Microsoft's response to the CMA's information request [✄].
[627] Microsoft's response to the CMA's information request [✄].
[628] [✄] response to the CMA's information request [✄].
[629] 'What is 'compute', and why does it matter?, Computational Power and AI - AI Now Institute, page 32, accessed 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 25 August 2024. Lambda Cloud Deploys NVIDIA H100 Tensor Core GPUs. See also [✄] response to the CMA's information request [✄].
[630] Google's response to the CMA's information request [✄].
[631] AWS' response to the CMA's information request [✄].
[632] Microsoft's response to the CMA's information request [✄].

(a) In particular, Microsoft stated that other providers '[✂]';[633] and

(b) AWS said [✂] 'highly competitive with other cloud providers (…) [✂].[634]

3.424 In one cloud provider's internal documents, the prices of competitors' accelerated compute offerings and associated discounts were compared – the cloud provider considered when and whether to match the competitors' discounts.[635]

3.425 Another cloud provider discussed in internal documents how it had 'pursued competitive GenAI deals' that decreased the average selling price of its GPU instances by approximately [✂], with further reductions expected by the end of 2025 – the cloud provider stated that it believed 'these average selling price assumptions are reflective of the competitive price pressure seen in recent deals'.[636]

Barriers to entry and expansion

3.426 We have considered whether newer or smaller cloud providers will be able to enter or expand in the supply of accelerated compute to FM developers in such a way that impacts the competition in cloud services.

Cloud provider investments in the supply of accelerated compute

3.427 This section covers cloud providers' capex investments in supplying accelerated compute. Then, we discuss two such types of capex investments in detail: purchasing AI accelerator chips from third parties, and self-supply of AI accelerator chips.

3.428 The supply of AI accelerator chips (such as GPUs) has been unable to keep pace with increasing demand from FM developers over the past two years, resulting in a global shortage.[637]

3.429 Several cloud providers have highlighted that demand for accelerated compute exceeds supply.[638] AWS, Microsoft and Google said that they are making substantial investments to expand their accelerated compute capacity. In particular:

633 Microsoft's response to the CMA's information request [✂].
634 AWS' response to the CMA's information request [✂].
635 [✂] response to the CMA's information request [✂].
636 [✂] response to the CMA's information request [✂].
637 See section 'How cloud providers supply accelerated compute to FM developers'.
638 [✂] response to the CMA's information request [✂].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 30 September 2024.

(a) In a finance review submitted to Microsoft's Board of Directors in July 2023, Microsoft's server GPU capex was forecast to [✂] server CPU CapEx was forecasted to [✂].[639]

(b) In a 2023 strategy document, Google outlined how its GPU and TPU CapEx grew from [✂] in 2022 and was expected to grow again to [✂] in 2023 and [✂] in 2024.[640]

(c) In a November 2023 update to [✂], AWS said it was [✂].[641]

3.430 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.[642] Another cloud provider submitted to the CMA that 'barriers to meeting [FM developers' compute] demands relate primarily to capital required to install and operate the compute'.[643]

3.431 In the context of significant capex investment 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 [✂].[644]

(b) In a November 2023 AWS update to [✂] discussing generative AI, AWS said that [✂], it is also [✂].[645]

3.432 AWS and Microsoft also noted the [✂] margins resulting from the significant capex associated with AI.

(a) In a March 2023 update [✂], AWS noted that [✂].[646]

(b) A September 2023 CFO update to Microsoft's Board said that [✂].[647]

Investments in AI accelerator chips from third parties

3.433 Cloud providers are making investments to increase their accelerated compute capacity including by procuring AI accelerator chips, such as GPUs.[648]

---

[639] Microsoft's response to the CMA's information request [✂].
[640] Google's response to the CMA's information request [✂].
[641] AWS' response to the CMA's information request [✂].
[642] [✂] response to the CMA's information request [✂].
[643] [✂] response to the CMA's information request [✂].
[644] Microsoft's response to the CMA's information request [✂].
[645] AWS' response to the CMA's information request [✂].
[646] AWS' response to the CMA's information request [✂].
[647] Microsoft's response to the CMA's information request [✂].
[648] [✂] response to the CMA's information request [✂].

3.434 Nvidia is the current market leader in the supply of data centre GPUs, which are used for FM development and evidence from cloud providers shows that the majority of their accelerated compute capacity is powered by Nvidia's GPUs.[649]

3.435 AWS and Microsoft have emphasised 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, [✂]. [650] In a July 2023 internal document, AWS noted that 'NVIDIA GPU costs represented about [✂] of its total server cost'. AWS noted that [✂].[651]

(b) Microsoft submitted that it had spent approximately [✂] on GPUs for model training and inference globally in 2023,[652] [✂].[653]

3.436 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'.[654]

Cloud provider self-supply of AI accelerator chips

3.437 AWS, Microsoft and Google are investing in developing their own proprietary AI accelerator chips for internal use and providing to cloud customers.

3.438 AWS and Microsoft have noted the large investments required to develop custom AI accelerator chips. For example:

(a) AWS said that it had invested a total of [✂] 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.[655] AWS said that it [✂].[656] AWS said that it is also working to meet the growing demand for its Trainium and Inferentia chips.[657]

(b) Microsoft said that between July 2021 and June 2024 (including forecast spend), it expects to have spent in total [✂] in developing its custom AI

[649] Responses to the CMA's information requests [✂].
[650] AWS' response to the CMA's information request [✂].
[651] AWS' response to the CMA's information request [✂].
[652] [✂] these GPUs were used for model training.
[653] Microsoft's response to the CMA's information request [✂].
[654] [✂] response to the CMA's information request [✂].
[655] AWS' response to the CMA's information request [✂].
[656] AWS' response to the CMA's information request [✂].
[657] AWS' earnings conference for Q1 of the FY24, April 30th 2024, Amazon.com announces first quarter results, accessed 30 September 2024.

accelerator chip, Maia.[658] Microsoft estimated that AWS, Microsoft and Google had invested between $139 billion and $323 billion 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.[659]

3.439 Self-supply also requires development of the accompanying software (that is used to program the AI accelerator chips), which requires significant investment.[660] As one cloud provider stated [✂], 'the biggest asset for NVIDIA is its superior software environment, which gives developers and scientists the flexibility to innovate and optimise performance'. [✂].[661]

Technical barriers to switching

3.440 We have considered whether there are any particular technical barriers for FM developers to switch between providers of accelerated compute.

3.441 AI workloads on AWS, Azure and Google Cloud Platform are typically managed through abstracted container orchestration services, such as Kubernetes, for which the underlying compute is increasingly shifting towards being serverless. Cloud providers have identified several reasons for why customers may prefer to use managed and serverless services for large-scale AI workloads, including ease of use and reduced costs.[662]

3.442 AWS, Microsoft and Google submitted that the use of open-source container orchestration services, such as Kubernetes, can support switching and the ability to multi-cloud because open-source software enables porting or interoperability between different cloud providers. However, because these open-source containers are being used within cloud providers' managed systems, their use may nevertheless create a lock-in risk for FM developers seeking to access accelerated compute, which would be sustained in future.

Partnerships between cloud providers and FM developers

3.443 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.

658 Microsoft's response to the CMA's information request [✂]. [✂].
659 Microsoft's response to the CMA's information request [✂].
660 Notes of meetings with [✂].
661 [✂] response to the CMA's information request [✂].
662 Responses to the CMA's information requests [✂].

3.444 The Table 3.15 below provides an overview of the current or recent partnerships between FM developers and AWS, Microsoft, Google and Oracle.

Table 3.15: Partnerships between cloud providers and AI developers

AWS	Microsoft	Google	Oracle
• Anthropic	• AdeptAI	• AI21 Labs	• Cohere
• Hugging Face	• Falcon LLM	• Anthropic	• OpenAI
• AI21 Labs	• Hugging Face	• Character.AI	
• Cohere	• Meta	• Cohere	
• Falcon LLM	• Mistral AI	• Contextual AI	
• Meta	• Open AI	• Essential AI	
• Mistral AI	• Perplexity AI	• Hugging Face	
• Runway AI		• Midjourney	
• Stability AI		• Mistral AI	
		• Runway	
		• Stability AI	

Source: AI Foundation Models technical update report (publishing.service.gov.uk) Figure 7; Responses to the CMA's information requests [✂].

3.445 Cloud providers submitted their views on partnerships:

(a) One cloud provider submitted 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.[663]

(b) 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].'[664]

(c) 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.[665] 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)'.[666]

(d) 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

663 [✂] response to the CMA's information request [✂].
664 [✂] response to the CMA's information request [✂].
665 Responses to the CMA's information requests [✂].
666 [✂] response to the CMA's information request [✂].

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'.[667]

(e) 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 procure and provide (due to tight supply and/or required size of investment) the large amount of required compute capacity'.[668]

3.446 Overall, providers had 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.

Our assessment

3.447 With regard to cloud providers' supply of accelerated compute to FM developers, we have provisionally found that:

(a) FM developers are an increasingly important revenue source for cloud providers that supply accelerated compute because demand for accelerated compute is growing amongst FM developers and cloud providers do not forecast a slowdown in the near future.

(b) AWS, Microsoft and Google all have a strong position in the market for the supply of accelerated compute, each supplying several FM developers as well as developing their own FMs.

(c) As yet, it is unclear if and how cloud providers' relative strengths in the supply of accelerated compute will affect competition in the supply of standard compute. This is because the supply of accelerated compute is not currently substitutable for IaaS based on standard compute infrastructure, due to their different technical specifications and use cases.

(d) Smaller specialised providers of accelerated compute have emerged to supply FM developers. However, they have much smaller capacity and so

667 [✂] response to the CMA's information request [✂].
668 [✂] response to the CMA's information request [✂].

their entry has not had a strong effect on AWS, Microsoft or Google in the provision of accelerated compute.

(e) There are significant barriers to entry and expansion in the supply of accelerated compute. AWS, Microsoft and Google are investing significant sums to procure accelerator chips from third parties – predominantly from Nvidia – or they self-supply.

(f) Partnerships between larger cloud providers and FM developers are widespread. These partnerships could negatively affect competition if they restrict an FM developer from choosing other cloud providers as suppliers of accelerated compute. Partnerships may also lead to benefits for customers by, for example, promoting innovation in the development and deployment of accelerator chips.[669] However, as noted by one provider, providers are likely to be able to achieve such benefits through means other than partnerships.[670]

Cloud providers' supply of access to FMs to their customers

3.448 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 in their own products and services.

3.449 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:

(a) how cloud providers compete in providing customers with access to FMs;

(b) whether this may affect competition in the supply of cloud services; and

(c) providers' relative strengths in the provision of access to FMs.[671]

How cloud providers compete to supply customers with access to FMs

3.450 Cloud providers compete to supply customers with access to FMs by:

(a) developing and commercialising their own FMs; and

(b) offering a range of third-party FMs through their managed platforms.

[669] Responses to the CMA's information requests [✂].
[670] [✂] response to the CMA's information request [✂].
[671] 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.

Providing access to first-party FMs

3.451 AWS, Microsoft and Google's FMs are often only available on their own cloud platforms. (Exceptions include 'open models' that cloud providers make available on platforms such as HuggingFace). For example:

(a) AWS launched its Titan FMs on its Bedrock service for general availability in late September 2023, [✂].[672] In August 2024, AWS [✂].[673]

(b) Microsoft has released versions of its Phi and Orca FMs that are smaller in size than competitors, meaning they can be deployed in a more efficient way.[674] These FMs are currently commercially available on Azure AI Studio and/or Azure Machine Learning Services (and, in some cases, available on an open basis on Hugging Face).[675]

(c) Google's proprietary FMs are called Gemini, and its latest version, Gemini 1.5, was rolled out in February 2024. Gemini FMs are only available to developers through API access on Google's managed platform, Vertex AI.[676] Additionally, Google has also developed Gemma, a family of smaller FMs that have been released as 'open models' and are distributed on a range of platforms.[677]

Providing access to third party FMs

3.452 AWS, Microsoft and Google offer customers access to a range of third party FMs through their managed platforms as set out in the table above.[678] These include Microsoft being the only provider to offer access to OpenAI's FMs.[679]

3.453 The effort to expand access to third-party FMs is driven by 'strong customer demand for using multiple models from different providers' according to a cloud provider's [✂]. [✂] noted that 'more than [✂] of [✂] 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 'expanding model choice [✂].[680]

[672] In a presentation to [✂], AWS noted that [✂]. AWS response to the CMA's information request [✂].
[673] In an August 2024 [✂], AWS [✂]. AWS response to the CMA's information request [✂].
[674] In its Q1 FY24 earnings conference, Microsoft described its recently launched Phi-3 model as 'the most capable and cost-effective Small Language Model (SLM) available in the market', noting that Phi-3 is already being trialled by companies such as PwC. See: Microsoft's earnings conference for Q1 FY24, Microsoft FY24 third quarter earnings conference call, accessed 30 September 2024.
[675] Microsoft's response to the CMA's information request [✂].
[676] Google's response to the CMA's information request [✂].
[677] Gemma: Google introduces new state-of-the-art open models.
[678] Table 3.14 lists a selection of the third party FMs that were available on some cloud providers' managed platforms in September 2024.
[679] Responses to the CMA's information requests [✂].
[680] [✂] response to the CMA's information request [✂].

3.454 Evidence from cloud providers' internal documents and earnings conferences in Q1 FY24 demonstrates their efforts to expand the range of third party FMs they offer:

(a) In its earnings conference for Q1 FY24 [✂], AWS mentioned the continued expansion in the number of FMs offered,[681] through its managed Amazon Bedrock service.[682] [✂].[683]

(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.[684]

(c) In its earnings conference for Q1 FY24, Google said that it now offers more than 130 models through its cloud platform, including its own first-party models, open-source models and third-party models.[685]

3.455 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 [✂].[686]

(b) AWS' partnership with Anthropic allows AWS to offer its customers early access to fine-tuning capabilities for Anthropic's models [✂].[687]

(c) Google's partnership with [✂] FM developers [✂] permit Google to optimise their FMs with its cloud infrastructure.[688]

Effect on competition in the supply of cloud services

3.456 In this section, we consider whether providing customers with access to FMs on cloud platforms is important in competing for overall cloud services.

---

[681] [✂].
[682] AWS' earnings conference for Q1 of FY24, April 30th 2024, accessed 30 September 2024; AWS' response to the CMA's information request [✂].
[683] [✂] response to the CMA's information request [✂].
[684] Microsoft's earnings conference for Q3 of FY24, April 25th 2024, Microsoft FY24 third quarter earnings conference, accessed 30 September 2024.
[685] Google's earnings conference for Q1 of FY24, April 25th 2024, Alphabet first quarter 2024 earnings conference call.
[686] [✂] response to the CMA's information request [✂].
[687] [✂] response to the CMA's information request [✂].
[688] [✂] response to the CMA's information request [✂].

Cloud providers' views

3.457 Cloud providers have submitted different views on the importance of offering access to FMs in competing for cloud customers. The three main cloud providers said that it is not a crucial factor in competing for customers' cloud workloads.

(a) One cloud provider 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, the cloud provider 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.[689]

(b) Another cloud provider submitted that, although it expects that AI will become an important feature of its cloud offering, it will not be the driving force of competition for public cloud services in the foreseeable future.[690] That provider 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'.[691]

(c) Another cloud provider 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'.[692] The cloud provider separately submitted to the CMA that FMs are relevant 'to a customer's choice of not only AI/FM platform, but also in turn wider cloud services' and '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 – and customers may choose a different AI/FM platform which offers access to FMs that we do not'. The cloud provider added that 'customers may purchase other PaaS services as complementary to their access to an FM, such as data and analytics tools' and that 'customers who purchase access to FMs typically also purchase a wider range of basic cloud infrastructure services to support their FM use (such as storage, networking and compute).'[693]

[689] [✂] response to the CMA's information request [✂].
[690] In particular, this cloud provider [✂] referenced [✂]. [✂] response to the CMA's information request [✂].
[691] This cloud provider [✂] 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 [✂]. [✂] response to the CMA's information request [✂].
[692] This cloud provider [✂] 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'. [✂] response to the CMA's information request [✂].
[693] [✂] response to the CMA's information request [✂].

3.458 Two other cloud providers submitted that supplying access to FMs is becoming increasingly important in competing for cloud customers:

(a) 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 'customers are looking for more and more AI enabled services' and 'an AI offering is very important to compete'.[694]

(b) 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]'.[695]

3.459 Although [✂] and [✂] have submitted that access to FMs is not currently an important factor in competing for customers' cloud workloads, their internal documents demonstrate that providing access to FMs is becoming a strategic focus for them. Internal documents also show that the supply of access to FMs is contributing to these cloud providers' global revenue growth, as set out below.

3.460 AWS, Microsoft and Google's internal documents outline that the usage and revenue has grown rapidly for the new managed platforms they have launched to provide access to FMs. This growth trend – which cloud providers expect to accelerate in the future – is having a positive direct effect on their global revenues. In particular:

(a) In an August 2024 planning document, AWS projected that the global (annual run rate) of Bedrock would be [✂] and its global customer spend run rate [✂]. The same document projected Bedrock's global revenue to reach [✂].[696] For context, in an August 2023 planning document, AWS projected that its total global revenue would be [✂].[697]

(b) A March 2024 Microsoft Board update estimated the global revenue from Azure AI platform (excluding OpenAI) to be [✂]. The same document predicted Microsoft's 'AI net revenue',[698] would [✂].[699]

694 [✂] response to the CMA's information request [✂].
695 [✂] response to the CMA's information request [✂].
696 AWS' response to the CMA's information request [✂].
697 AWS' response to the CMA's information request [✂].
698 Microsoft AI revenue includes [✂].
699 Microsoft's response to the CMA's information request [✂].

(c) A 2024 Google internal document identified GenAI as [✂] with Vertex AI as its primary offering.[700] In a 2024 competitor analysis of Google, Microsoft noted that [✂].[701]

3.461 Additionally, cloud providers' documents and public statements discuss how rising demand for access to FMs translates into increased demand for other cloud services, indirectly impacting cloud providers' revenue:

(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.'[702] In Microsoft's Q2 FY24 earnings call it noted that it now has 53,000 Azure AI customers, over one-third of which were new to Azure over the past 12 months and that strong demand for its Microsoft Cloud offerings, including AI services, drove better than expected growth in large, long-term Azure contracts.[703]

(b) In an October 2023 internal presentation discussing Azure AI positioning, Microsoft noted that it would continue to [✂].[704]

(c) In a September 2023 competitor analysis of Azure/OpenAI, Google expressed its aim to [✂].[705]

(d) One cloud provider's internal document stated that 'investing significantly in AI will influence overall Cloud vendor preference, and open doors to new accounts that would otherwise go to [other cloud providers].'[706]

(e) One cloud provider submitted that 'there is already evidence of cloud customers expressly choosing Azure in order to access OpenAI's Flagship FMs through Azure OpenAI Service'. The same cloud provider gave two examples of customers that had 'historically used non-Azure clouds' now reportedly spending 'millions per month' on Azure OpenAI Service.[707]

Customer views

3.462 Evidence from customers indicates that being able to access FMs is emerging as a potential future driver of their choice of cloud service provider. However, for most

[700] Google's response to the CMA's information request [✂].
[701] Microsoft's response to the CMA's information request [✂].
[702] Microsoft's earnings conference for Q3 of FY24, 25th April 2024, accessed 30 September 2024.
[703] Microsoft's earnings conference for Q2 FY24, 30th January 2024, accessed 2 December 2024, pages 4 and 16.
[704] Microsoft's response to the CMA's information request [✂].
[705] Google's response to the CMA's information request [✂].
[706] [✂] response to the CMA's information request [✂].
[707] The Information has reported that TikTok is spending $20m per month on Azure OpenAI Service to access OpenAI's FMs' and 'Intuit is also spending millions per month on Azure OpenAI Service – both of which are customers that historically used non-Azure clouds (eg Intuit previously used AWS as its primary cloud provider)'. [✂] response to the CMA's information request [✂]; TikTok Spending Drove Microsoft's Booming AI Business.

customers it was not a key consideration when they last chose their main cloud provider.

3.463 We asked large customers of public cloud providers to rate the importance of providers' AI capabilities when choosing their main cloud provider.[708] We found that most customers currently regard access to FMs as a material but not decisive factor.[709] However, many large customers said that AI capabilities are expected to become increasingly important in the future.[710]

(a) Two of these customers said that, although they view AI capabilities as extremely important, their choice of a cloud provider predated the rise of AI and, therefore, was not influenced by it.[711]

(b) Two other customers submitted that AI capabilities are increasingly becoming a source of differentiation between cloud providers, as cloud providers are 'leapfrogging' each other in this space.[712]

(c) One customer, who described the availability of AI services as a 'critical factor in choosing a particular Cloud provider',[713] said that [✂].[714]

3.464 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 PaaS solution and compete strongly for this AI PaaS business.[715] 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 which they are looking at and expecting to make use of in the near future. Regardless of the context, these participants all said that AI is likely to be the most significant change to public cloud services over the next five years or so'.[716]

3.465 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 submitted 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

[708] Although AI capabilities were not defined, we have inferred from customers' responses that they are referring to access to FMs and not AI services or accelerated compute. This inference is further supported by the fact the standard cloud customers surveyed were not FM developers and were therefore not likely to be referring to access to accelerated compute. Furthermore, at the time of gathering this evidence many AI services were yet to be developed. However, we note that different customers may have interpreted AI capabilities differently and we have taken that into account when assessing this evidence.
[709] The average rating for the importance of this factor was close to 2.5/5
[710] Responses to the CMA's information requests [✂].
[711] Responses to the CMA's information requests [✂].
[712] Responses to the CMA's information requests [✂].
[713] [✂] response to the CMA's information request [✂].
[714] As noted earlier, it is unclear whether this statement refers to the provision of FM-enabled services by providers or access to FMs. [✂] response to the CMA's information request [✂].
[715] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), page 89.
[716] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), para. 8.1.1.

using third-party FM tools integrated and hosted within the same cloud service provider as the datastore you are using'.[717]

3.466 Overall, providing customers with access to FMs is impacting cloud providers' global revenues and some of them consider it to be an important factor in competing for cloud customers. The internal documents from some providers set out above also show that they expect the revenue from their managed platforms to [✂],[718] indicating that the importance of offering this service could increase.

3.467 However, evidence from customers is more mixed, with most customers submitting that providers' AI capabilities are not an important consideration driving their choice of a provider, and some customers noting this may change moving forward. As different customers may have interpreted differently the meaning of 'AI capabilities', we have been cautious of placing too much weight on this evidence.

Impact of supply of access to FMs on competition for other cloud services

3.468 As supplying access to FMs becomes more important in competing for cloud services customers, providers' relative strength in supplying access to FMs may impact their strength in the supply of cloud services, particularly if accessing FMs from a cloud provider leads customers to buy additional cloud services from them.

3.469 We consider below the likelihood of customers adopting a multi-cloud strategy to access FMs across different providers.

Providers' submissions

3.470 AWS and Microsoft submitted that customers' growing interest in accessing FMs may lead to a higher use of multi-cloud architectures:

(a) AWS submitted 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.[719]

(b) Microsoft submitted that 'Generative AI applications are particularly suited to multi-cloud strategies as the cost of processing data by querying a 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

[717] [✂] response to the CMA's information request [✂].
[718] Responses to the CMA's information requests [✂].
[719] 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: [✂] response to the CMA's information request [✂].

provider's cloud] while integrating AI features through an independent FM developer (…), or instead through another cloud provider's service'.[720]

(c) Microsoft also stated that, due to the fast pace at which FMs are developing, customers are [✂].[721]

3.471 In contrast, other cloud providers submitted that there may be barriers to adopting a multi-cloud strategy for accessing FMs:

(a) One cloud provider submitted 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'.[722] [✂] also submitted that if a cloud provider were 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'.[723]

(b) One cloud provider submitted 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.[724]

(c) 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'.[725]

(d) A different cloud provider submitted that co-locating customers' workloads and data with the FMs they want to use will likely benefit customers.[726]

Customer views

3.472 Some customers said that, where necessary, they can access the AI services from providers other than their main one.[727]

---

[720] Microsoft's response to the CMA's information request [✂].
[721] Microsoft's response to the CMA's information request [✂].
[722] [✂] response to the CMA's information request [✂].
[723] [✂] response to the CMA's information request [✂]. [✂] submission to the CMA [✂].
[724] [✂] response to the CMA's information request [✂].
[725] [✂] response to the CMA's information request [✂].
[726] [✂] response to the CMA's information request [✂].
[727] Responses to the CMA's information requests [✂].

3.473 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.

3.474 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.[728] The report noted that, moving forward, the growing appeal of multi-cloud strategies may force customers to confront barriers to multi-cloud that, in their current set up, they considered trivial or irrelevant to them and their business.[729]

3.475 While these 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:

(a) 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'.[730]

(b) Another 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'.[731]

3.476 One customer said that accessing OpenAI's FMs from Azure 'has limited the choice of other cloud vendors for integration reasons'.[732] Similarly, another customer said that 'it is easier to implement OpenAI/Microsoft FMs because we use Microsoft cloud service provider'.[733]

Internal documents

3.477 Providers' internal documents discuss how the supply of access to FMs may raise barriers to switching or multi-cloud amongst customers:

[728] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraphs 8.1.2-8.1.3
[729] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 8.1.4
[730] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 8.1.7
[731] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 8.1.8
[732] [✂] response to the CMA's information request [✂].
[733] [✂] response to the CMA's information request [✂].

(a) In an undated competitive analysis of the supply of FMs, Microsoft identified [✂].[734] Additionally, an October 2023 presentation on Microsoft's [✂].[735]

(b) Similarly, an August 2024 AWS planning document for its ML infrastructure and GenAI businesses noted that [✂].[736]

Cloud providers' relative strength in the provision of access to FMs

3.478 We have assessed evidence from customers and cloud providers' internal documents regarding providers' relative strength in the supply of access to FMs.

Customer views

3.479 Customer evidence indicates that they perceive the competitive landscape to be quite dynamic. For example:

(a) Although participants in Jigsaw's research 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.[737]

(b) Some large customers told us that cloud providers are 'leapfrogging' each other in the provision of access to FMs, suggesting that there is no provider that is consistently 'ahead' in this space.[738] This perception of dynamism contrasts with customers' views regarding other cloud services – some customers have noted that in other respects the main cloud providers' capabilities had become broadly similar over time.[739]

Microsoft strength in the provision of access to FMs

3.480 Microsoft's internal documents indicate that it views its own strengths in the provision of access to FMs as being:

(a) Its [✂]. [✂].[740] [✂].[741] [✂].[742]

(b) Its provision of [✂].[743]

---

[734] Microsoft's response to the CMA's information request [✂].
[735] Microsoft's response to the CMA's information request [✂].
[736] AWS' response to the CMA's information request [✂].
[737] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 8.1.6.
[738] Responses to the CMA's information requests [✂].
[739] Responses to the CMA's information requests [✂]. CMA note of meeting with [✂].
[740] Microsoft's response to the CMA's information request [✂].
[741] Microsoft's response to the CMA's information request [✂].
[742] Microsoft's response to the CMA's information request [✂].
[743] Microsoft's response to the CMA's information request [✂].

3.481 FM developers and cloud providers' revenues from the supply of first-party FMs and ancillary services demonstrate that OpenAI's FMs account for a [✂] share of FM revenues in 2023 and up to June 2024.[744]

3.482 Anthropic and Google are OpenAI's next largest competitors, in terms of revenue from the provision of FMs.[745] Customers also considered OpenAI's FMs as the best FMs currently available, with Anthropic, Google and Meta's FMs being the closest competitors.[746]

3.483 The global revenue from Azure AI platform is [✂].[747] In its earnings conference for Q4 FY24, Microsoft's CEO stated that '[Azure's] share gains accelerated this year, driven by AI' and, as of July 2024, '[Azure had] over 60,000 Azure AI customers, up nearly 60% year-over-year and average spend per customer continues to grow'.[748]

AWS' strength in the provision of access to FMs

3.484 AWS' internal documents indicate that AWS views its own strengths in the provision of access to FMs as being:

(a) Model choice - [✂].[749] [✂].[750] In the same document, AWS highlighted that [✂].[751]

(b) Its emphasis on 'strong security [✂].[752]

3.485 In its September 2023 update to Amazon's Board of Directors, AWS stated that, although [✂].[753]

3.486 However, in an August 2024 planning document, AWS noted [✂].[754]

Google's strength in the provision of access to FMs

3.487 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 FMs, Gemini, which is often benchmarked against [✂] models.[755]

[744] [✂]; Responses to the CMA's information requests [✂].
[745] Responses to the CMA's information requests [✂].
[746] 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 CMA's information requests [✂].
[747] Microsoft's response to the CMA's information request [✂].
[748] Microsoft Fiscal Year 2024 Fourth Quarter Earnings Conference Call, accessed 10 October 2024.
[749] AWS' response to the CMA's information request [✂].
[750] AWS' response to the CMA's information request [✂].
[751] AWS' response to the CMA's information request [✂].
[752] AWS' response to the CMA's information request [✂].
[753] AWS' response to the CMA's information request [✂].
[754] AWS' response to the CMA's information request [✂].
[755] Google's response to the CMA's information request [✂].

3.488 Regarding its third-party FM offering, Google noted in an internal document discussing the 'strategic rationale' for its partnership with Anthropic that [✂].[756]

3.489 Microsoft submitted 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.[757]

Our assessment

3.490 Cloud providers' supply to customers of access to FMs is nascent and is evolving rapidly.

3.491 The increase in customer demand for FMs has been driving growth in cloud providers' revenues in two ways: directly through increased earnings from their managed platforms, and indirectly by promoting customers' use of other cloud services from the provider who offers access to FMs.

3.492 Providing access to FMs has emerged as a potential future driver of customers' choice of cloud service provider. Cloud providers strengthen their ability 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.

3.493 Partnerships between FM developers and cloud providers could play an important role in shaping the 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.

3.494 The impact that cloud providers' relative strengths in supplying customers with access to FMs may have on their strength in the supply of cloud services will also depend on whether customers can access FMs through a multi-cloud approach and the extent to which cloud providers can attract additional (or retain existing) revenue on cloud services from customers interested in accessing a particular FM. Customers may be incentivised to multi-cloud where they want to access FMs hosted in a cloud other than that of their main provider. However, any technical or commercial barriers to multi-cloud may prevent this from happening.

3.495 Cloud providers have differentiated strengths in the supply to customers of access to FMs. However, customers do not think that the current competitive landscape in respect of FMs is necessarily fixed.

---

[756] Google's response to the CMA's information request [✂].
[757] Microsoft's response to the CMA's information request [✂]. Microsoft's submission to the CMA [✂].

Provisional conclusions

3.496 We have provisionally found that there are separate markets for IaaS based on standard compute infrastructure (excluding IaaS based on accelerated compute infrastructure) and PaaS, each having a European (UK and EEA) geographic scope.

3.497 We have also provisionally found that SaaS services and alternative IT models (traditional IT and private cloud) do not form part of the same markets as IaaS and PaaS.

3.498 We have looked at evidence on market shares and shares of flow of new business across IaaS, PaaS and IaaS and PaaS combined. In interpreting these shares, we account for the following considerations:

(a) Market shares for the IaaS market and, separately, market shares for the PaaS market are, theoretically, the most relevant shares for our competitive assessment. That is because they align with our provisional findings on market definition.

(b) However, due to different ways providers and market research companies classified revenues into IaaS and PaaS, the calculations we used for the individual IaaS and PaaS markets may not accurately reflect our product market definition. For this reason, we have also looked at the market share IaaS and PaaS combined. This does not change our overall conclusions.

(c) We have also discussed other methodological points that affect our interpretation of the results. For example, we consider that shares of supply by revenue for IaaS may overstate the strength of smaller providers for which we do not have direct data. In addition, revenues associated with IaaS often do not distinguish between standard and accelerated compute; however, we understand the contribution of accelerated compute to the total revenues to be small.

3.499 In relation to the IaaS market (excluding accelerated compute infrastructure), we have provisionally found that it is highly concentrated and is likely to remain so in the near future. The two largest providers, AWS and Microsoft, have held large market shares for a sustained period of time. This is in contrast to all other suppliers which are much smaller. The position of AWS and Microsoft is unlikely to change in the near future.

3.500 AWS and Microsoft are the cloud providers with largest shares of overall revenue growth. On this basis, our provisional view is that the market is likely to remain concentrated and AWS and Microsoft to remain significantly larger than other providers in the IaaS market.

3.501 In relation to the PaaS market, we have provisionally found that it is concentrated, albeit to a lesser extent than in the IaaS market and is likely to remain so in the near future. AWS, Microsoft and Google are the largest providers of PaaS services although their market shares are smaller than in IaaS due to the presence of several ISVs.

3.502 Despite the relative lower concentration in the PaaS market, ISVs do not provide a strong competitive constraint on vertically integrated cloud providers (in particular, AWS, Microsoft and Google) due to their dependency on the three main providers' infrastructure. Indeed, any sale ISVs make is also, typically, an IaaS sale for one of AWS, Microsoft or Google, as IaaS is an input to PaaS.

3.503 In terms of market outcomes, we have provisionally found that AWS and Microsoft's cloud businesses (Microsoft Cloud & Enterprise and Microsoft Azure) have had sustained profits (as measured by the return on capital employed) substantially above our estimated cost of capital for a number of years.

3.504 Our provisional view is that evidence on market shares and profitability indicate that AWS and Microsoft each holds a strong position in the IaaS and PaaS markets. Google is the largest of the remaining competitors, but it remains significantly smaller than them. Other providers of either IaaS or PaaS services are smaller and provide only a weak competitive constraint on AWS, Microsoft and Google. Forward-looking metrics suggest this market structure and outcomes are likely to endure.

3.505 While customers recognise that cloud services offer both quality and innovation to them, we consider that a more competitive market would have sustained better market outcomes. These include more consistently competitive prices, as well as improvements in quality and innovation

3.506 We sought to understand the potential impact of AI on how competition works in cloud services. We looked at the potential impacts on competition through two lenses: the extent to which competitive conditions in IaaS based on accelerated compute, including the partnerships formed by cloud providers and FM developers, could affect the market for IaaS based on standard compute, and the rising importance of access to FMs by customers of cloud services.

3.507 We have found that partnerships between larger cloud providers and FM developers are widespread and that they may play an important role in shaping the competitive conditions in the supply of accelerated compute to FM developers and in the supply of access to FMs to other customers.

3.508 In relation to the first lens, although we note that AWS, Microsoft and Google each has a strong position in the supply of IaaS based on accelerated compute, the lack of substitutability with IaaS based on standard compute and the current uncertainty as to precisely how commercial and technical links between accelerated compute and different parts of the standard compute supply chain might evolve, have led us to provisionally find that there is currently no significant direct impact from IaaS based on accelerated compute on competition in IaaS based on standard compute.

3.509 In relation to the second lens, we provisionally consider that providing access to FMs has emerged as a potential future driver of customers' choice of cloud service provider. However, customers do not think that the current competitive landscape in respect of FMs is fixed. Based on the evidence we have considered, we have provisionally found that access to FMs by customers of cloud services is not currently a strong driver of customer choice and that the extent to which it will become a driver of choice in the future is uncertain.

3.510 In view of our provisional findings that AWS and Microsoft each hold significant unilateral market power in cloud services, and that each has a strong position in the market for IaaS based on accelerated compute, it is important that their potential influence on the development of FM-related markets does not give rise to competition concerns. We consider that the application of the CMA's competition and consumer protection principles for FM development and deployment as set out in its AI Foundation Models reports are relevant to ensuring that competition works well as this area of cloud services continues to evolve.[758] These principles are helpful as factors that the CMA can use to monitor developments and consider any potential need for use of its powers to address any harms to competition as they emerge.

3.511 For instance, if FM developers are restricted in accessing accelerated compute, this could undermine competition in the supply of FMs; if FM developers are restricted in how and where they can commercialise their FMs or if cloud providers are restricted in being able to provide access to FMs that emerge as important, this may limit choice of FMs for cloud customers. Any exclusivity in partnerships between FM developers and cloud providers or restrictions on the development of open and fair competition between FM developers, or the downstream deployment of FMs, could exacerbate these risks.

3.512 Due to the pace of change in these areas, we will suggest that the CMA should continue to monitor market developments and to consider if or when any intervention is needed to ensure competition works well.[759]

[758] In particular, the principles of access, diversity and choice. Access – ongoing ready access to inputs, such as accelerated compute; Diversity – sustained diversity of business models and FM types; Choice – sufficient choice for businesses and consumers to they can decide how to use FMs. Figure 3 in AI Foundation Models: Update Paper - GOV.UK.
[759] AI Foundation Models: initial review - GOV.UK

4. Barriers to entry and expansion

- 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.

- We have provisionally found substantial barriers to entry and expansion in the provision of cloud services, in particular 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, whereby the larger providers have comparatively lower 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. This disincentivises IaaS market entry and expansion.

- Furthermore, given the scale of investment and expansion that large cloud providers have made to date in IaaS, any new entrant (or company seeking to expand) would need to invest substantially more than the existing large providers in order to close the gap in a timely way. The levels of investment that AWS and Microsoft are expecting to make in the coming years may raise these barriers even higher.

- While we recognise that investments by cloud providers may 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 supplier and ISVs value access to a wider user base. Customers also place importance on a provider's reputation and this contributes to overall barriers to entry and expansion in cloud services.

- AWS and Microsoft appear to be the largest providers to the public sector with similar positions to those they have in the cloud services markets overall. Their leading positions with public sector customers are likely driven by the features we have identified elsewhere in this provisional report and so covered by our wider assessment. We have not seen any evidence that that public sector procurement practices are harming competition in the market, but we propose suggesting that the UK government should promote best practice in procurement as public sector use of cloud services continues to grow.

4.1 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 may be harming competition.[760] 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.[761]

4.2 This chapter sets out our assessment of barriers to entry and expansion in the cloud services markets. To do this, we assess 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.

4.3 We also assess 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

4.4 Economies of scale, in combination with sunk investment costs, can constitute a barrier to entry or expansion.[762] Entry on a large scale will often entail a high risk (that sunk investment costs may not be recovered) because it will generally be successful only if the firm can expand the total market significantly, or substantially replace one or more existing firms.[763]

4.5 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.

4.6 We focus on large, unavoidable, 'sunk' costs because they could deter new entry into the market, as new entrants would need to make significant upfront investments in infrastructure without any certainty of return. If their plans to enter

---

[760] CC3 (Revised), Guidelines for market investigations: Their role, procedures, assessment and remedies, paragraph 205.
[761] CC3 (Revised), paragraph 207.
[762] CC3 (Revised), paragraph 212.
[763] CC3 (Revised), paragraph 213.

the market failed, they would be unable to recoup the initial investment, further increasing the risk of entry for them.

4.7 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.[764]

4.8 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.

Stakeholders’ views

4.9 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;[765]

(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, Alibaba, Cloudsky, Skytap, Digital Ocean and OVHcloud as evidence of competition in the cloud services markets;[766] and

(a) new entrants can defer fixed costs to third party data centre providers and by doing so incur costs that are proportionate to their size.[767]

Our assessment

4.10 Cloud providers primarily use their fixed asset infrastructure to deliver IaaS.

4.11 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.

[764] Six advantages of cloud computing - Overview of Amazon Web Services, accessed 6 November 2024.
[765] AWS’ Economic response to the CMA’s competitive Landscape Working Paper (prepared by CRA) dated 2 August 2024, paragraph 116.
[766] AWS’ Economic response to the CMA’s competitive Landscape Working Paper (prepared by CRA) dated 2 August 2024, paragraph 89.
[767] AWS’ Economic response to the CMA’s competitive Landscape Working Paper (prepared by CRA) dated 2 August 2024, paragraph 90.

4.12 However, competition in PaaS using third party IaaS likely involves sourcing IaaS from firms that are also competing in the supply of PaaS. This has the potential to limit the scope for PaaS providers to be fully effective competitors as:

(a) IaaS providers may be 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.

4.13 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.

4.14 Cloud providers group their fixed assets under three broad headings:[768]

(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.[769]

4.15 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.[770] 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 likely be less than the initial investment, as the value of the asset decreases over time.

[768] Responses to the CMA's information requests [✂].
[769] These include central processing units (CPUs), graphics processing units (GPUs), dynamic random access memory (DRAM), Motherboards, solid state drives (SSDs) and hard disk drivers (HHDs).
[770] This is standard accounting practice.

4.16 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 costs a cloud provider would incur disassembling and selling the assets.

4.17 We have reviewed the carrying value of the global fixed assets of AWS, Google, Oracle Cloud and IBM Cloud as at 31 December 2022 and for Microsoft Azure as at 30 June 2023.[771] The evidence shows that AWS and Microsoft have invested significantly more in global infrastructure than smaller cloud providers.[772] For example:[773]

(a) AWS had invested in cloud fixed assets with a carrying value of $[✂] as at 31 December 2022;[774]

(b) Microsoft had invested in cloud fixed assets with a carrying value of $[✂] as at 30 June 2023;[775]

(c) Oracle had invested in cloud fixed assets with a carrying value of $[✂] as at 31 December 2022;[776] and

(d) IBM had invested in cloud fixed assets with a carrying value of $[✂].[777]

4.18 Public statements and internal documents from AWS and Microsoft indicate that they expect to increase their investments in cloud fixed assets in the next couple of years:[778]

(a) [✂].[779]

(b) [✂].[780]

4.19 In addition to their investments in fixed assets, AWS and Microsoft have also made acquisitions to improve their network infrastructure.[781] This is further evidence of significant capital investment, which could deter entry and act as a barrier to expansion by smaller providers.

---

771 The financial year end for AWS, Google, Oracle and IBM is 31 December. The financial year end for Microsoft is 30 June.
772 See detail in Appendix G.
773 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.
774 AWS' response to the CMA's information request [✂].
775 Microsoft's response to the CMA's information request [✂].
776 Oracle's response to the CMA's information request [✂].
777 IBM's response to the CMA's information request [✂].
778 Amazon share pop on earnings beat, cloud growth, accessed on 6 November 2024; Microsoft FY24 Third Quarter Earnings Conference Call, accessed 6 November 2024. [✂] response to the CMA's information request [✂].
779 [✂] response to the CMA's information request [✂].
780 [✂] response to the CMA's information request [✂].
781 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.

Co-locating vs owning data centres

4.20 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.[782]

(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.[783]

(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.[784]

4.21 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 AWS, Microsoft 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.[785] It recently committed to invest £8 billion between 2024 and 2028 to build and operate data centres in the UK.[786]

(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.[787]

(c) Google said that all of its data centres in the UK are currently owned and operated by third parties.[788] It recently announced plans to invest $1 billion to build its first UK data centre.[789]

(d) Oracle said that it does not typically build or operate data centres but it leases them (globally).[790]

---

782 AWS' response to the CMA's information request [✂]; Google's response to the CMA's information request [✂].
783 AWS' response to the CMA's information request [✂].
784 Google's response to the CMA's information request [✂].
785 AWS' response to the CMA's information request [✂].
786 AWS UK – Driving Digital Acceleration, accessed on 7 October 2024.
787 In 2023, Microsoft began construction on 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.
788 Google's response to the CMA's information request [✂].
789 Google starts building £790m site in Hertfordshire - BBC News, accessed on 29 October 2024.
790 Oracle's response to the CMA's information request [✂].

4.22 The evidence we have seen shows that new entrants (and existing providers) can, in some circumstances, use co-location to enter the IaaS market and to expand into new regions.

4.23 However, even if a new entrant enters the IaaS market, the sheer scale of investment by AWS, Microsoft and to a lesser extent Google, means that any new entrant would not be able to compete directly with these providers, 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.

4.24 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.

4.25 This shows that entry, even on a small scale, requires investment and involves considerable risk.

4.26 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 tend to be firms offering GPU capacity to deliver Gen-AI (eg Voltage Park), clouds limited by geography (eg Cloudsky) or storage only services (eg Wasabi).

Our assessment

4.27 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.

4.28 The size of the investments made by the largest cloud providers are very large relative to those of smaller providers: the investments planned by AWS and Microsoft 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 begins to close the gap with AWS and Microsoft in IaaS provision, a rival provider would have to make even larger annual investments.[791]

4.29 While it is, in many markets, feasible for smaller rivals to constrain larger ones without matching their base of assets, the assets we have discussed in this section - the data centres, networking 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.

4.30 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.[792]

4.31 On this basis, the size of AWS' and Microsoft's investments and holdings of assets represent substantial barriers to entry for any provider that does not have the capital available to make investments of a similar magnitude.

4.32 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. However, entry on a small scale means that any new entrant would not be able to compete on an equal footing with 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 large 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.

4.33 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.

4.34 While there is some evidence of entry and expansion into niche areas within the IaaS market, we do not consider this to 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.

---

[791] 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. As a result, we do not have the data to compare Google's cloud-related assets with those of AWS and Microsoft.
[792] CC3 (Revised), paragraph 213.

Economies of scale

4.35 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.[793]

4.36 In this section, we set out our assessment of whether large 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.

4.37 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.

Stakeholders' views

4.38 Below we set out stakeholders' views that we have received on economies of scale.

4.39 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 more limited than we considered in our working paper on the competitive landscape.[794]

4.40 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

[793] CC3 (Revised), paragraph 212.
[794] AWS' economic response to the CMA's competitive landscape working paper, dated 23 May 2024, paragraphs 92 and 93.

factors above, it is unsurprising that a small number of super-providers have come to dominate these markets.[795]

Purchasing efficiencies

4.41 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.

4.42 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;[796]

(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;[797] and

(c) IBM said that procurement is shared across business units and divisions, allowing for purchase discounts and more efficient supply chain management.[798]

4.43 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).

4.44 Given that the cloud-related fixed asset bases of AWS and Microsoft are significantly larger than the smaller cloud providers and that AWS and Microsoft's 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.

4.45 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.

795 [✂] submission to the CMA [✂].
796 AWS' response to the CMA's information request [✂].
797 Microsoft's response to the CMA's information request [✂].
798 IBM's response to the CMA's information request [✂].

Operating efficiencies

4.46 Below, we consider 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.

4.47 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.[799]

4.48 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

4.49 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.

4.50 We have analysed the weighted average Power Usage Effectiveness (PUE) [800] of AWS, Microsoft Azure, Google Cloud, Oracle Cloud and IBM Cloud. Our analysis shows that:[801]

[799] Microsoft 2023 Annual Report, accessed 20 September 2024.
[800] 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 20 September 2024.
[801] See detail in Appendix J.

(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;[802] and

(b) AWS, IBM and Microsoft have [✂].[803]

4.51 We also received 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.[804]

(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.[805]

(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.[806]

(d) A global data centre provider, said that larger data centres achieve economies of scale through more efficient use of energy.[807]

(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.[808]

4.52 Smaller cloud providers may be 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 unlikely to benefit from other energy efficiencies to the same extent as larger providers. For example, when compared to smaller providers, AWS, Microsoft and to a lesser extent Google, have the ability to:

(a) negotiate bigger volume discounts through their supply contracts;

[802] See detail in Appendix J.
[803] [✂].
[804] Google's response to the CMA's information request [✂].
[805] Google's response to the CMA's information request [✂].
[806] Microsoft's response to the CMA's information request [✂].
[807] Note of meeting with [✂].
[808] Note of meeting with [✂].

(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.[809]

Data centre capacities

4.53 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.[810] We have also considered the number and size of the data centres of six cloud providers: AWS, Microsoft, Google, Oracle, IBM and OVHcloud.[811]

4.54 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.[812]

4.55 The analysis of capacity shows that [✂] 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.[813]

4.56 There are differences in composition of the global data centre networks between the different cloud providers. However, in general, the larger cloud providers tend to have more larger data centres and smaller providers tend to have mainly smaller data centres.[814]

Utilisation

4.57 Cloud providers aim for high levels of utilisation of their data centres in order to maximise cost efficiencies.[815]

[809] 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.
[810] See details in Appendix J. Sources: Data from cloud providers.
[811] See analysis in Appendix J. Sources: Data from cloud providers.
[812] What to consider when selecting a region for your workloads, accessed on 7 October 2024.
[813] Responses to the CMA's information requests [✂].
[814] See details in Appendix J.
[815] [✂] response to the CMA's information request [✂].

4.58 We sought to gather data from cloud providers to understand their levels of utilisation, however not all the data that we received allows for comparisons between companies.

4.59 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.[816] AWS' utilisation rate for 2023 of [✂]% is [✂]% [✂] than that of IBM, who notes that its current average level of both daily and monthly utilisation is approximately [✂]% and that it manages its capacity so as not to exceed [✂]% utilisation.[817]

4.60 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.

4.61 We have also analysed the utilisation rates that cloud providers consider to be 'high',[818] and are finding that:

(a) the providers' view 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.[819]

4.62 The evidence that we have received indicates that [✂] utilisation rates compared to smaller providers and that the high utilisation rates that [✂].

Investment in research and development

4.63 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 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.[820]

[816] We have 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 [✂].
[817] IBM's response to the CMA's information request [✂].
[818] See details in Appendix J.
[819] See details in Appendix J, Sources: Data from cloud providers.
[820] CC3 (Revised), paragraph 183.

4.64 Spend on research and development 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 research and development is significant, it could be difficult for smaller firms to achieve and so could contribute to barriers to entry and expansion.

4.65 We have considered the following evidence on cloud providers' spend on research and development:

(a) Microsoft said that it spent $27bn on research and development in financial year 2023 and that the other large cloud providers made similar investments in research and development. It also said that high levels of research and development spend are required to attract new customers and limit the number of customers switching to rivals.[821]

(b) IBM also said that it incurred significant annual expenditure on research and development exceeding $[✂] across its global technology businesses and that investment results in innovation benefitting products both in and out of cloud.[822]

(c) Google said that its internal research and development entails a degree of fixed and sunk costs. It also said that it is not necessary to make significant research and development investments or incur costs to begin to offer basic services (compute, storage and networking).[823]

(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.[824] [✂].[825]

4.66 In summary, the size of the investments made by the largest cloud providers in research and development 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 in order to achieve similar outcomes.

[821] Microsoft's response to the CMA's information request [✂].
[822] IBM's response to the CMA's information request [✂].
[823] Google's response to the CMA's information request [✂].
[824] AWS' response to the CMA's information request [✂].
[825] AWS' response to the CMA's information request [✂]. [✂].

Our assessment

4.67 There appear to be material economies of scale in the provision of cloud services, in particular for IaaS. These include that larger cloud providers generally:

(a) achieve purchasing efficiencies from bulk discounts on necessary equipment;

(b) achieve 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) are making material investments in R&D which they can then spread over a wider business.

4.68 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 consider that economies of scale act as material barriers to entry and expansion.

Size of cloud providers' product portfolio

4.69 We have considered whether having a large portfolio of cloud services may give a cloud provider advantages over its rivals.

4.70 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.

4.71 In assessing cloud providers' product portfolios, we look at both first and third party services, as well as both IaaS and PaaS.

4.72 We break down product portfolios by discussing:

(a) the importance of range of services;

(b) economies of scope; and

(c) network effects

4.73 While we do not assess it in this section, we recognise that there may be economies of scope across cloud and non-cloud services, rather than just across cloud services.

Importance of range of services

4.74 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.

4.75 We have gathered customer views on the importance of range of first and third party products.

4.76 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.

4.77 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.[826]

4.78 Customers rating it 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.[827]

(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.[828]

4.79 Some customers said that the range of cloud infrastructure services is less important because most providers offer the same services.[829]

4.80 However, these customers might have considered only the three main providers when answering those questions.

4.81 The Jigsaw report shows that one of the key reasons why AWS, Microsoft, 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

[826] CMA analysis of responses to the CMA's information requests [✂].
[827] Responses to the CMA's information requests [✂].
[828] Responses to the CMA's information requests [✂].
[829] Responses to the CMA's information requests [✂].

cloud providers, such as IBM and Oracle, are perceived as being for secondary use or only for certain use cases.[830]

Our assessment of the importance of range of services

4.82 We consider that the range of first party products is an important factor for customers when choosing which cloud provider(s) to use.

Economies of scope

4.83 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.[831] Economies of scope might be relevant if, for example, R&D and operations spend can be spread over a wider range of services.

4.84 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.

4.85 Cloud providers said that these efficiencies can 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.

4.86 In relation to the efficiencies that can 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.[832]

(b) Oracle said that cloud providers benefit 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

[830] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraphs 3.4.4 and 3.4.8, page 27.
[831] CC3 (Revised), paragraph 214.
[832] Microsoft's response to the CMA's information request [✂].

fundamental services of compute, storage, networking, identity, and security. It also said that it has benefited from increasing the range of services.[833]

(c) AWS said that an increase in the number of services offered does not necessarily increase overall efficiency for a cloud provider.[834] It also said it benefitted from simplifying operations and sharing learnings 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 discovers a vulnerability or bug in the code, they will 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 do not need to build those tools from the ground up to launch a new service and can pass on benefits to customers in the form of a consistent user experience.[835]

(d) IBM said that cloud providers can benefit from efficiencies as they increase 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 allows 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.[836]

(e) Google said that any efficiencies that may be derived from developing a broader range of services are, in its experience, limited (albeit there may be exceptions).[837]

4.87 In relation to the inefficiencies that may arise:

(a) AWS said that many of the services it offers are discrete and distinct, and the learnings and functionalities that apply to one may not necessarily translate or be useful to others. It also said that increasing the number of service offerings introduces inefficiencies due to increased complexity. It said that customers generally expect 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.[838]

---

833 Oracle's response to the CMA's information request [✂].
834 AWS' response to the CMA's information request [✂].
835 AWS' response to the CMA's information request [✂].
836 IBM's response to the CMA's information request [✂].
837 Google's response to the CMA's information request [✂].
838 AWS' response to the CMA's information request [✂].

(b) Microsoft said that having a large service portfolio requires 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, entities that focus on just one or a limited set of cloud solutions also have their own efficiencies that can be generated by such specialisation.[839]

Our assessment

4.88 There may be some economies of scope in supplying a range of cloud services, but this may 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

4.89 If customers value having access to a large portfolio of 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.

4.90 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.

4.91 To assess the strength of any network effects, we collected evidence from ISVs and customers.

[839] Microsoft's response to the CMA's information request [✂].

Evidence from ISVs

4.92 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 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.[840]

(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.[841]

(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.[842]

(d) Some ISVs also noted that there are significant costs associated with running their services to additional public cloud providers beyond the largest ones.[843]

Evidence from customers

4.93 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.[844]

4.94 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.

[840] Note of meeting with [✂].
[841] Note of meeting with [✂].
[842] [✂] response to Ofcom's information request [✂]; Ofcom note of meeting with [✂].
[843] Responses to Ofcom's information requests [✂].
[844] 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.

4.95 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;[845] and

(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.[846]

4.96 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.[847]

Our assessment

4.97 Range of services is important to customers, and ISVs appear to value a broad customer base.

4.98 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.

4.99 If the concentration in the market kept increasing, such that the incentives on ISVs to support smaller cloud providers decreased because the incremental customer base would not justify the investment, this could result in a larger gap between the services on the largest cloud providers relative to others as a result of ISVs reducing the breadth of suppliers these other providers support.

4.100 We consider that this could compound barriers to entry and expansion as smaller cloud providers would find it harder to offer such a wide range of services.

Other potential barriers to entry and expansion

Cloud credits

4.101 Cloud credits are a form of discounting by cloud providers. They work as follows:

[845] Responses to the CMA's information requests [✂].
[846] Responses to the CMA's information requests [✂].
[847] Responses to the CMA's information requests [✂].

(a) Cloud providers offer customer credits for free as an incentive to switch clouds or use their cloud 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 cloud credits that they will offer by type of customer eg a start-up, an SME, an AI start-up, etc.

(d) Customers are aware of the amount of cloud credits they will receive in advance of using a cloud service - 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.

4.102 Cloud credits are offered by a range of cloud providers:

(a) AWS offers start-ups up to $100,000 in AWS activate credits and up to an additional $300,000 in AI specific credits.[848]

(b) Microsoft offers start-ups up to $150,000 in Azure credits.[849]

(c) Google offers start-ups up to $200,000 in cloud credits and AI start-ups up to $350,000 in cloud credits.[850]

(d) Oracle gives start-ups the ability to purchase discounted credits to scale.[851]

(e) OVHcloud offers start-ups up to €100,000 in cloud credits.[852]

(f) Some smaller cloud providers, such as Civo, also offer cloud credits.[853]

4.103 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.

[848] Get AWS Activate Credits - AWS Startups, accessed on 7 October 2024.
[849] Unlocking Azure credits as your startup grows, accessed on 7 October 2024.
[850] Google for Startups Cloud Program | Google Cloud, accessed on 7 October 2024.
[851] Free cloud is just the beginning, accessed on 7 October 2024.
[852] Startup Program, accessed on 7 October 2024.
[853] Grow with the Civo Startup Program, accessed on 7 October 2024.

Stakeholders' views

4.104 Two smaller cloud providers and former UK Cloud employees have raised concerns that credits offered to customers by the largest cloud providers are anti-competitive and act as a barrier to entry and expansion.[854]

4.105 Civo said that:

(a) The largest cloud providers have the 'deepest pockets' which allows them to offer significant cloud credits to customers, which act as a loss leader.

(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.

(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 provider offered a customer of another small cloud provider $500k to transfer its business.

(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.

(e) Customers are unlikely to migrate back to the smaller providers after they have utilised the cloud credits offered by the largest providers, as they are 'locked-in' to the large cloud providers and it is costly and time consuming for customers to switch.[855]

4.106 A different small cloud provider said AWS and Microsoft offer $100,000 and $150,000 in cloud credits to start-ups and this effectively means that smaller providers can no longer compete with the larger providers for this business.[856]

4.107 The former UK Cloud employees said:

(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.

854 Note of meeting with [✂]; [✂] submission to the CMA [✂]; 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. Open Cloud Coalition members include Google Cloud.
855 Note of meeting with Civo [✂].
856 [✂] submission to the CMA [✂].

(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.[857]

4.108 The Startup Coalition,[858] 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.[859]

4.109 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).[860]

(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.[861]

(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.[862]

(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.[863]

(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

[857] Former UKCloud employees response to issues statement.
[858] Google is a supporter of the Startup Coalition and offers cloud credits. Source: https://startupcoalition.io/join-us/ accessed on 7 October 2024.
[859] Startup coalition's response to our Issues Statement: Startup Coalition response to the CMA's Issues Statement
[860] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 6.1.17.
[861] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 3.2.5.
[862] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 3.3.3.
[863] 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.[864]

(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.[865] Cloud providers offer cloud credits to assist enterprise customers migrating to their cloud.[866]

Our assessment

4.110 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.[867]

4.111 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.[868] Similarly, the Jigsaw research noted that some start-up customers planned to migrate away from their current cloud provider once they had used their cloud credits.[869]

4.112 While customers view cloud credits positively, we recognise that cloud credits could have a negative impact on smaller cloud providers who submitted that they are unable to match the significant amount of cloud credits offered by large cloud providers and that this impacts their ability to compete with them particularly for start-up customers. They said that this acts as a barrier to their expansion.

4.113 The value of discounts offered to customers through cloud credits is relatively small compared to cloud providers' other types of discounting: 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.[870]

4.114 Our provisional view is that cloud credits benefit smaller customers, particularly start-ups and that they are small in value when compared to other discounts

[864] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 6.1.18.
[865] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 6.1.19.
[866] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 6.1.19.
[867] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 3.2.5.
[868] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 6.1.18.
[869] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 3.5.5.
[870] 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.

offered by cloud providers such as committed spend agreements and reserved instances. We consider that the impact of cloud credits is not sufficiently material to create a barrier to entry or expansion or result in harm to competition in the cloud services markets.

4.115 However, any harm arising from pricing and discounting by large cloud providers could become a potential commercial barrier for smaller providers in the future.

Reputational barriers

4.116 One small cloud provider said that there are reputational barriers to entry and expansion as customers perceive the large cloud providers to be superior and are seen as a safe choice for a Chief Information Officer when choosing a cloud provider.[871]

4.117 Some customers also commented that they consider the large cloud providers to be more credible and capable than smaller cloud providers.[872]

4.118 The Jigsaw report states: 'The main providers are seen as AWS, Microsoft 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.'[873]

4.119 It also said that AWS is commonly seen by customers as a 'first mover'. There was a sense that other public cloud providers are catching up but there is still a reputational advantage.[874]

Our assessment

4.120 Our provisional view is that there may be some reputational barriers to entry and expansion in the cloud services markets.

Regulatory barriers

4.121 In this section we consider barriers to entry and expansion that could arise from regulations and from public sector procurement.

4.122 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

[871] Note of meeting with [✂].
[872] Notes of meetings with [✂]; [✂] response to the CMA's information request [✂].
[873] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 1.3.6.
[874] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 3.4.10.

Network and Information Systems Regulations) and data privacy (for example UK General Data Protection Regulation).[875]

4.123 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.[876]

4.124 Generally, cloud providers were also unable to provide accurately the costs of compliance and how this might differ globally and in the UK.

4.125 Many cloud providers,[877] commented on their overall regulatory burden, including European Union legislation (most notably the Data Act),[878] 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.

4.126 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.[879]

(b) All providers recognised the tailored offerings they were required to offer financial services firms.[880]

4.127 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.[881]

(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

---

[875] AWS' response to the CMA's information request [✂]; Google's response to the CMA's information request [✂]; IBM's response to the CMA's information request [✂].
[876] Google's response to the CMA's information request [✂].
[877] Responses to the CMA's information requests [✂].
[878] Data Act - Shaping Europe's digital future, accessed on 21 March 2024.
[879] 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.
[880] Responses to the CMA's information requests [✂].
[881] Data centres to be given massive boost and protections from cyber criminals and IT blackouts, accessed on 7 October 2024

Third Party (CTP) Services provided to the financial sector.[882] 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.

(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 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.[883] .

(d) The Information Commissioner's Office oversees compliance with UK data protection law,[884] and will consult on updated draft guidance on cloud computing in early 2025.[885]

Our assessment

4.128 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

4.129 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.[886]

4.130 We have received submissions and evidence raising issues around public sector procurement policy and implementation.[887]

[882] PS24/16: Operational resilience: Critical third parties to the UK financial sector - FCA, accessed on 19 November 2024.
[883] CP26/23 - Operational resilience: Critical third parties to the UK financial sector - Bank of England, accessed on 14 February 2024.
Cost benefit analysis, accessed on 7 November 2024.
[884] Cloud computing guidance for organisations, accessed on 22 October 2024.
[885] Our plans for new and updated guidance - ICO.
[886] More detail is included in Appendix K.
[887] 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; [✂] response to the CMA's information request [✂].Open Cloud Coalition submission to the CMA dated 18 December 2024. Open Cloud Coalition members include Google Cloud.

4.131 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.[888] 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.[889]

4.132 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 this.[890] In particular its latest policy Cloud Compute 2 requires competitive tendering.[891] In fact, most cloud providers we contacted said public sector customers already tender competitively).[892]

4.133 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'.[893] However, overall, we have not seen any evidence that public sector procurement practices are harming competition.

4.134 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 there is some evidence that their joint share may be even higher in the public sector. We consider that their leading positions amongst public sector customers are likely driven by the features we have identified elsewhere in this provisional report, and therefore some of the concerns raised with us will be addressed in our wider assessment.

4.135 While we have provisionally not found any harm to competition arising from public sector procurement, we propose suggesting that the Cabinet Office and the Crown Commercial Service continue to improve data collection on public sector procurement of cloud services so that outcomes can be better assessed, and any risks that public procurement strategies act as barriers to entry and expansion to smaller cloud providers or act to increase concentration in cloud services are visible and can be acted upon.[894]

4.136 We also propose suggesting that the UK government should promote and share best practice, and address inconsistencies in public bodies' cloud procurement

---

888 [✂] submission to the CMA [✂].
889 Note of meeting with [✂].
890 One Government Cloud Strategy and Cloud First policy.
891 Note of meeting with [✂].
892 Responses to Ofcom's information requests [✂].
893 NHS Cloud Strategy adoption plan - NHS England, Digital Cloud Strategic Roadmap for Defence.
894 This would be consistent with the NAO's recent report on efficiency in government procurement, that recommended improvements to data collection and making better use of that data to improve government decision-making; National Audit Office, Efficiency in government procurement of common goods and services, accessed 4 September 2024. Efficiency in government procurement of common goods and services.

strategies where necessary as public sector use of cloud services continues to grow.

Provisional conclusions

4.137 We have provisionally found substantial barriers to entry and expansion in the provision of cloud services, in particular for IaaS.

4.138 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, whereby the larger providers have comparatively lower ongoing costs.

4.139 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. This is disincentivising IaaS market entry and expansion.

4.140 Furthermore, given the scale of investment and expansion that large cloud providers, have made to date in IaaS, any new entrant would need to invest substantially more than the large existing suppliers in order to close the gap in a timely way. The levels of investment that AWS and Microsoft are expecting to make in the coming years may raise these barriers even higher.

4.141 While we recognise that 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.

4.142 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 and ISVs value access to a wider user base.

5. Barriers to switching and multi-cloud

- This chapter presents our assessment of potential barriers that cloud services customers face when switching provider and/or using multiple clouds. We assess providers' incentives to facilitate switching and multi-cloud and the specific roles played by technical barriers and egress fees.
- Customers face both commercial and technical barriers when seeking to multi-cloud or switch their cloud provider and many think that the costs of switching or using multi-cloud outweigh the benefits.
- While some customers are able to multi-cloud due to their willingness to adopt a workaround, or because their desired integration between clouds was minimal or switch if they are willing to invest sufficient time and money to do so, technical barriers affect many customers' ability to multi-cloud and switch.
- We have looked at technical barriers to switching and multi-cloud. Technical differentiation of features and interfaces in core and ancillary cloud services means that customers cannot easily compare or substitute products. This harms customers' ability to multi-cloud and/or switch clouds. Further disincentives include latency, a lack of transferable skills across clouds and insufficient transparency on how to mitigate or overcome technical barriers.
- There are some mitigations to these technical barriers from customers and cloud providers, but these only partly mitigate the technical barriers to multi-cloud and switching experienced by customers or their effect on competition.
- We have considered whether the charging of 'egress fees' for transferring data for the purposes of switching and/or multi-cloud harms competition. The presence and relevance of these fees reduces the ability of, and/or incentives for, customers to switch or multi-cloud; they may also reduce the incentives of suppliers to compete for their rivals' customers. We have not seen sufficient evidence that these fees fund investment and innovation, nor is it clear that they result in passed-on cost savings to customers or deter inefficient egress usage by them.
- We have provisionally found that technical barriers and egress fees are both barriers to switching and multi-cloud.

Introduction

5.1 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.[895] This also applies to the use of multiple providers (ie multi-cloud) and that represents an important source of competitive pressure on providers.

5.2 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.[896]

5.3 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.

5.4 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.

5.5 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 switching or multi-cloud are too high for some customers or for certain workloads to be moved to other providers.

5.6 We consider that the costs of switching and/or multi-cloud warrant close scrutiny as they directly affect customers' ability and incentives to choose the provider(s) that best fits their needs. As we have seen in the Competitive landscape chapter the supply of cloud services is concentrated, levels of switching are low and multi-cloud is adopted to varying degrees. This can lead to market power being enjoyed by a small number of providers.[897] In this context, customers' ability to switch and multi-cloud is particularly important.

5.7 The CMA may consider a wide range of potential barriers to switching or using multiple providers, including inconvenience, administrative obstacles, the presence

[895] CC3 (Revised), Guidelines for market investigations: Their role, procedures assessment and remedies, paragraph 296.
[896] CC3 (Revised), paragraphs 297 and 316.
[897] CC3 (Revised), paragraph 189.

of network effects locking customers into existing standards and large one-off costs or investments required to switch.[898]

5.8 In this market investigation we have focused on technical and commercial barriers to switching and multi-cloud. In particular, we have looked at:

(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,[899] 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.

5.9 Before assessing these barriers in turn, we consider the providers' overarching submissions on their incentives to facilitate switching and multi-cloud.

Providers' incentives to facilitate switching and multi-cloud

Evidence from providers

5.10 AWS, Microsoft, Google all told us that they have incentives to enable switching and/or customers' use of multi-cloud.[900]

(a) AWS submitted that IT providers are heavily commercially incentivised to ensure that customers are able to multi-cloud and switch between different IT providers, if and when they wish to do so. It said that limiting customers' ability to switch workloads or multi-cloud by creating technical barriers could result in a customer moving all their workloads to another IT provider because, since IT providers compete on a workload-by-workload basis, if a customer cannot host a third-party service on AWS or cannot have an AWS service interoperate successfully with a third-party service it wishes to use, the customer will simply switch the workload away from AWS to another IT provider (or choose another IT provider for the specific workload in the first place).[901]

(b) AWS submitted to Ofcom that because customers use multiple different options to meet their varying IT needs, AWS must make it easy for customers to migrate all or part of their workloads on and off AWS' services.[902] AWS

[898] CC3 (Revised), paragraph 318.
[899] Latency measures the time it takes to transfer data between public clouds.
[900] 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.
[901] AWS' response to the CMA's updated issues statement and working papers, paragraphs 13 and 15.
[902] AWS' response to Ofcom's information request [✂].

told us it has invested heavily in containerisation technology,[903] despite containers making it easier for customers to switch workloads away from AWS, as customers expect to use such technology in conjunction with the IT provider of their choice, both to move applications to and from AWS' infrastructure.[904] We explore more on this topic in the Technical barriers section below.

(c) 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.[905]

(d) 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 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. [906]

(e) 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.[907]

(f) 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).[908]

[903] Containerisation is a software system that isolates a software application, separating it from the operating system. It can be used to help ensure applications are portable, by bundling them with their dependencies, or to separate multiple applications running on the same.
[904] AWS' response to the CMA's updated issues statement and working papers, paragraph 15.
[905] AWS' response to the CMA's information request [✂].
[906] AWS' response to the CMA's Technical barriers working paper dated 06 June 2024, paragraph 75.
[907] Microsoft's response to the Issues Statement, 17 October 2023, paragraphs 26.
[908] Microsoft's response to the Technical barriers working paper dated 06 June 2024, paragraphs 21-25.

(g) Google submitted that, as a ‘challenger’ cloud provider, it is in Google’s interest to make switching and multi-clouding of workloads as easy as possible for potential, new and existing customers. It submitted that it offers the same benefits and support to any customer needing help with their migration irrespective of (i) customer size or industry; (ii) whether they are migrating all or part of their workloads to Google Cloud; or (iii) whether they are moving to Google Cloud from another cloud provider or are using cloud services for the very first time.[909]

5.11 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.[910]

(b) [✂] 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.[911]

5.12 Providers also made submissions in relation to the influence of open-source associations on providers’ incentives to facilitate switching and multi-cloud. We assess these in the section on technical barriers below.

[909] Google’s response to the CMA’s information request [✂].
[910] 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.
[911] [✂] submission to the CMA [✂].

Our assessment

5.13 When providers consider whether to reduce barriers to switching and multi-cloud, they face potential benefits and costs. These benefits and costs are acknowledged by our guidelines.[912] 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.

5.14 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.[913]

5.15 The size of these benefits and costs may be affected by various factors. For instance, 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 in terms of their existing customer base. In contrast, the costs of reducing switching barriers for existing customers may be more significant for large providers.

5.16 Moreover, technologies that facilitate migration from traditional IT may also be used to facilitate switching or multi-cloud between public cloud providers:

(a) 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.[914] This adds complexity to the incentives of the providers:

5.17 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.

5.18 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.

5.19 AWS and Microsoft have also submitted examples of where they have responded to specific customer requests for features related to interoperability and portability.[915]

912 CC3 (Revised), paragraph 317.
913 CC3 (Revised), paragraph 296.
914 [✂] response to Ofcom's information request [✂].
915 AWS' response to the CMA's information request [✂]; Microsoft's response to the CMA's information request [✂].

5.20 However, we observe that many such customer requests were 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 any benefits.

5.21 We therefore consider that this evidence does not necessarily demonstrate a strong incentive for AWS and Microsoft to facilitate interoperability and portability.

5.22 In general, the benefit to a cloud provider of making it easier for customers to switch may not confer a competitive advantage or result in increased sales if rivals replicate these changes. So providers may recognise that their efforts to reduce switching costs may increase customers' bargaining power and reduce their profitability. In concentrated markets, firms may recognise this, reducing their incentive to proactively reduce switching costs.

5.23 Overall, we consider that cloud providers' incentives to facilitate switching and multi-cloud are not clear-cut, and that this applies to both workloads migrating to the cloud from other IT models and cloud-native workloads.

5.24 We have instead assessed the extent to which switching barriers exist and the significance of those switching barriers on customers' behaviour and competition.

5.25 The rest of the chapter considers barriers to switching and multi-cloud related to:

(a) technical barriers; and

(b) egress fees.

Technical barriers

Introduction

5.26 This chapter presents our analysis of the potential impact of technical barriers on public cloud customers' ability and incentive to switch between clouds and/or use multiple clouds.

5.27 We consider that there are a number of sources of technical barriers: they can be grouped into those that relate to the way cloud services are designed and those that do not. We note that over time, new sources of technical barriers may emerge or become apparent.

5.28 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.

5.29 Other technical barriers include:

(a) latency: the time it takes to transfer data between public clouds. 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.

5.30 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.

5.31 Broadly, customers may incur two types of additional costs when trying to use multiple clouds for their workloads. We term these 'multi-cloud costs'. They may include:

(a) Operational costs: Customers may incur additional costs when operating workloads on more than one public cloud, independently 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 may incur additional costs when enabling workloads on multiple public clouds to communicate. This may include setting up connections over a network using service APIs, as well as any ongoing management and operation of any integrations.

5.32 Customers may also incur various costs when switching workloads that exist on one public cloud (origin cloud) to another (target cloud). We term these 'switching costs'. They may include:

(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 will likely also incur temporary multi-cloud costs as set out above.

5.33 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 5.1:The costs associated with different customer activities

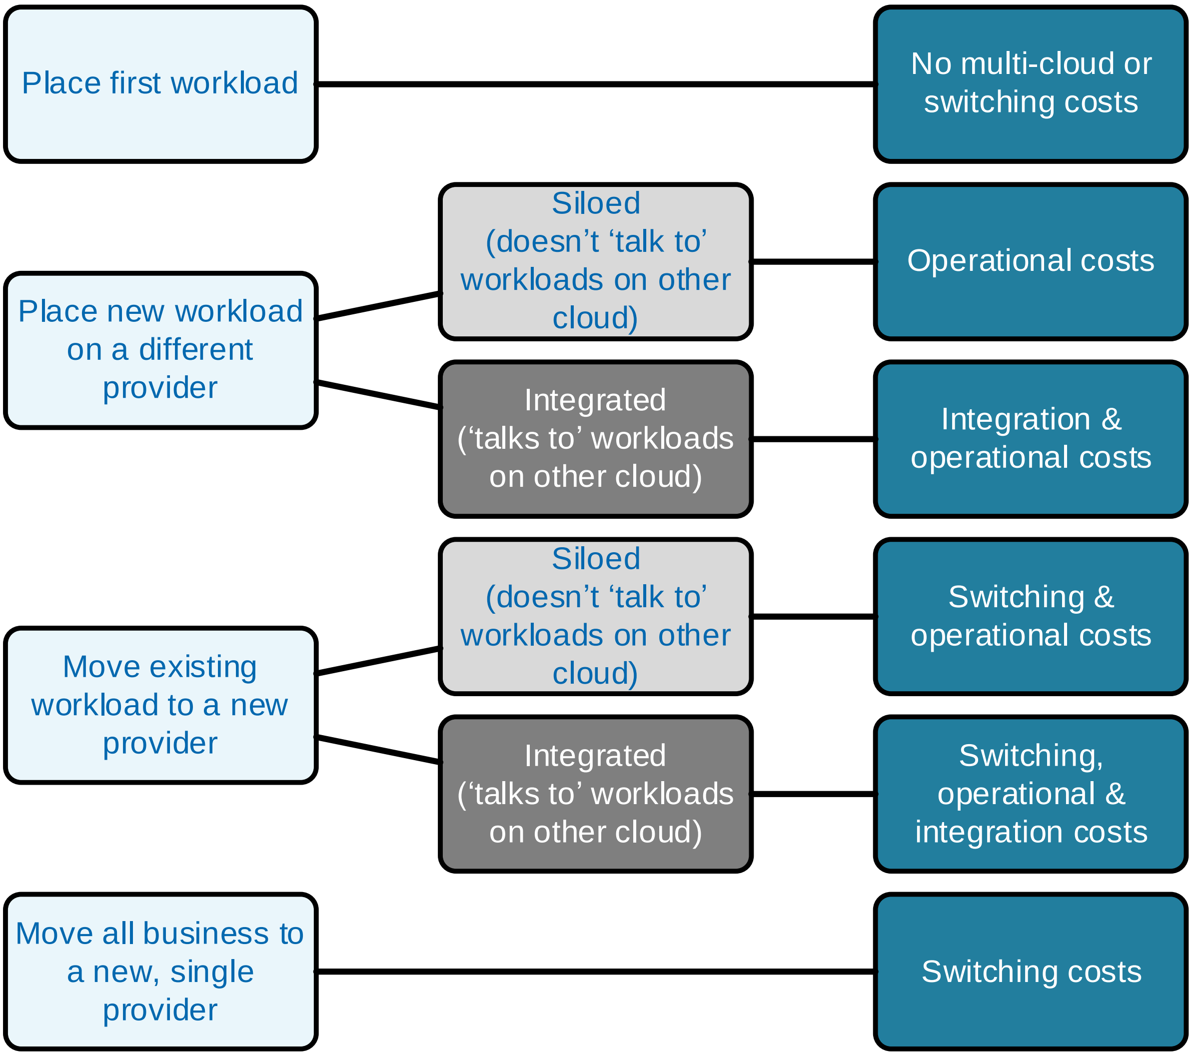

Source: CMA

5.34 This section sets out:

(a) our conceptual framework for assessing the potential impact of technical barriers on customers' ability to switch and multi-cloud; and

(b) our analysis of:

(i) the existence of technical barriers to switching and multi-cloud and the impact of these barriers on customer behaviour;

(ii) the technical barriers to switching and multi-cloud separately in relation to:

(1) core services;[916]

(2) ancillary services and tools;[917]

(3) other technical factors; and

(iii) technical mitigations that may reduce the technical barriers to switching and multi-cloud.

Conceptual framework

5.35 We have considered whether, and to what extent, technical barriers prevent or restrict the ability of customers to:

(a) adopt and use a multi-cloud architecture, particularly integrated multi-cloud; and/or

(b) switch, particularly between public clouds.

5.36 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 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.

(c) Customers who migrated their workloads from on-premises may also experience technical barriers differently to customers who built their

[916] 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.
[917] 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.

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.

5.37 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.[918]

5.38 Even if technical barriers only affect a proportion of customers, this could still have a significant effect on competition overall. We have therefore considered the potential effects of all technical barriers, regardless of whether they impact all customers. We also considered the cumulative impact of several potential costs or points of friction that may arise from different aspects of a customer's workload and its operation in the cloud.

5.39 We received responses from cloud providers which set out their view that technical barriers, or most technical barriers are inherent or intrinsic, arising naturally from innovation and competition: AWS and Microsoft 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 in the ways we describe is unrealistic.[919]

5.40 Our view is that some technical barriers may be inherent or intrinsic,[920] 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 Latency section 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 Transparency section below).

---

[918] Responses to the CMA's information requests [✂]. [✂].
[919] 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
[920] We use the term 'intrinsic' rather than 'inherent' to match our guidance. See for instance CC3 (Revised), paragraph 320.

5.41 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.

5.42 Instead, we have sought to assess the related question of whether the mitigations currently available to customers can address the technical barriers, given their potential costs and benefits, and whether the effect of any technical barriers that may to some extent be intrinsic to cloud services can nevertheless be removed or reduced.

5.43 Moreover, we consider that 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.

5.44 We also do not consider it appropriate or necessary to ascertain the motive behind these design choices,[921] as our focus is instead on their effect and whether the barrier constitutes a feature of the market that affects competition.

5.45 We consider that, where certain features of cloud services may be intrinsic to the extent that the relevant market cannot realistically be envisioned without them, we can nevertheless consider what effect these have in shaping competition. This is the approach we have taken in our assessment.

Technical barriers to switching and multi-cloud

5.46 We have considered evidence on the existence of technical barriers to switching and multi-cloud where we have been provided evidence applicable to both.

5.47 We have focused on customers' experiences of the technical aspects of attempting to switch between clouds and/or use multiple clouds. Evidence from other stakeholders, such as cloud providers, ISVs and industry bodies is also included where relevant.

Overall customer views on switching and multi-cloud

5.48 The Jigsaw research amongst customers found that most participants 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

[921] 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.

infrastructure.[922] Many participants 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.[923]

5.49 The Jigsaw report found that participants 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.[924]

5.50 The most common response from participants, 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 the companies' 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.[925]

5.51 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.[926] Jigsaw also found that there are two key factors that make it more difficult to overcome the technical barriers: a skills gap and a resource gap.[927]

5.52 A report from a market research firm states that a multi-cloud architecture substantially increases complexity and creates operational challenges noting: '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.'[928]

5.53 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.[929] 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

[922] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 4.1.2.
[923] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 4.1.4.
[924] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 4.1.3.
[925] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 4.1.8.
[926] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 4.2.1.
[927] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 4.3.
[928] [✂] response to the CMA's information request [✂].
[929] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 4.4.1.

participants to invest this time, effort and resources to switch clouds or implement a multi-cloud architecture.[930]

5.54 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.[931] 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.[932]

Multi-cloud costs

5.55 In this section, we consider 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 also consider any operational costs relating to these architectures.

Multi-cloud operational costs

5.56 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.[933] 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 difficulty in building 'high availability' architectures, decreased scalability and resilience; more patterns and standards to maintain; and increased effort for governance and audit.[934]

5.57 This customer estimated that just the direct cost of maintaining two clouds added an overhead of approximately 5% to its total cloud spend.[935] 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.[936]

Multi-cloud integration costs

5.58 The evidence relating to the costs that customers incur when integrating multiple public clouds is mixed, consistent with our view that customers are heterogenous.

930 Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 4.5.9.
931 Responses to the CMA's information requests [✂]; Notes of meetings with [✂]; [✂] submission to Ofcom [✂].
932 Responses to the CMA's information requests [✂].
933 Responses to the CMA's information requests [✂]; Notes of meetings with [✂].
934 [✂] response to the CMA's information request [✂].
935 [✂] response to the CMA's information request [✂].
936 Note of meeting with [✂].

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.[937]

5.59 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.[938] 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.[939]

5.60 Responses from other customers showed that they experienced few barriers to integration across multiple public clouds.[940] Reasons given for this included the availability of open APIs[941] 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.

5.61 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.[942] 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.[943]

5.62 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.[944] 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.[945]

5.63 Some customers said that multi-cloud is a part of their business strategy,[946] and others said whilst they do not currently have a use-case for integrated multi-cloud, they may do so in the future.[947] 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.[948]

---

[937] Note of meeting with [✂].
[938] Responses to the CMA's information requests [✂]; Notes of meetings with [✂].
[939] Responses to the CMA's information requests [✂]; Note of meeting with [✂].
[940] Responses to the CMA's information requests [✂]; Notes of meetings with [✂].
[941] An open API is a free, publicly available application programming interface (API) that allows developers to access software and data.
[942] Responses to the CMA's information requests [✂]; Note of meeting with [✂].
[943] Responses to the CMA's information requests [✂].
[944] Between applications: Responses to the CMA's information requests [✂]; Notes of meetings with [✂]. Within applications: Responses to the CMA's information requests [✂]; Notes of meetings with [✂].
[945] Note of meeting with [✂].
[946] Responses to the CMA's information requests [✂].
[947] [✂] response to the CMA's information request [✂]; Note of meeting with [✂].
[948] [✂] response to the CMA's information request [✂];

5.64 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.[949] 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.[950]

5.65 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.[951] 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.

5.66 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.[952]

Cloud Providers' views on the costs of multi-cloud

5.67 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-native[953] 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.[954] 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.[955]

(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

949 Notes of meetings with [✂]; [✂] response to the CMA's information request [✂].
950 Note of meeting with [✂].
951 Note of meeting with [✂].
952 [✂] response to the CMA's information request [✂].
953 'Cloud-native workloads' refers to workloads created on the cloud, not migrated from on-premises.
954 Google's response to the CMA's information request [✂].
955 Google response to the Egress Fees Working paper, paragraph 44.

to ensure multi-cloud is accessible and free from anticompetitive barriers at all stages of a company's IT modernisation timeline.[956]

5.68 However, AWS and Microsoft 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. 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.[957]

(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.[958] 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; increased risk of downtime.[959] It said however that such customers do have credible alternatives to switch to, should they need.[960]

5.69 Microsoft also 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.[961]

5.70 Additionally Microsoft said that, despite the fact that customer research shows that the core services of AWS, Microsoft and Google are all roughly equivalent (as each innovates and the others catch up), the CMA's working paper creates a hypothesis that there are 'best of breed' features that customers want to mix and match, whether in siloed or integrated multi-cloud applications.[962]

956 Oracle Response to Working Papers and updated Issues Statement, pages 3 and 4.
957 AWS' response to the CMA's information request [✂].
958 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.
959 Microsoft response to the Technical Barriers working paper, paragraphs 50.
960 Microsoft response to the Technical Barriers working paper, paragraphs 44.
961 Microsoft response to the Technical Barriers working paper, paragraphs 49.
962 Microsoft response to the Technical Barriers working paper, paragraph 27.

Switching costs

5.71 In this section, we consider the potential switching costs for customers moving workloads from one public cloud to another.

5.72 Many customers anticipate or experience significant costs to switch public clouds.[963] 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’,[964] or in relative terms - eg some customers described technical barriers as the main barrier to switching.[965]

5.73 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.[966]

5.74 Some of these customers indicated that these costs had stopped them from switching or considering switching.[967] 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.[968] This customer also said that a switch would cost a similar amount to the initial migration.[969]

5.75 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.[970] The Jigsaw report also found this.[971]

5.76 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 revenue 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.[972]

[963] Responses to the CMA’s information requests [✂]; Notes of meetings with [✂].
[964] Note of meeting with [✂].
[965] Responses to the CMA’s information requests [✂]; Notes of meetings with [✂].
[966] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 1.3.13.
[967] Notes of meetings with [✂].
[968] Note of meeting with [✂].
[969] [✂] response to the CMA’s information request [✂].
[970] Responses to the CMA’s information requests [✂]; Note of meeting with [✂].
[971] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 4.5.8.
[972] Note of meeting with [✂].

5.77 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.[973]

Operational concerns arising from switching

5.78 We have also seen a range of operational concerns that arise when customers consider switching between public clouds.[974]

(a) For example, DLG, a 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.[975]

(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.[976]

5.79 These views are consistent with the findings of the Jigsaw report.[977]

Cloud providers' views on technical barriers to both switching and multi-cloud

5.80 AWS, Microsoft 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, 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 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.[978]

(b) AWS said that its customer feedback indicates that a switch between cloud services providers would cost a similar (if not less) amount to the initial migration from on-premises IT infrastructure to a cloud services provider. AWS' view is that this further supports that these costs are inherent to IT

[973] Centerprise's response to the CMA's information request [✂].
[974] Responses to the CMA's information requests [✂].
[975] DLG's response to the CMA's information request [✂].
[976] [✂] response to the CMA's information request [✂].
[977] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 4.2.1.
[978] 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.

services, despite the mitigation efforts undertaken by cloud services providers.[979]

(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.[980]

(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.[981] It also said that the Egress fees working paper and Jigsaw report reinforce the findings of Ofcom's market test that technical and licensing challenges are the main barriers to switching.[982]

5.81 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.[983]

(b) IBM said that it agrees that customers face technical challenges related to switching [✂]. IBM said that portability is key and that it allows customers access to the best services for their use cases and preserves cloud providers' incentives to innovate.[984] 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 [✂].[985]

[979] AWS' response to the Updated issues statement and working papers dated 23 May 2024 and 6 June 2024, paragraph 63.
[980] Microsoft's response to the Technical barriers working paper dated 06 June 2024, paragraph 1.1.
[981] Google's response to the Technical barriers working paper dated 06 June 2024, paragraphs 4-5 and 9.
[982] Google's response to the Egress fees working paper dated 23 May 2024, paragraph 36.
[983] Oracle's response to the Updated issues statement and working papers dated 23 May 2024 and 06 June 2024, pages 3 and 4.
[984] IBM's response to the Technical barriers working paper dated 06 June 2024, page 1.
[985] IBM's response to the Technical barriers working paper dated 06 June 2024, page 1.

(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.[986]

5.82 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.[987]

(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. 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.[988]

(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.

(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.[989] 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.[990]

(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

[986] Civo's response to the CMA's information request [✂].
[987] AWS' response to the CMA's information request [✂].
[988] AWS' response to the Updated issues statement and working papers dated 23 May 2024 and 6 June 2024, paragraph 64.
[989] Microsoft's response to the Issues Statement dated 17 October 2023, paragraphs 23-25.
[990] Microsoft's response to the Technical barriers working paper dated 06 June 2024, paragraph 7 (a).

better performance, lower costs, higher security and a smaller environmental footprint for similar services. It also said that it is differentiated from other cloud providers in part because it optimises for speed and performance.[991]

(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'.[992]

5.83 Other cloud providers expressed different views despite the benefits that some technical differentiation can have.

(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.[993]

(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.[994]

(c) OVHcloud said that the justification for technical differentiation of cloud infrastructure services and cloud ancillary services through proprietary technologies is less clear.[995]

5.84 Some cloud providers said that the evidence on technical barriers[996] was mixed, and shared the view that this is because the heterogenous nature of cloud customers.

(a) AWS said that the mixed and contradictory views set out in the working paper 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.[997]

(b) Google said that, whilst some customers have described potential challenges to switching, other than in respect of IAM services (where they have 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

[991] Oracle's response to the Issues Statement dated 17 October 2023, pages 1-2.
[992] IBM response to the Technical Barriers working paper, page 3.
[993] Google's response to the Issues Statement dated 17 October 2023, paragraphs 3, 10-11 and 18-19.
[994] IBM's response to the Issues Statement dated 17 October 2023, paragraph 3.2.
[995] OVHcloud's response to the Issues Statement dated 17 October 2023, page 4.
[996] As set out in the Technical Barriers Working Paper.
[997] 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.[998]

(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.[999]

5.85 Microsoft said that the CMA's evidence of harm is anecdotal, selective and insufficiently conclusive. In particular, it said that the evidence presented in the working paper is almost entirely anecdotal and is often 'mixed' as to the impact of the identified barriers. Microsoft said that the CMA relies almost exclusively on interviews it conducted directly with customers, unweighted for the customer's technical sophistication, cloud spend, where on the spectrum it is between cloud-native, recent cloud migrant and long-time cloud user, whether and how long ago it attempted or seriously considered switching or multi-cloud. Microsoft also said that the CMA's evidence is a backward-looking static snapshot that ignores the iterative dynamism of cloud services.[1000]

5.86 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.[1001]

Our assessment

5.87 The way that cloud services have developed in the past, as well as elements of competition and innovation, may have led to some differentiation in cloud services. Whilst these differences may have associated benefits for customers they also create technical barriers and we consider what effect these have on competition.

5.88 Furthermore, even if cloud services 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. 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 innovative new products and services which are only available from some suppliers.

5.89 In relation to Microsoft's views, 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:

[998] Google response to the Technical Barriers Working Paper dated 06 June 2024, paragraphs 4-5, 9.
[999] IBM's response to the Technical barriers working paper dated 06 June 2024, page 1.
[1000] Microsoft's response to the Technical barriers working paper dated 06 June 2024, paragraphs 52-54.
[1001] Microsoft's response to the Issues Statement dated 17 October 2023, paragraphs 26-32.

(a) we 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) our qualitative findings that we gathered from customers directly are consistent with the Jigsaw report and other market research, which together represent the experience of a large number of customers; 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. We note that our qualitative approach also provided significant benefits, such as the ability to get a deeper understanding of stakeholders' views.

5.90 We consider that it is possible for it to be easier for customers to switch to cloud services than away from them, regardless of whether the services are designed for 'lopsided portability'.

5.91 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.

5.92 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.

5.93 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.

5.94 Due to customers' 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.

5.95 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.

5.96 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 that some customers would not switch even if they were forgoing significant benefits in failing to do so.

5.97 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

5.98 We have considered whether technical differentiation of the features and interfaces of core services is itself a technical barrier to switching and multi-cloud.

5.99 As described above, core services are the main IaaS and PaaS services that contribute to the key objectives of the workload.

5.100 Differentiation of features are differences in the functionality of similar core services hosted on different public clouds, and differentiation of interfaces are differences in the interfaces (eg protocols or APIs) of comparable core services hosted on different public clouds and provided by any supplier, including cloud providers and ISVs.

Cloud providers' views

5.101 All cloud providers that commented on this topic, including AWS and Microsoft, said that there are differences in the features and interfaces of core services.[1002]

5.102 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 providers, customers choose their specific IT providers because of such differences.[1003]

5.103 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

[1002] 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; OVH Cloud'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 [✂]; Responses to the CMA's information requests [✂].
[1003] AWS' response to the Technical barriers working paper dated 06 June 2024, paragraphs 16, 19 and 22.

is easier than moving PaaS applications, the CMA overestimates the difficulty of moving PaaS applications from one cloud services provider to another.[1004]

5.104 AWS said that the feature and interface differentiation identified in the working paper 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.[1005]

5.105 AWS also said that the evidence on differentiation of features and interfaces set out in the Technical barriers working paper is mixed and cannot support a conclusion that such differentiation results in any AEC. It said that the working paper itself includes several pieces of feedback which illustrate that feature differentiation and interface differentiation do not raise any concerns.[1006]

5.106 As set out above, other cloud providers, including Microsoft and Google, made general comments on the inherent nature of technical barriers and on the evidence of harm from technical barriers which are applicable to core services.

5.107 AWS also said that none of the evidence included in our working paper points towards a need for regulatory intervention with respect to feature or interface differentiation, as evidence shows that feature and interface differentiation does not have a significant impact on customers' ability to switch or use multiple clouds because sufficient mitigations exist.[1007]

5.108 Other cloud providers, including Microsoft and Google made more general comments on mitigations to technical barriers which are applicable to core services. We set out our assessment of technical mitigations later in this chapter.

Customers' views

5.109 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.[1008]

5.110 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.[1009] There were also some

[1004] AWS' response to the Technical barriers working paper dated 06 June 2024, paragraph 63.
[1005] AWS' response to the Technical barriers working paper dated 06 June 2024, paragraphs 14 and 16.
[1006] AWS' response to the Technical barriers working paper dated 06 June 2024, paragraphs 17 and 20.
[1007] AWS' response to the Technical barriers working paper dated 06 June 2024, paragraphs 19 and 22.
[1008] Responses to the CMA's information requests [✂]; Notes of meetings with [✂].
[1009] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraphs 4.1.9.

customers that said that differences in features or interfaces of core services did not affect their ability to switch or multi-cloud.[1010]

5.111 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.[1011]

Differentiation of features

5.112 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.[1012] For example, Sky Betting and Gaming 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.[1013] 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.[1014]

5.113 Some customers also said that differences in features also increase the cost of using multiple clouds.[1015] For example, a customer said that there are circumstances where multiple providers offer identical solutions but that there are also situations where functional and non-functional differences exist which complicate the wider use of multiple clouds.[1016]

5.114 Customers were particularly concerned with the differentiation of features of PaaS services.[1017] For example 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.[1018]

5.115 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.[1019]

5.116 Customers also mentioned specific types of PaaS services which are particularly challenging to switch away from, such as data warehouse and analytics

[1010] Responses to the CMA's information requests [✂]; Notes of meetings with [✂].
[1011] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraphs 4.1.13; Responses to the CMA's information requests [✂]; Notes of meetings with [✂].
[1012] Responses to the CMA's information requests [✂]; Notes of meetings with [✂].
[1013] Sky Betting and Gaming's response to the CMA's information request [✂].
[1014] Responses to the CMA's information requests [✂].
[1015] Responses to the CMA's information requests [✂]; Notes of meetings with [✂].
[1016] [✂] response to the CMA's information request [✂].
[1017] Responses to the CMA's information requests [✂]; Notes of meetings with [✂].
[1018] [✂] response to the CMA's information request [✂].
[1019] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraphs 4.1.3.

services.[1020] 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.[1021]

Differentiation of interfaces

5.117 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.[1022] 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.[1023]

5.118 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.[1024] For example, a customer said that different interfaces would mean Infrastructure as Code[1025] and other software managing applications must be re-written when switching clouds.[1026]

5.119 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.[1027]

5.120 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.[1028] 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.[1029]

Other stakeholders

5.121 ISVs, suppliers of professional services and other stakeholders also said that there are differences in the features and of interfaces of core infrastructure services.

1020 Responses to the CMA's information requests [✂]; Notes of meetings with [✂].
1021 Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraphs 4.1.14 and 4.5.2.
1022 Responses to the CMA's information requests [✂]; Note of meeting with [✂].
1023 [✂] response to the CMA's information request [✂].
1024 Responses to the CMA's information requests [✂]; Note of meeting with [✂].
1025 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.
1026 Note of meeting with [✂].
1027 Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraphs 4.1.9 and 4.5.2.
1028 Responses to the CMA's information requests [✂]; Notes of meetings with [✂].
1029 [✂] response to the CMA's information request [✂].

Similar to customers, these other parties generally said that such differentiation hinders customers' ability to switch and multi-cloud.[1030]

Our assessment

5.122 The evidence described above shows that, 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.

5.123 Based on evidence from customers and other parties, we consider that differentiation of features of core services hinders at least some customers' ability to use multiple clouds.

5.124 The evidence more broadly shows 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.

5.125 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.

5.126 We also consider that 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.

5.127 In terms of the differentiation of interfaces of core services, we consider that it hinders 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 and required creating or using workarounds. This was also echoed in the Jigsaw report.

5.128 In addition, based on evidence from some customers and other parties, we consider that differentiation of interfaces of core services hinders at least some customers' ability to switch clouds.

[1030] [✂] submission to the CMA [✂]; Notes of meetings with [✂].

5.129 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.

5.130 We recognise that there are some technical mitigations to the current differentiation of features and interfaces of core services but the evidence shows that these mitigations only overcome the technical barriers to a limited extent. We discuss these below at Technical mitigations.

Ancillary services and tools

5.131 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. [1031]

5.132 Ancillary services and tools provide functions that support the operation and management of core services. They include IAM, billing, observability, and provisioning and orchestration.[1032]

5.133 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.

5.134 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

5.135 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 separately, without integration and automation, which could 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,

[1031] 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.
[1032] 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.

the customer may need to make additional changes to its applications, data and associated tools so that they can work and perform well on the target cloud.[1033]

Cloud providers' views

5.136 AWS and Google said that generally ancillary services and tools are differentiated.[1034]

5.137 AWS said that technical barriers related to differentiation of features and interfaces of ancillary services and tools are inherent to IT services (and cannot be resolved through regulatory intervention) and it does not support that there is any AEC requiring regulatory intervention.[1035]

5.138 Other cloud providers, including Microsoft and Google, made general comments (see above) on the inherent nature of technical barriers which are also applicable to ancillary services and tools.

5.139 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.[1036] It also said that 'monitoring is not necessarily a "dependency" that prevents containers from promoting portability'.[1037]

5.140 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.[1038]

[1033] 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.
[1034] 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.
[1035] AWS' response to the CMA's Technical barriers working paper dated 06 June 2024, paragraph 31.
[1036] AWS' response to the CMA's Technical barriers working paper dated 06 June 2024, paragraphs 40 and 41.
[1037] AWS' response to the CMA's Technical barriers working paper dated 06 June 2024, paragraph 40 and 66(b).
[1038] IBM's response to the Issues Statement, 17 October 2023, page 3; IBM's response to the CMA's information request [✂].

5.141 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.[1039]

5.142 A party said that when a customer uses proprietary ancillary services this can create lock in even if the customer uses otherwise portable services, such as IaaS services.[1040] It said that ensuring that access to cloud providers' services is protected and enhanced is essential to improving competition in the market. [✂][1041] [✂].[1042]

Customer views

5.143 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.[1043] For example:

(a) One customer said that differences in APIs of ancillary services and tools require reworking of deployment pipelines, tooling or code.[1044]

(b) Another customer told us 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.[1045]

(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.[1046]

5.144 However, some other customers told us that differences between ancillary services and tools do not necessarily impact their ability to switch or multi-cloud.[1047]

(a) For example, one customer told us that the differentiation of ancillary services neither makes it more difficult/costly to switch cloud provider nor

[1039] Note of meeting with IBM [✂].
[1040] Note of meeting with [✂]; [✂] submission to the CMA [✂].
[1041] See for example [✂].
[1042] [✂] submission to the CMA [✂].
[1043] Responses to the CMA's information requests [✂].
[1044] [✂] response to the CMA's information request [✂].
[1045] [✂] response to the CMA's information request [✂].
[1046] Note of meeting with [✂].
[1047] Responses to the CMA's information requests [✂].

more difficult/costly to integrate multiple public clouds (but did not elaborate).[1048]

Other stakeholder views

5.145 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,[1049] 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.[1050]

(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.[1051]

(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 the lock-in of developers to clouds, through the cost and time of retraining on these tools.[1052]

[1048] [✂] response to the CMA's information request [✂]. 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.
[1049] 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.
[1050] Note of meeting with [✂].
[1051] Ofcom note of meeting with [✂].
[1052] Note of meeting with [✂].

(d) Linux Foundation told us that that using proprietary services or add-on features can lead to barriers to switching and integrated multi-cloud.[1053]

5.146 In addition, market research states that [✂].[1054]

5.147 Some other parties 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.[1055]

(b) Another ISV said that cloud-agnostic provisioning tools can be used to simultaneously set up cloud services from multiple cloud providers.[1056]

(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.[1057]

Our assessment

5.148 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.

5.149 Based on evidence from customers and cloud providers, we consider that the differentiation of ancillary services and tools hinders customers' ability to switch and multi-cloud.

5.150 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).

5.151 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 can be used across different clouds.[1058] However, our evidence on technical mitigations shows that these are not fully effective. We discuss these below at Technical mitigations.

[1053] Note of meeting with the Linux Foundation [✂].
[1054] [✂] response to the CMA's information request [✂].
[1055] Okta's response to the CMA's information request [✂].
[1056] Note of meeting with [✂].
[1057] Note of meeting with the Linux Foundation [✂].
[1058] Note of meeting with [✂].

5.152 Below, we discuss IAM as a specific area of ancillary services and tools because the differentiation of IAM was a particular area of focus and concern among the evidence we received.

Identity and Access Management (IAM)

5.153 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.[1059]

5.154 As noted above, we heard from various stakeholders that IAM is especially important for customers who switch or multi-cloud, which is why we consider it separately here.

5.155 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.[1060] Customers are able to choose other IAM services and tools for additional functionalities and/or for multi-cloud purposes. 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.[1061]

Cloud providers' views

5.156 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.[1062]

5.157 AWS said that it supports a number of industry-standard IAM protocols which allows customers to store their identities 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.[1063]

[1059] 'What is Cloud Identity and Access Management?', accessed 2 May 2024.
[1060] Note of meeting with Okta [✂].
[1061] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 3.5.5 (b).
[1062] AWS' response to the CMA's updated issues statement and working papers dated 23 May 2024 and 6 June 2024, paragraph 66(b).
[1063] AWS' submission to the CMA [✂].

5.158 AWS shared a list of product feature requests made by customers in the last 12 months in relation to multi-cloud functionalities, and some of these requests involved technical challenges with IAM.[1064]

5.159 Google told us 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.[1065]

5.160 Another cloud provider 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.[1066]

5.161 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 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.[1067] 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.[1068]

5.162 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.[1069]

[1064] AWS' response to the CMA's information request [✂].
[1065] Google's response to the CMA's Technical barriers working paper dated 06 June 2024, paragraph 12.
[1066] [✂] response to the CMA's information request [✂].
[1067] Oracle's response to the CMA's information request [✂].
[1068] Oracle's response to the CMA's information request [✂].
[1069] Civo's response to the CMA's information requests [✂].

5.163 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.[1070]

5.164 Some cloud providers submitted specific comments on identity authentication as one of the main functions of IAM:

(a) AWS said that technical costs associated with identity authentication (and single sign-on) are low.[1071]

(b) Microsoft told us that most enterprises using multiple cloud services deploy some kind of single sign-on.[1072] Similarly, Microsoft said that the work needed for identity authentication is a scenario that is well understood by enterprises and commonly solved. It said that multiple IAM service providers compete on the quality and ease of implementation of their IAM services.[1073]

(c) We heard from cloud providers that, in general, they adopt standard IAM protocols such as SCIM, SAML and OICD,[1074] which facilitate this type of integration.[1075]

(d) 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.[1076] We understand that these nuances can add to the technical costs involved for customers who set up this type of integration.

Customer views

5.165 Many customers said that IAM is an important consideration and/or poses technical challenges when switching between clouds or adopting a multi-cloud architecture.[1077] Most of these customers said that the differentiation of IAM across clouds was either the main reason or a significant reason for these challenges.[1078]

5.166 Two customers, ASDA and WPP, suggested that for them IAM may not pose technical challenges to switching or multi-cloud of associated workloads between

---

[1070] Microsoft's response to the CMA's information requests [✄].
[1071] AWS' response to the CMA's Technical barriers working paper dated 06 June 2024, paragraphs 32-36.
[1072] Microsoft's submission to the CMA dated 8 August 2024, CMA Cloud Services MIR - Identity and Access Management Follow-up, page 2.
[1073] Microsoft's submission to the CMA dated 8 August 2024, CMA Cloud Services MIR - Identity and Access Management Follow-up, page 2.
[1074] System for Cross-domain Identity Management (SCIM) is a standard protocol that enables the management of user accounts, groups, and some level of 'access' across different digital (and cloud) services. Security Assertion Markup Language (SAML) and Open ID Connect (OIDC) are standard protocols that standardise the process of authenticating and authorising users when they sign in to access digital (and cloud) services.
[1075] [✄] submission to the CMA [✄]; Responses to the CMA's information requests [✄].
[1076] [✄] response to the CMA's information request [✄].
[1077] Responses to the CMA's information requests [✄]; Notes of meetings with [✄]; [✄] submission to Ofcom [✄].
[1078] Responses to the CMA's information requests [✄]; Notes of meetings with [✄]; [✄] submission to Ofcom [✄].

clouds, since they would be able to do so without also switching their current IAM service and the cloud it is hosted on.[1079]

5.167 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.[1080] 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.[1081]

(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.[1082]

5.168 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.[1083]

5.169 Some evidence from customers is specifically related to identity authentication as one of the main functions of IAM. From the evidence we obtained from customers[1084] 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.[1085]

5.170 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.[1086]

1079 ASDA's response to the CMA's information request [✂]; WPP's response to the CMA's information request [✂].
1080 Responses to the CMA's information requests [✂]; Notes of meetings with [✂].
1081 [✂] response to the CMA's information request [✂].
1082 Note of meeting with [✂].
1083 Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 4.1.9 (b).
1084 Responses to the CMA's information requests [✂]; Note of meeting with [✂]; [✂] response to Ofcom's information request [✂].
1085 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 [✂]).
1086 Responses to the CMA's information requests [✂]; Notes of meetings with [✂].

(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.[1087]

Other stakeholder views

5.171 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.[1088] 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.[1089]

(b) Another ISV, Okta, said that customers who use an IAM service offered by an ISV, such as Okta, are unable to substitute a cloud provider's 'native' IAM service completely. It 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. It also said that 'single sign-on' can simplify managing a multi-cloud architecture.[1090]

5.172 Market research 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, 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.[1091]

1087 Ofcom note of meeting with [✂].
1088 [✂] response to the CMA's information request [✂].
1089 Note of meeting with [✂].
1090 Note of a meeting with Okta [✂].
1091 [✂] response to the CMA's information request [✂].

Microsoft's IAM services and tools

5.173 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 services and tools and other Microsoft software.

5.174 The design and implementation of Active Directory may be relevant to customers' ability to switch and multi-cloud, for example in a situation where a customer is using a hybrid cloud architecture with an integration between a public cloud IAM service and an on-premises instance of Active Directory. 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.[1092]

Microsoft Entra ID

5.175 We have seen evidence of differences in how aspects of identity authentication (including so-called 'federation'[1093]) works for Entra ID compared to other cloud providers' IAM services, which in some cases can restrict customers' choice of primary IAM service:

(a) A cloud provider told us that Microsoft does not support inbound SCIM[1094] for Entra ID, which means that SCIM cannot be used to provision identities maintained primarily in another IAM service (such as Okta, Ping, or CyberArk) into Entra ID. It said that in this case, cloud providers or ISVs who develop IAM services or tools must build custom integrations (using Microsoft's Graph API) which are specific to Entra ID, costlier to build and increase friction for customer adoption.[1095]

(b) Another cloud provider told us that synchronisation between Entra ID and other IAM services is only available in one direction (from Microsoft to the other IAM service). It said that if a user is created in its IAM service [✂] then this cannot be automatically synced across to Microsoft, and as a result, Entra ID often remains a customer's primary IAM service.[1096]

(c) 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

[1092] Okta's response to the CMA's information request [✂].
[1093] 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).
[1094] As noted above, System for Cross-domain Identity Management (SCIM) is a standard protocol that enables the management of user accounts, groups, and some level of 'access' across different digital (and cloud) services.
[1095] Responses to the CMA's information requests [✂].
[1096] [✂] response to the CMA's information request [✂].

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.[1097]

(d) 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.[1098]

5.176 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.

5.177 It said [✂].[1099]

5.178 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.[1100] 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.[1101]

5.179 Microsoft also said that it is not correct that [✂], which is another cloud provider's IAM service, cannot replicate changes into Entra ID. It said that there is nothing inherent in the features or functionality of Entra ID that means the provisioning must be one-way, and instead it has developed this integration in a certain way because it addresses customers' needs in this specific scenario.[1102]

Technical integration between Entra ID and Active Directory

5.180 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.[1103]

5.181 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

1097 [✂] submission to the CMA [✂].
1098 [✂] response to the CMA's information request [✂].
1099 Microsoft's response to the CMA's information request [✂].
1100 Microsoft's response to the CMA's information request [✂].
1101 Microsoft's submission to the CMA dated 8 August 2024, CMA Cloud Services MIR - Identity and Access Management Follow-up.
1102 Microsoft's response to the CMA's information request [✂].
1103 Google's response to the Technical barriers working paper dated 06 June 2024, Section II.

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[1104]) 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.[1105]

5.182 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 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.[1106]

5.183 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.[1107]

5.184 It said that Entra ID (on the other hand) is entirely interoperable with Active Directory and Microsoft has developed integrations between Entra ID and Active Directory's Kerberos-based authentication mechanism, which are not available to third parties. That same cloud provider said that as a result, most customers with Microsoft-related workloads continue using their on-premises Active Directory with Entra ID.[1108]

5.185 It said the effects of these restrictions, combined with Microsoft's software licensing restrictions, can also be seen in the CMA's market data. It said that this technical challenge, together with licensing conditions make it uneconomical for customers to switch Microsoft-related workloads to other cloud providers.[1109]

1104 Active Directory's Group Policy functionality allows IT administrators to centrally configure settings on Windows computers, for example in a corporate network ('Group Policy overview – Microsoft Learn', accessed 10 September 2024).
1105 Google's responses to the CMA's information requests [✂]; Google's response to the Technical barriers working paper dated 06 June 2024, paragraph 16.
1106 [✂] response to the CMA's information request [✂].
1107 [✂] response to the CMA's information request [✂].
1108 [✂] responses to the CMA's information requests [✂]. As an example, the same cloud provider said that Azure users can authenticate to an Azure-hosted Windows virtual machine by using their Entra ID credentials over Remote Desktop Protocol (RDP), which is possible because of a Kerberos extension that Microsoft has developed. The same cloud provider said that the same extension points are either not available or not documented for third parties, which prevents third parties from implementing a similar feature. It said that as a result, it is not possible for users to use a third-party IAM service to sign in to such a virtual machine over RDP without incurring additional friction or compromising security.
1109 [✂] response to the CMA's information requests [✂]; [✂] submission to the CMA [✂].

5.186 Some customers indicated that they do use both Active Directory and Entra ID in a hybrid cloud architecture.[1110] 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.[1111]

5.187 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, which enables complete interoperability. IT said that multiple IAM providers (such as Okta) support comprehensive integration with Active Directory.[1112]

Technical integration between Entra ID and other Microsoft software and services using Entra ID

5.188 We have also seen evidence suggesting that, for technical reasons, customers who use other Microsoft software and services may choose Entra ID over competing IAM services and that 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.[1113]

(a) Google said that 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. It said that this is consistent with its own experience and is another clear example of Microsoft leveraging its market power in software markets into new cloud markets (from Microsoft 365 into cloud IAM services).[1114]

(b) A cloud provider said that 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.[1115]

(c) A cloud provider said that 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.[1116]

1110 Responses to the CMA's information requests [✂].
1111 [✂] response to the CMA's information requests [✂].
1112 Microsoft's submission to the CMA dated 8 August 2024, CMA Cloud Services MIR - Identity and Access Management Follow-up.
1113 Google's response to the Technical barriers working paper dated 06 June 2024, paragraph 24; Submissions to the CMA [✂]; Responses to the CMA's information requests [✂]; Note of meeting with [✂].
1114 Google's response to the Technical barriers working paper dated 06 June 2024, paragraph 24.
1115 [✂] submission to the CMA [✂].
1116 [✂] response to the CMA's information request [✂].

(d) 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 interoperability between Entra ID and third-party identity management products.[1117]

(e) 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.[1118]

5.189 One customer said that using Entra ID does not necessarily incentivise or force them to use other Azure cloud services.[1119]

5.190 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.[1120]

5.191 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.[1121]

Our assessment

5.192 We consider that IAM is an important area for customers' ability to switch and integrate multiple public clouds.

5.193 We recognise that 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 and potentially to a lesser extent, switch between clouds.

5.194 However, the evidence set out above shows that the design and implementation of IAM services and tools hinders customers' ability to switch and multi-cloud. 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.

[1117] CISPE's submission to the CMA [✂].
[1118] [✂] submission to the CMA [✂]; Note of meeting with [✂].
[1119] [✂] response to Ofcom's information request [✂].
[1120] Microsoft's submission to the CMA dated 8 August 2024, CMA Cloud Services MIR - Identity and Access Management Follow-up, page 2.
[1121] Microsoft's submission to the CMA dated 8 August 2024, CMA Cloud Services MIR - Identity and Access Management Follow-up, page 4.

5.195 We recognise that AWS Cedar (and comparable open-source projects, such as Open Policy Agent's Rego[1122]) may help customers who develop software applications involving IAM-related features. We also note that it has been adopted by at least one ISV, in a service that helps customers develop software applications.[1123]

5.196 However, we consider that AWS Cedar is not sufficient to mitigate the technical barriers associated with the design and implementation of cloud providers' IAM services and tools.

5.197 For example, Okta explained 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 authorization 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.[1124] As noted above, Okta also told us that previous attempts to solve other IAM-related issues have not been successful.

5.198 We also recognise 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.[1125] It is therefore unlikely to significantly help mitigate challenges for customers who switch between IAM services or integrate multiple public clouds (discussed above).

5.199 We have also considered evidence on whether there are additional specific technical challenges with Microsoft's IAM services and tools. We note that we have done so in a context in which we have found customers to be experiencing technical barriers in relation to IAM more generally. As we have not gathered sufficient evidence to assess whether there is currently a material impact from these specific technical challenges, we have provisionally found that Microsoft's IAM services and tools do not create additional technical barriers affecting the ability of customers to switch or multi-cloud through additional mechanisms beyond those experienced in relation to other IAM services and tools.

[1122] 'Policy-based control for cloud native environments – Open Policy Agent', accessed 17 October 2024.
[1123] 'Permit.io Cedar Implementation Q&A: Everything you need to know', accessed 2 October 2024.
[1124] Note of meeting with Okta [✂]; Okta submission to the CMA [✂]. AuthZEN is an industry working group.
[1125] 'Get the best out of Amazon Verified Permissions by using fine-grained authorization methods – AWS Security Blog', accessed 31 October 2024, and 'Amazon Verified Permissions increases default quotas for authorization APIs – AWS', accessed 31 October 2024.

Other ancillary services and tools

5.200 In addition to IAM, there are other categories of ancillary services and tools, including billing, observability and Continuous Integration/Continuous Deployment (CI/CD).[1126]

5.201 We have seen some evidence on these other types of ancillary services and tools although this has not been sufficient for us to form a conclusion on the extent to which they represent barriers to switching and multi-cloud.[1127]

(a) In relation to billing, a customer said that cloud billing is 'painful and difficult to work with'. It said 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.[1128]

(b) In relation to observability, while there is ongoing work on standards like OpenTelemetry,[1129] two customers, Co-operative Group and [✂], suggested that differentiation is making switching and multi-cloud difficult, noting the need for an uplift in skills, re-engineering of the application/platform and all of the security and observability required for the service.[1130]

(c) In relation to CI/CD, an ISV stated that a customer switching cloud provider might have to expend significant effort redeveloping its CI/CD setup.[1131]

Other technical factors

5.202 We have considered whether certain other technical factors are also barriers to switching and multi-cloud, specifically data latency, skills and transparency. These factors arise separately from the purely technical details of core and ancillary service design and implementation, but nevertheless can have implications on the ability of customers to switch and multi-cloud.

Latency

5.203 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

[1126] CI/CD services and tools are part of a software development method, in which changes to the project's code or configuration files are automatically processed and integrated into the 'master copy' of the project code or configuration file.
[1127] Technical barriers working paper, paragraphs 6.62-6.88.
[1128] Note of meeting with [✂].
[1129] In the Technical barriers working paper we said that there are current efforts to standardise some aspects of observability tools around a software package and associated specification named OpenTelemetry. See Technical barriers working paper, paragraphs 6.70-6.79.
[1130] Co-Operative Group's response to the CMA's information request [✂]; Notes of meetings with [✂].
[1131] Note of meeting with [✂].

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.

5.204 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 Providers' views

5.205 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.[1132] 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.[1133]

5.206 Similarly, AWS said that some customers require ultra-low latency and sometimes only one cloud services 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 third-party infrastructure can meet its needs. In such circumstances, there is not much cloud services providers can do, other than offering infrastructure that is even closer, which may not be feasible.[1134]

5.207 In response to our technical barriers working paper, cloud providers highlighted existing mitigations for latency.[1135] For example, Microsoft said that direct interconnects between cloud providers' data centres already exist and do not resolve problems for latency-sensitive applications.[1136]

5.208 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, [✂].[1137]

[1132] 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 [✂].
[1133] AWS' response to the CMA's information request [✂].
[1134] AWS' response to the CMA's Technical barriers working paper dated 06 June 2024, paragraph 45.
[1135] 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).
[1136] Microsoft's response to the CMA's Technical barriers working paper dated 06 June 2024, paragraph 78.
[1137] [✂] submission to the CMA [✂].

Customer views

5.209 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.[1138] This is because transferring data between separate clouds means it may have to traverse intermediate steps such as over the public internet.

5.210 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.[1139]

5.211 However, some customers did not see latency as an impediment to operating a multi-cloud architecture.[1140] 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.[1141]

Other stakeholder views

5.212 Two ISVs said that latency can be an obstacle to multi-cloud because it can impact the performance and functionality of customers' applications.[1142]

Our assessment

5.213 The evidence set out above 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.

5.214 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-

[1138] Responses to the CMA's information requests [✂].
[1139] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 4.1.15.
[1140] Responses to the CMA's information requests [✂].
[1141] Notes of meetings with [✂].
[1142] Notes of meetings with [✂].

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

5.215 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 providers' views

5.216 Cloud providers told us about the training they offer to customers.[1143]

(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.[1144]

(b) Microsoft submitted that, to the extent that there are skill gaps, professional services firms have emerged to provide support for developers.[1145]

(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 industry. It said that engineers certified on other clouds can easily become certified on GCP.[1146]

5.217 However, a small cloud provider said that large 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 large 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.[1147]

Customer views

5.218 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.[1148]

[1143] 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.
[1144] AWS' response to the CMA's information request [✂].
[1145] Microsoft's response to the CMA's Technical barriers working paper dated 06 June 2024, paragraph 86.
[1146] Google's response to the CMA's Technical barriers working paper dated 06 June 2024, paragraph 26(a).
[1147] Cloud Provider 1's response to the Issues Statement dated 17 October 2023, section 3.
[1148] Responses to the CMA's information requests [✂]; [✂] response to Ofcom's information request [✂].

5.219 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.[1149]

5.220 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.[1150] These costs include the recruitment difficulties that can arise from having to hire new cloud engineers to operate on different clouds.[1151]

5.221 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.[1152]

5.222 Some customers said that the training offered by large cloud providers is an important consideration and can be an enabler of multi-cloud.[1153]

5.223 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.[1154]

Other stakeholder views

5.224 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.[1155]

5.225 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.[1156]

1149 Responses to the CMA's information requests [✂]; Note of meeting with [✂].
1150 [✂] response to the CMA's information request [✂]; Notes of meetings with [✂].
1151 Responses to the CMA's information requests [✂]. [✂] submission to the CMA [✂].
1152 Responses to the CMA's information requests [✂].
1153 Responses to the CMA's information requests [✂].
1154 Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraphs 4.3.4-4.3.14 and 4.5.4.
1155 Note of meeting with [✂].
1156 Note of meeting with [✂].

5.226 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.[1157] One (KPMG) said that it is the biggest migration challenge facing customers.[1158]

5.227 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.[1159]

5.228 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 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.[1160]

Our assessment

5.229 We consider 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.

5.230 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.

5.231 We also considered 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.

5.232 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

[1157] Note of meeting with KPMG [✂]; Note of meeting with [✂].
[1158] Note of meeting with KPMG [✂].
[1159] Note of meeting with KPMG [✂].
[1160] Note of meeting with [✂].

of this, it appears that this is not sufficient to overcome the effect of the lack of transferable skills on switching and multi-cloud.

Transparency

5.233 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 providers' views

5.234 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:

(a) AWS said that it takes many active steps to inform and educate its customers, explaining the programming language behind various tools that can be used to build on AWS and documenting the changes to the underlying open source of its managed open-source services.[1161] AWS also said that it regularly publishes blog posts dedicated to the topic of switching away from its public cloud. It also said it provides free courses and guidance on how to move workloads to or from AWS through the AWS Migration Acceleration Program.[1162]

(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.[1163]

(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.[1164] Google offers public training and guidance to switch between cloud services or pursue a multi-cloud strategy with Google.[1165]

Customer views

5.235 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.[1166] This

[1161] AWS' response to the CMA's Technical Barriers Working Paper dated 06 June 2024, paragraph 53.
[1162] AWS' response to the CMA's information request [✂]; AWS' submission to the CMA [✂].
[1163] Microsoft's response to the CMA's information request [✂].
[1164] Google's response to the CMA's Technical barriers working paper dated 06 June 2024, paragraph 26(c).
[1165] Google's response to the CMA's information request [✂].
[1166] Responses to the CMA's information requests [✂].

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.[1167] Another customer also said that some providers are more transparent than others and where transparency is lower, this can make integration more challenging.[1168]

5.236 In contrast, some customers said that the available documentation is usually sufficient.[1169]

Other stakeholder views

5.237 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.[1170]

5.238 This is consistent with market research which said that subtle differences in capabilities or configuration which may exist below the threshold of awareness can be frustrating for customers when portability is attempted.[1171]

5.239 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.[1172]

Our assessment

5.240 We consider 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

5.241 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

[1167] Responses to the CMA's information requests [✂].
[1168] [✂] response to the CMA's information request [✂].
[1169] Responses to the CMA's information requests [✂].
[1170] Accenture's response to the CMA's information request [✂].
[1171] [✂] response to the CMA's information request [✂].
[1172] Accenture's response to the CMA's information request [✂].

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.

5.242 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 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.[1173]

(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.

Customers and other market participants

5.243 We set out below the evidence from customers and other market participants on the technical mitigations that support customers' ability to switch and multi-cloud.

IaC

5.244 We heard from a non-profit organisation that IaC tools like Terraform have emerged to overcome the absence of standardised cloud provider APIs.[1174]

[1173] 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.
[1174] Note of meeting with [✂].

5.245 However, we also heard from an ISV that provides IaC tools that an IaC 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 IaC tools only support the provisioning of services and not their runtime.[1175]

5.246 Some customers and the Jigsaw report said that IaC platforms enable multi-cloud deployments by standardising the deployment of infrastructure across cloud providers.[1176] Additionally, some customers and ISVs and a provider of professional services said IaC tools improve a customer's ability to manage their multi-cloud architecture.[1177]

5.247 However, market research and the Jigsaw report stated that IaC tools can introduce their own lock-in.[1178] 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.[1179]

5.248 Finally, 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.[1180] Additionally, it said that that the IaC vendor lock-in can be overcome,[1181] but competing with IaC offered by cloud providers is harder since those tools are offered for free to customers.[1182]

Open-source software

5.249 Two ISVs, Blue Prism and Splunk, and a customer, Vodafone, said the use of open-source software improves the scope for portability across clouds.[1183] One ISV mentioned limiting the number of proprietary services it used to build its product to increase the scope of portability of the product.[1184] Similarly, two customers said that using generic IaaS increased their ability to switch between different clouds.[1185]

5.250 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

---

1175 Note of meeting with [✂].
1176 Notes of meetings with [✂]; Responses to the CMA's information requests [✂]; Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 4.4.7.
1177 Notes of meetings with [✂]; [✂] response to the CMA's information request [✂].
1178 Responses to the CMA's information requests [✂]; Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 4.4.17.
1179 Notes of meetings with [✂]; [✂] response to the CMA's information request [✂]; Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 4.4.17.
1180 Note of meeting with [✂].
1181 Note of meeting with [✂].
1182 Note of meeting with [✂].
1183 Notes of meetings with Blue Prism, Splunk and Vodafone [✂].
1184 [✂] submission to the CMA [✂].
1185 Responses to the CMA's information requests [✂].

solutions and potential financial loss due to loss of access to critical features.[1186] Additionally, a supplier of professional services said that while working with open-source software, customers should understand that they are trading service portability for continuous management and support and that they must rely on a community for understanding and using the tool.[1187]

5.251 We heard from a non-profit organisation that cloud providers make their managed open-source services interoperable between platforms.[1188] It also said that the large engineering costs associated with forking OSS is a major deterrent to such practices.[1189] 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.[1190]

5.252 Finally, KPMG said customers typically use PaaS for their high-value workloads since it is more economically viable than for low value workloads.[1191] 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.[1192]

Containers and Kubernetes

5.253 Some customers said that containers improve the portability of their workloads[1193] with one customer saying it helped them manage applications across multiple clouds.[1194] 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.[1195]

5.254 Similarly, customers said that Kubernetes[1196] helped streamline application development for portability and provided open APIs that support workload portability.[1197]

5.255 However, some customers, an ISV and a supplier of professional services said that containers cannot effectively overcome the technical barriers to switching and

[1186] Responses to the CMA's information requests [✂]; Notes of meetings with [✂].
[1187] Note of meeting with [✂].
[1188] [✂] submission to the CMA [✂].
[1189] [✂] submission to the CMA [✂].
[1190] [✂] submission to the CMA [✂].
[1191] Note of meeting with KPMG [✂].
[1192] Accenture's response to the CMA's information request [✂].
[1193] Responses to the CMA's information requests [✂]; Note of meeting with [✂].
[1194] [✂] response to the CMA's information request [✂].
[1195] Accenture's response to the CMA's information request [✂]; Container orchestration platforms provide services and tools for automating container orchestration.
[1196] Kubernetes is a portable, extensible, open-source platform for managing containerised workloads and services, that facilitates declarative configuration and automation.
[1197] Responses to the CMA's information requests [✂].

multi-cloud.[1198] Consistent with this, a market research firm said that containers cannot significantly improve portability.[1199]

5.256 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'.[1200]

5.257 Similarly, market research said 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.[1201]

5.258 Separately, a customer, Asda, said that using Kubernetes increases the management costs and requires significant upfront investment,[1202] while an ISV, Blue Prism, said using Kubernetes can increase a customer's cost to switching and multi-cloud.[1203]

Cloud-agnostic services and tools

5.259 We heard from a customer that there is a cloud-agnostic tool for every type of tool it needs,[1204] while an ISV said it uses a cloud-agnostic billing tool, Flexera, to manage its billing processes across the three major cloud providers.[1205]

5.260 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.[1206]

5.261 However, a customer, a market research firm 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.[1207] For example, a customer said that for a hypothetical application

[1198] Notes of meetings with [✂].
[1199] [✂] response to the CMA's information request [✂].
[1200] Note of meeting with KPMG [✂].
[1201] [✂] response to the CMA's information request [✂].
[1202] Asda's response to the CMA's information request [✂].
[1203] Note of meeting with Blue Prism [✂].
[1204] Note of meeting with [✂].
[1205] Note of meeting with [✂].
[1206] Note of meeting with Capgemini [✂].
[1207] Responses to the CMA's information requests [✂].

workload that costs £1 million to develop, it might cost £200,000 to ensure that the workload could be easily switched between public clouds.[1208]

5.262 Additionally, an ISV said that cloud-agnostic tools are not as integrated into the a provider's cloud as the proprietary tools are,[1209] while a market research firm said that cloud agnostic tools will rarely have the same breadth and depth of functionality across all clouds.[1210] Finally, a supplier of professional services and a non-profit organisation said that using cloud-agnostic ISVs can still lead to vendor lock-in.[1211]

Adaptors

5.263 On the use of adaptors, 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.[1212] Separately, 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'.[1213]

Cloud providers' views

IaC

5.264 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 IaC tools needed to perform single cloud deployments. It said that, even with the time needed to learn IaC 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 IaC tools like Terraform can become integrated into customers' development pipelines, adding steps to transition from one IaC tool to another for certain use cases, they can ultimately save customers time and costs and promote and facilitate multi-clouding.[1214]

5.265 We have seen in [✂].[1215]

1208 Note of meeting with [✂].
1209 Note of meeting with [✂].
1210 [✂] response to the CMA's information request [✂].
1211 Notes of meetings with [✂].
1212 Note of meeting with [✂].
1213 [✂] response to the CMA's information request [✂].
1214 AWS' submission to the CMA [✂].
1215 [✂] response to the CMA's information request [✂].

5.266 We also understand that some cloud providers have made investments to support the growth of IaC services and tools [✂].[1216]

5.267 However, Civo said that an IaC tool can abstract some but not all differences between clouds. [1217]

Open-source software

5.268 Cloud providers utilise 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.

5.269 In general, cloud providers said that they support open-source software:

(a) AWS and Microsoft said they support open-source software that facilitates switching and provided the examples of Linux and Kubernetes.[1218] 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.[1219]

(b) Google said its cloud supports a variety of open standards, protocols and compatibility with open-source frameworks.[1220]

(c) IBM and Oracle said they support open-source software.[1221]

5.270 In response to the technical barriers working paper, 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.[1222]

5.271 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

1216 Responses to the CMA’s information requests [✂]; Transcript of hearing with [✂].
1217 Note of meeting with Civo [✂].
1218 AWS' response to the Technical barriers working paper dated 06 June 2024, paragraph 19(a); Microsoft’s response to the CMA’s information request [✂].
1219 Microsoft's response to the Technical barriers working paper dated 06 June 2024, paragraph 35.
1220 Google’s response to the CMA’s information request [✂].
1221 IBM’s response to the CMA’s information request [✂]; Oracle’s response to the CMA’s information request [✂].
1222 AWS' confidential response to the Technical barriers working paper dated 06 June 2024, paragraph 64.

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.[1223] AWS also said it makes many of its APIs and SDKs publicly available under open-source licences,[1224] 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.[1225]

(b) 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.[1226] Microsoft said it helps with multi-cloud and switching by using common protocols and formats and providing APIs and SDKs for developers.[1227]

(c) Google said its APIs are openly accessible to all and its SDK is available to developers;[1228]

(d) IBM and Oracle said they also provide API and SDK access to third parties.[1229]

(e) 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.[1230]

5.272 Microsoft and Google said that they make investments in open-source projects.[1231]

1223 AWS' response to the CMA's information request [✂].
1224 AWS' submissions to the CMA [✂].
1225 AWS' response to the Technical barriers working paper dated 06 June 2024, paragraph 66(a).
1226 Microsoft's response to the CMA's information request [✂].
1227 Microsoft's response to the CMA's information request [✂].
1228 Google's response to the CMA's information request [✂].
1229 IBM's response to the CMA's information request [✂]; Oracle's response to the CMA's information request [✂].
1230 [✂] response to the CMA's information request [✂].
1231 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.

(a) Microsoft highlighted investments in tools like Radius,[1232] Paraglider,[1233] and Dapr[1234] that it said support application portability and multi-cloud network management.[1235]

(b) Google highlighted its development of Kubernetes.[1236]

Containers and Kubernetes

5.273 In response to evidence about the effectiveness of containers as a mitigation in the technical barriers working paper, 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 also said that the observations that container workloads are sticky, are limited to those customers that used managed container services like ECS and EKS, for which the customers were compensated with management and orchestration capabilities.[1237]

5.274 However, one cloud provider said containerised workloads have a greater 'attach power' and managed workloads are more sticky.[1238]

5.275 In addition, 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.[1239]

5.276 We understand that AWS, Microsoft and Google have made investments to improve the quality of their container and Kubernetes offerings.[1240]

Cloud-agnostic services and tools

5.277 A party said that the scope for unliteral effectiveness of tools used to mitigated technical barriers is limited. [✂]. [✂] said that, [✂], these tools need to

[1232] Microsoft says 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.
[1233] Microsoft says 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.
[1234] Microsoft says 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
[1235] Microsoft response to the Technical Barriers working paper, paragraph 32.
[1236] 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.
[1237] AWS' response to the Technical barriers working paper dated 06 June 2024, paragraphs 66a and 66b.
[1238] [✂] response to the CMA's information request [✂].
[1239] Responses to the CMA's information requests [✂].
[1240] AWS' response to the CMA's information request [✂]; Transcript of hearing with Microsoft [✂]; Google's response to the CMA's information request [✂].

interoperate with various functionalities of the relevant cloud service providers, including APIs. [✂].[1241]

Adaptors

5.278 AWS said that customers can procure adaptors from cloud services providers, often for free or as part of the cloud services 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.[1242]

5.279 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 license, making it open to all. [1243]

Services designed to facilitate switching and multi-cloud

5.280 Cloud providers provided examples of the services that they have developed which they say specifically facilitate switching and multi-cloud for customers.

(a) AWS said that it has developed Amazon BedRock[1244] and Amazon Q[1245] which allow customers to use different foundational models and integrate data across providers with an AI-based personal assistant.[1246] AWS also submitted that 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. [1247] AWS also said that it has invested in solutions like AWS Application Migration Services, AWS Database Migration services, Kubernetes and direct connections to on-premises environments

---

[1241] [✂] submission to the CMA [✂].
[1242] AWS' response to the Technical barriers working paper dated 06 June 2024, paragraph 57.
[1243] AWS' response to the Technical barriers working paper dated 06 June 2024, paragraph 59.
[1244] Amazon Bedrock provides customers with access to foundation models ('FMs'), allows them to experiment with high-performing FMs from a range of AI companies, enabling organisations to access and experiment with different FMs, and integrate the model best suited to their needs into their applications.
[1245] Amazon Q is a generative AI-powered assistant, which customers can use to generate guidance tailored to their business through a simple conversational interface and supports over 40 data source connectors, with 75% of these being connectors to third-party services.
[1246] AWS' response to the Updated issues statement and working papers dated 23 May 2024 and 6 June 2024, paragraph 21(b); AWS' submission to the CMA [✂]. [✂].
[1247] AWS response to the CMA's Technical barriers working paper dated 06 June 2024, paragraph 5.

from AWS, point-to-point connections and site-to-site VPNs, among other investments to assist customers in switching.[1248]

(b) Microsoft said some of its services are designed to facilitate interoperability between clouds and provided the example of Azure Arc.[1249] 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.[1250] Microsoft said it has developed Azure Fabric[1251] to support multi-cloud querying and analytics.[1252]

(c) 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.[1253] Additionally, Google has also worked with third parties like Apache Casandra, Datadog and VMWare to facilitate integrations with its own services.[1254]

(d) 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.[1255]

(e) IBM said that it designs product to be interoperable across clouds and that its cloud offering is based on a hybrid multi-cloud approach, meaning that IBM provides cloud-related services largely irrespective of the customers' choice of cloud service provider.[1256]

Our assessment

5.281 We asked customers for views on 'enablers' of switching and multi-cloud. As set out above, in general, the mitigations that customers mentioned most frequently

---

1248 AWS' response to the Updated issues statement and working papers dated 23 May 2024 and 6 June 2024, paragraph 9.
1249 'Azure Arc overview - Azure Arc - Microsoft Learn', accessed 23 September 2024; Microsoft's response to the CMA's information request [✂].
1250 Microsoft's response to the CMA's information request [✂].
1251 'What is Microsoft Fabric - Microsoft Fabric - Microsoft Learn', accessed 23 September 2024.
1252 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 [✂].
1253 Google's response to the Technical barriers working paper dated 06 June 2024, paragraph 6.
1254 Google's response to the CMA's information request [✂].
1255 Oracle's response to the CMA's information request [✂].
1256 IBM's response to the CMA's information request [✂].

were their use of Infrastructure as Code, self-managed open-source software (using IaaS), containers and Kubernetes, cloud-agnostic services and tools and adaptors. These types of mitigations are typically developed by ISVs or the open-source community to support switching and multi-cloud. We view these mitigations as customer-led, in that customers shoulder most of the implementation effort and cost, which can affect customers' decisions to use them.

5.282 The evidence relating to these 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 effective in doing so, for example because the costs outweighed the benefits, customers were locked-in to the ISVs that supply these mitigations, the breadth of functionality was reduced and scaling was more difficult.

5.283 The cloud providers submitted evidence relating to these customer-led mitigations, but also mentioned a range of other potential mitigations that were generally not mentioned by customers. We view these other mitigations as 'provider-led', meaning services and tools designed by cloud providers to facilitate switching and multi-cloud. Such services and tools may emerge where cloud providers' incentives align with customers' needs.

5.284 We have examined the range of provider-led mitigations further and found that these mitigations appear to vary in quality. We found that some provider-led mitigations sufficiently addressed a particular and narrow problem, some seemed insufficient and for some we couldn't conclude on their effectiveness. For example, we found that:

(a) AWS Glue and Microsoft Fabric allow customers to query and analyse data across data sources located in different cloud environments but these solutions only address multi-cloud data analytics. Similarly, Azure Arc and AWS Outposts help customers manage and govern multi-cloud deployments but only for servers that are 'Arc-enabled' or compatible with AWS Outposts. Overall, these services do not address broader customer needs when it comes to multi-cloud.

(b) AWS has invested in tools like Amazon Bedrock and Amazon Q that support the use of open-source AI models. However, we note that both are managed services deployed primarily on AWS and do not support switching and multi-cloud for customers.

(c) 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 we cannot conclude on their practical effectiveness and usability.

5.285 Overall, we consider that cloud providers, ISVs and open-source communities have developed some mitigations to the technical barriers to switching and multi-cloud. However, the current patchwork of mitigations, which often address particular instances of particular barriers, mitigates those technical barriers to a limited extent but is not sufficient to mitigate the effects of those technical barriers on switching and multi-cloud.

Provisional conclusions

5.286 Our provisional view is that there are substantial technical barriers facing customers who wish to switch or use multiple clouds.

5.287 Some customers can and do successfully multi-cloud but we have provisionally found that 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.

5.288 Whilst customers are able to switch if they are willing to expend all the required resources to do so, we have provisionally found that 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.

5.289 We have provisionally found that 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.

5.290 There are some mitigations to some of these technical barriers, including those from cloud providers and third parties and customers own mitigations. However, these only mitigate the technical barriers to switching and multi-cloud experienced by customers or their effect on competition to a limited extent.

Egress fees

Introduction

5.291 We have considered whether, and to what extent, egress fees are a barrier to switching and multi-cloud. This section presents our analysis of the impact of egress fees on switching and multi-cloud and on competition among cloud providers.

5.292 We set out in turn:

(a) background information on egress fees;

(b) our conceptual framework for assessing the impact of egress fees on competition; and

(c) our analysis and provisional views of:

(i) the relevance of egress fees in customers' switching and multi-cloud decisions; and

(ii) potential benefits in relation to egress fees.

5.293 We have focussed 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, unless stated otherwise.

Background

5.294 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').

5.295 As set out in Figure 5.2, 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 5.2: Types of data transfers

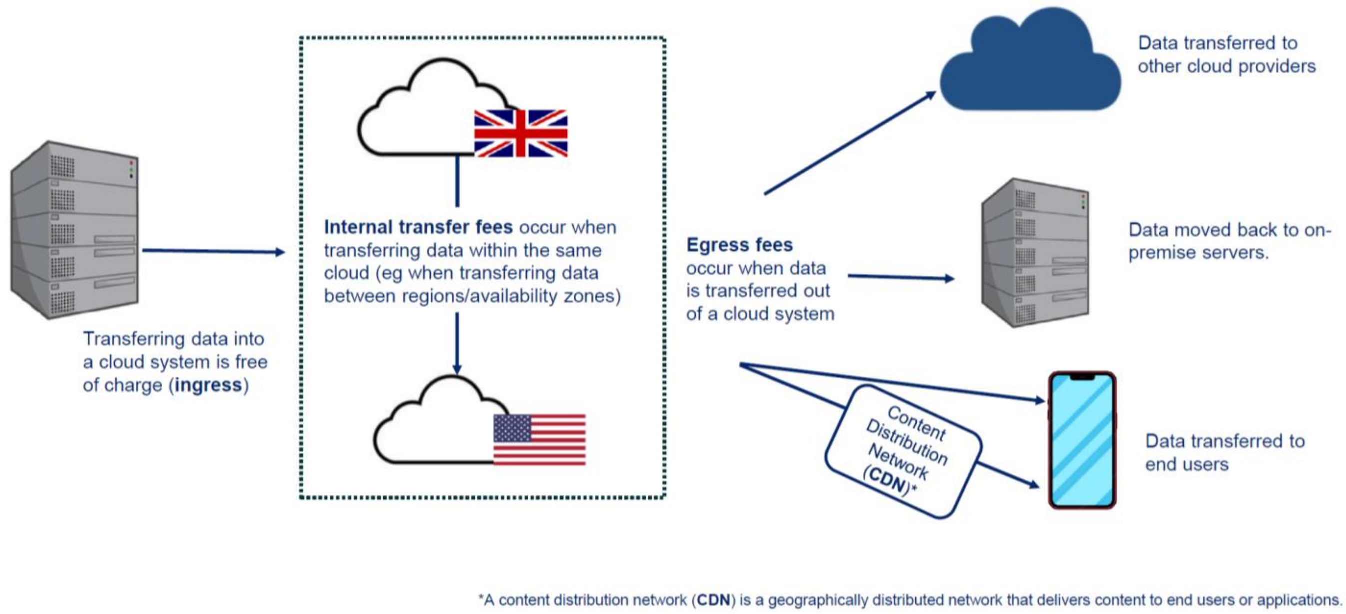

Source: CMA analysis of a Google submission to Ofcom [✂].[1257]

5.296 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.[1258] This can occur when:

(a) customers move their data from their cloud to their on-premises data centres;

(b) customers move data between different cloud providers as part of a multi-cloud architecture or as part of switching; and

(c) customers deliver content to end users or applications, either directly or via content distribution networks ('CDN') (eg Cloudflare).

5.297 We refer to the wider set of fees charged by cloud providers for data transfer services (ie internal transfer fees and egress fees) as data transfer fees.

5.298 The cost of egressing data 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

[1257] 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 provisional report.
[1258] The terminology and the definition of these fees are not consistently used by cloud providers.

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.

5.299 Our analysis covers egress fees incurred when transferring data to other cloud providers, as customers are likely to consider these fees when switching cloud provider or using multiple clouds. However, we have been told by more than one cloud provider that they cannot:[1259]

(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

5.300 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.

5.301 Egress fees represent a cost to customers both when switching and when using multiple clouds:

(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 (except if 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 such a case, both costs would need to be considered.

5.302 Where the 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

[1259] Responses to the CMA's information requests [✂].

when alternative cloud providers would otherwise have a better offer (in terms of price, quality, range of features/capabilities, or other considerations).

5.303 We have assessed the following issues:

(a) First, we consider 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 also consider the effect of free switching programmes on our analysis of switching costs.

(b) Second, we consider whether there are potential benefits, such as efficiencies, which may remove or mitigate any competitive harm arising from egress fees. In particular we have considered cloud providers' submissions that costs recovered from egress fees 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 on customers' switching and multi-cloud decisions

5.304 As set out at the start of this chapter, barriers to switching suppliers or barriers to using multiple suppliers 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.[1260]

5.305 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 may be expected to make customers less responsive to the efforts of rivals to compete by improving their offer.

5.306 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.[1261] When transferring as a 'one-off', the impact

[1260] CC3 (Revised), paragraphs 297 and 316.
[1261] 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.

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 may be larger.

5.307 In view of the above, we have considered the extent to which egress fees may be relevant to customers' switching and/or multi-cloud decisions. In particular, we have considered evidence on:

(a) cloud providers' views on the relevance of egress fees;

(b) the cost of egressing data under hypothetical scenarios of one-off switching and multi-cloud;

(c) the impact of recently announced free switching programmes on our analysis;

(d) customers' views on the impact of egress fees on their decisions to switch and/or multi-cloud; and

(e) cloud providers' internal documents relevant to egress fees.

5.308 We focus much of our analysis on AWS, Microsoft 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,[1262] we consider that the findings of our analysis would also apply more widely to any cloud providers that charge egress fees, as any level of egress fees may 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.

Cloud providers' views

5.309 Cloud providers have submitted views on the relevance of egress fees. We have grouped these by theme below:

[1262] 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 24 September 2024.

(a) The observed prevalence of egress fees;

(b) customers' levels of concern about egress fees;

(c) trends in egress fees; and

(d) justification for egress fees in the context of free ingress.

The observed prevalence of egress fees

5.310 AWS, Microsoft and Google told us the proportion of their customers who paid egress fees for any purpose in 2023. These varied widely.

5.311 AWS submitted that [✂] of UK customers with at least $1,000 of annual spend paid egress fees[1263] 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. [1264]

5.312 A separate analysis by AWS which included customers that spent under $1,000, estimated that approximately [✂] of UK customers paid egress fees in 2023. AWS noted that this may still overestimate the share of UK customers paying DTO fees because it excludes certain customers, namely customers [✂] and customers where AWS has not been able to match [✂].[1265]

(a) 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 whose low spend may be due to experimenting with AWS services.

(b) This submission also includes customers with a total data egress volume of under 1,200GB in 2023 as falling within AWS' free tier. We note that this does not include customers who have exceeded the 100GB free tier in any given month (and therefore paid egress fees) as having paid egress fees in 2023. We would expect customers who have paid egress fees at any point in a given year to be included.

5.313 Microsoft submitted that [✂] 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, unweighted by yearly cloud spend. When weighted by cloud spend, this rose to [✂].[1266]

[1263] 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.
[1264] AWS' response to the CMA's information request [✂].
[1265] AWS' response to the CMA's information request [✂].
[1266] Microsoft's response to the CMA's information request [✂].

5.314 Google submitted that just over [✂] of its UK customers with a cloud spend of over $1,000 in 2023 paid egress fees, unweighted by yearly cloud spend. When weighted by cloud spend, this rose to [✂]. 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 course 'serving' egress), it is to be expected that most customers pay egress fees.[1267]

5.315 We consider that these analyses are not wholly representative of customers' eligibility to pay egress fees. Some customers are likely to see egress fees as being too high and therefore avoid needing to pay them, reducing the overall incidence of egress fees despite the fact they have still had an impact on customers' decision-making. Further, these analyses do not tell us about the observed prevalence of egress fees specifically for switching and/or multi-cloud. As such, we do not place any weight on these analyses.

Customers' levels of concern about egress fees

5.316 Google submitted that the majority of customer evidence in our Egress fees working paper and in the Jigsaw report shows that egress fees are not a meaningful barrier to switching.[1268]

5.317 AWS submitted that the evidence gathered on customer views indicates 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]'.[1269] AWS further submitted that survey evidence does not support the claim that egress fees distort competition.[1270]

5.318 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.[1271]

[1267] Google's response to the CMA's information request [✂].
[1268] Google's response to Egress Fees Working Paper, paragraph 31.
[1269] AWS' response to the CMA's updated issues statement and working papers, paragraph 35(d).
[1270] AWS' submission to the CMA [✂].
[1271] IBM's response to the CMA's working paper on egress fees, page 1.

5.319 Microsoft submitted that 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.[1272]

5.320 When assessing customers' views, we consider that egress fees:

(a) do not have to be the main concern, but can be one of a number of concerns and still be meaningful, and

(b) do not have to be a relevant factor for a majority of customers to be considered a meaningful barrier to switching and/or multi-cloud. Their influence for a substantial minority of customers is sufficient.

Trends in egress fees

5.321 Google submitted that the effective price[1273] 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.[1274]

5.322 AWS submitted that its internet egress fees had fallen by over 30% globally between 2018 and 2022.[1275]

5.323 Microsoft submitted that competition has driven egress fees down and that quality-adjusted egress fees have fallen for customers in the UK, while the volume of data transferred had [✂]. Microsoft also submitted that nominal effective prices for Microsoft's egress services had remained stable despite increases in product quality and infrastructure since 2018.[1276]

5.324 We do not consider that effective prices for egress falling is enough to conclude that the presence of egress fees cannot constitute a barrier to switching and/or multi-cloud. Additionally, 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.

Justification for egress fees in the context of free ingress

5.325 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

[1272] Microsoft's response to the Competitive landscape, Committed spend agreements and Egress fees working papers, paragraph 89.
[1273] Here, Google uses 'effective price' to mean price adjusted for quality. Even when nominal prices remain the same, effective prices can reduce if the product in question increases in quality.
[1274] Google's response to Egress fees working paper, paragraph 15.
[1275] AWS' response to the CMA's updated issues statement and working papers, paragraph 35(a).
[1276] Microsoft's submission to the CMA [✂].

assets.[1277] This raises the question as to whether egress fees have been set too high, given that data ingress is free despite the fact that both services use the same fixed assets and network capacity, and therefore represent common costs incurred by a cloud provider.[1278]

5.326 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.[1279]

5.327 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.[1280]

5.328 Microsoft submitted that, historically, cloud providers charged for both egress and ingress. As competition in the cloud market intensified, Microsoft (and other cloud providers) moved to a model of free ingress but more significant charges for egress. Microsoft further 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 the cloud, while mitigating overall costs for integrating multi-cloud solutions and for customer switching.[1281]

5.329 Oracle submitted that if egress fees for all cloud providers were based on a cost-recovery model, the costs of egress fees 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.[1282]

5.330 We consider that these submissions do not 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. Regardless of the validity of the explanations for egress being charged while ingress is not, we consider that these points do not impact our overall conclusion.

[1277] Responses to the CMA's information requests [✂].
[1278] See Appendix Q for information on common vs egress-specific data transfer costs.
[1279] AWS' response to the CMA's information request [✂].
[1280] Google's response to Egress fees working paper, paragraph 5, footnote 3.
[1281] Microsoft's response to the Competitive landscape, Committed spend agreements and Egress fees working papers, paragraph 85.
[1282] Oracle's response to the Updated Issues Statement and working papers, page 2.

5.331 Nevertheless, we note that all of AWS, Microsoft and Google previously charged for both ingress and egress fees. [✂] and [✂] describe ingress fees as having been waived in response to competitive pressure.[1283] A [✂] internal document shows [✂].[1284] Further, Microsoft's submission that data ingress being free facilitates migration to the cloud suggests that data egress being paid hinders migration out of the cloud.

5.332 Overall, we consider that while data ingress being free could indicate that egress fees are not a necessity, this does not contribute to our assessment of whether egress fees are likely to harm competition.

Analysis of hypothetical scenarios

5.333 When assessing the impact of egress fees on switching and/or multi-cloud, we cannot rely on data on actual spend on egress fees as the level of spending is likely to be 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, we consider that by using the evidence on the actual levels of spend on egress fees, we would risk underestimating the potential for egress fees to have an impact on customers' switching and multi-cloud decisions.

5.334 Our solution to this problem is to consider 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 hypothetical scenarios of switching and of multi-cloud architectures, which we set out below.

Switching

5.335 To assess the extent to which potential egress fees would represent a non-negligible proportion of customers' annual cloud spend, we have considered the following evidence:

(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.

5.336 This analysis is based on list prices for data egress via the internet (ie 'standard tier' egress). While we recognise that each cloud provider does offer various cost-mitigating alternatives (eg direct connect services), our understanding is that standard tier egress is the main method through which switching is executed. For

[1283] [✂] note this was also [✂]. [✂] response to the CMA's information request [✂]; [✂] submission to Ofcom [✂].
[1284] [✂] response to the CMA's information request [✂].

example, we have seen evidence from one provider to suggest that switching out of its cloud is most commonly undertaken using standard tier egress.[1285]

5.337 AWS submitted that [✂].[1286] We discuss this further in Appendix M.

Cloud providers' views on hypothetical 'one-off' switching costs

5.338 AWS, Microsoft and Google submitted that egress fees incurred by switching customers were a small percentage of these customers' 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 does not explicitly explain how.[1287] AWS submitted that this is a level that is unlikely to deter customers from switching;[1288]

(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;[1289]

(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.[1290] 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);[1291] 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.

1285 [✂] response to the CMA's information request [✂].
1286 AWS' submission to the CMA [✂].
1287 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 10 September 2024).
1288 AWS' response to Ofcom's Interim Report, paragraph 17.
1289 AWS' submission to Ofcom [✂].
1290 Microsoft's response to Ofcom's Interim Report, paragraph 237.
1291 Microsoft response to the Issues Statement, paragraph 41.

Google further submitted that egress fees covering all types of egress traffic are still a small proportion of an average customer's spend.[1292]

5.339 Wasabi 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.[1293]

5.340 We 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 who considered egress fees low enough that they do not disincentivise switching. Such an approach will not capture cases where egress fees are higher than the expected benefits of switching and therefore act as a barrier to switching. As such, we place little weight on these submissions.

5.341 We further note that the 2017 IDC analysis referenced by both AWS and Microsoft used an individual 'mid-sized' customer with a specific workload as the base for its analysis.[1294] 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 market has evolved in the seven years since its publication.[1295] We therefore believe that this analysis is unlikely to be indicative of the egress fees the average customer may expect to be charged when switching.

Our analysis of hypothetical switching costs

5.342 We calculated the hypothetical '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 AWS, Microsoft 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;[1296]

(b) we used the list prices on cloud providers' websites for data transfer from the applicable UK region,[1297] via the internet;[1298]

[1292] Google's submission to Ofcom [✂].
[1293] Wasabi's response to the CMA's Working Papers, pages 1 to 2.
[1294] Switching of Cloud Services Providers - Publications Office of the EU), page 42.
[1295] See chapter on competitive landscape.
[1296] CMA analysis of AWS, Microsoft and Google data. AWS' response to the CMA's information request [✂]; Microsoft's response to the CMA's information request [✂]; Google's response to the CMA's information request [✂].
[1297] We have used the 'London' region for Google and AWS as this encompasses all of the UK.
[1298] 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 29 Jan 2024.

(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 for each customer in each applicable year.

5.343 Our analysis is laid out in detail in Appendix M. This shows that:

(a) egress fees paid to move all of a UK customer's data once from one cloud provider to another range from [0-5]% to [5-10]% of a customer's annual spend, with the smallest customers facing the highest egress costs;

(b) the majority of AWS', Microsoft's and Google's UK customers would have had to pay [✂] of their total annual spend in 2022 for a total switch. About [10-20]% of customers would have had to pay egress fees of more than [10-20]% of their total annual spend in a single year;

5.344 Oracle submitted that, in focusing on egress fees as a percentage of customer spend, we underestimate the importance of the absolute value of egress fees to the customer by failing to account for 'sticker shock' at the highest levels. Oracle explained that '[0-5%]' of spend sounds much smaller than $10 million (for example) just to switch cloud providers.'[1299] Nevertheless, we maintain our decision to focus on egress fees as a percentage of customer spend as these are more representative of the impact of egress fees on customers. For instance, the impact of $10 million to a small customer is not the same as the impact of $10 million to a large customer.

5.345 We acknowledge that costs can be spread over multiple years (amortised). A customer may incur switching costs once but incur the benefits over multiple years. As such, we consider analysis from single-year estimates alongside those where costs are spread over multiple years. For example, over a three-year contract,[1300] a customer might expect to pay egress fees of [0-5%] to [0-5%] of annual spend. For those customers who would pay egress fees of over 10% in a single year, egress fees for switching may make up more than [0-5%] of their total yearly cloud spend.

5.346 For [0-5]% - [5-10]% of customers, costs may exceed [10-20]% of annual cloud spend in a single year. If these customers were to spread their costs over an average contract, ie three years, costs may still exceed [0-5]% per year. This means that some customers would not have the incentive to switch to a rival cloud provider even if that provider offered a price that was [10-20]% better or

[1299] Oracle's response to the Updated Issues Statement and working papers, page 3.
[1300] This is the average length of a contract. CMA analysis of AWS and Microsoft data. AWS' response to the CMA's information request [✂]; Microsoft's response to the CMA's information request [✂].

equivalently better quality or functionality. These costs can also be materially higher for some customers.

5.347 More detail on our analysis of hypothetical scenarios is included in Appendix M.

Multi-cloud

5.348 Below, we analyse 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 model these costs under a variety of assumptions and break down our results by customers' spend.

5.349 Below, we consider:

(a) cloud providers' views on hypothetical multi-cloud costs, including a Google submission on the costs of multi-cloud; and

(b) our analysis of hypothetical multi-cloud costs.

Cloud providers' views on hypothetical multi-cloud costs

5.350 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.[1301]

5.351 AWS submitted that, when considering hypothetical scenarios related to multi-cloud, the CMA does not consider whether integrated multi-cloud is the most efficient outcome for customers.[1302] 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.[1303]

5.352 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.[1304]

5.353 Microsoft submitted that it has designed egress fee strategies to minimise the effects of egress fees on multi-cloud customers,[1305] but did not provide information on what these strategies entail.

[1301] Google's submission to the CMA [✂].
[1302] AWS' submission to the CMA [✂].
[1303] AWS' submission to the CMA [✂].
[1304] AWS' submission to the CMA [✂].
[1305] Microsoft's response to the Competitive landscape, Committed spend agreements and Egress fees working papers, paragraph 87.

5.354 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.[1306]

5.355 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.[1307] 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.[1308] We discuss this submission in further detail in Appendix M.

5.356 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.

5.357 We further consider that internal transfer fees are not a relevant benchmark against which to assess the potential impact of egress fees, as not all customers will incur them, eg if they do not operate across multiple regions. As such, we consider this analysis does not show that egress fees are not a barrier to multi-cloud.

Our analysis of hypothetical multi-cloud costs

5.358 In this analysis, we do not consider which cloud architecture provides the greatest benefits to customers, but instead focus on customers' ability to multi-cloud should they wish to. It is the ability to multi-cloud that allows customers to achieve benefits through multi-cloud. In that context, egress fees add some costs for the customer to adopt their preferred architecture.

1306 IBM's response to the CMA's information request [✂].
1307 Google's submission to the CMA [✂].
1308 Google's submission to the CMA [✂].

5.359 With respect to AWS, Microsoft 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;[1309]

(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);[1310] and

(c) we then performed a series of calculations to assess the hypothetical egress fees incurred 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.

5.360 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 hypothetically transfer.

5.361 We have chosen to focus on the percentage of total cloud spend instead of the absolute value, as percentages are more representative of the impact of egress fees on customers.

5.362 AWS, Microsoft and Google all submitted that our data transfer volumes are not reflective of customers' current architectures:

(a) AWS and Google submitted that our assumption that customers may need to transfer up to 100% of their data stored to multi-cloud is not justified by any consideration of how customers actually multi-cloud.[1311]

(b) Microsoft submitted that the hypothetical multi-cloud use case set out in the Egress fees working paper, which Microsoft characterised as one that requires significant data to be moving between clouds on a constant basis, does not exist in practice. Microsoft added that the examples of potential frictions created by multi-cloud provided are not realistic and mainstream.[1312]

[1309] CMA analysis of AWS, Microsoft and Google data. AWS' response to the CMA's information request [✂]; Microsoft's response to the CMA's information request [✂]; Google's response to the CMA's information request [✂].
[1310] 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 29 Jan 2024.
[1311] AWS' submission to the CMA [✂]; Google's submission to the CMA [✂].
[1312] Microsoft's response to the Competitive landscape, Committed spend agreements and Egress fees working papers, page 25.

(c) A cloud provider submitted that, in its own experience, customers who multi-cloud typically transfer less than 10% of their total stored data on a monthly basis.[1313] It further submitted that customers who multi-cloud have strong incentives – unrelated to egress fees, such as compliance, security and data localisation requirements – to keep the volume of data transferred for the purposes of multi-cloud to the minimum amount required.[1314]

(d) The same cloud provider submitted that our analysis fails to account for the different types of egress (and their related costs) customers can adopt.[1315]

5.363 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.

5.364 Further, our analysis only assesses egress fees for data transfers between clouds, as this is the only type of egress relevant to whether egress fees are a barrier to multi-cloud. Hence, we do not consider any other type of egress fees in our analysis (eg egress to end users, or 'serving' egress).

5.365 Overall, our analysis shows that any non-trivial amount of data transferred out of a customer's cloud is likely to lead to non-negligible egress fee costs as a percentage of total monthly cloud spend. As the degree of integration increases, the cost of egress as a percentage of monthly cloud spend increases.[1316] For example:

(a) for customers of any size, egressing 20% of their data can cost from [0-5%] to [10-20%] of their monthly cloud spend in egress fees alone. This increases to [5-10%] – [30-40%] when egressing 50% of data stored;

(b) for customers spending less than £10,000 in 2022, egressing 20% of their data costs at least [0-5%] of their monthly cloud spend; and

(c) for customers spending more than £20 million in 2022, egressing 20% of data costs at least [0-5%] of their monthly cloud spend.

5.366 Our analysis also shows that smaller customers may have a higher disincentive to increase the level of integration between clouds or multi-cloud as they must pay

[1313] [✂] submission to the CMA [✂].
[1314] [✂] submission to the CMA [✂].
[1315] [✂] submission to the CMA [✂].
[1316] CMA analysis of AWS, Microsoft and Google data. AWS' response to the CMA's information request [✂]; Microsoft's response to the CMA's information request [✂]; Google's response to the CMA's information request [✂].

more than larger customers, in terms of percentage of their annual spend, to transfer data between clouds at higher data transfer levels.[1317]

5.367 We consider that, while multi-cloud costs can be mitigated to some extent, some customers would still face high egress fees even after a mitigation.

Free switching programmes

5.368 The European Union's (EU) Data Act (the 'EU Data Act')[1318] 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'.[1319] The EU Data Act contains provisions relevant to egress fees for EU customers:[1320]

(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.[1321] From 12 January 2027, a full removal of switching charges is required.[1322]

(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.[1323]

5.369 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 Google, AWS and Microsoft.

[1317] CMA analysis of AWS, Microsoft and Google data. AWS' response to the CMA's information request [✂]; Microsoft's response to the CMA's information request [✂]; Google's response to the CMA's information request [✂].

[1318] 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').

[1319] 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.'

[1320] EU Data Act, Articles 29(1), 29(2) and 29(3).

[1321] EU Data Act, Articles 29(2) and 29(3).

[1322] EU Data Act, Article 29(1).

[1323] EU Data Act, Article 50.

Following the announcement of these free switching programmes, egress fees have been removed by Civo (a UK-based smaller cloud provider).[1324]

5.370 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.[1325]

5.371 Oracle said that [✂].[1326] We understand this to mean that Oracle [✂] already compliant with the current EU Data Act provisions that apply until January 2027 in relation to egress fees for switching.

5.372 We have considered whether the introduction of these free switching programmes means that we may be overestimating the hypothetical 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.

5.373 To assess this, we have gathered evidence on the nature and applicability of AWS', Google's and Microsoft'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.

Stakeholders' views

5.374 Google submitted that the launch of free switching programmes by Google, AWS and Microsoft has effectively taken the issue of egress fees as a potential barrier to switching off the table.[1327] 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 the 'only reasonable conclusion that can be drawn is that egress fees do not pose any – real or perceived – barrier to customer switching'.[1328]

5.375 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

[1324] Civo abolishes all egress fees to advance cloud simplicity strategy, accessed 24 September 2024.
[1325] IBM's response to the CMA's information request [✂].
[1326] Oracle's response to the CMA's information request [✂].
[1327] Google's response to Egress fees working paper, Annex response (b).
[1328] Google's response to Egress fees working paper, paragraph 42.

two potential concerns around egress fees identified in our Egress fees working paper.[1329]

5.376 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 [✂].[1330]

5.377 Google and AWS have submitted that [✂].[1331]

5.378 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.[1332]

5.379 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, and customers large enough to exceed the 100GB of data per month threshold are likely to require a multi-cloud strategy.[1333]

5.380 In its announcement about removing egress fees, Civo said that its abolishment 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.[1334]

5.381 Vodafone told us that the cloud providers' changed policies for switching do not address data egress charging in non-switching scenarios such as where an enterprise has demand to use a variety of cloud providers. Vodafone said it supports a wider application of data egress policy.[1335]

5.382 Google also 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.[1336]

1329 AWS' response to the Egress fees working paper, paragraph 32.
1330 Note of meeting with AWS [✂].
1331 AWS' response to the CMA's information request [✂]; Google's response to the CMA's information request [✂].
1332 Microsoft's response to the Competitive landscape, Committed spend agreements and Egress fees working papers, paragraphs 90-91.
1333 Oracle, consolidated commentary to the CMA's Updated Issues Statement, its working papers, and the Market investigation qualitative customer research final report, page 4.
1334 Civo abolishes all egress fees to advance cloud simplicity strategy, accessed 24 September 2024.
1335 Vodafone's response to the CMA's Working Papers, page 1.
1336 Google's response to Egress fees working paper, Annex response (b).

5.383 Similarly, 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.[1337]

Our analysis of relevance of free switching programmes to UK customers

5.384 We have assessed the scope and conditions of eligibility for AWS', Microsoft's and Google's newly introduced free switching programmes. In particular, we have 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.

5.385 Our analysis of the free switching programmes is set out in more detail in Appendix N. In this we also consider the uptake of the programmes, that AWS, Google and Microsoft consider to be low and reflective of the lack of egress fee barriers to switching. We find the uptake data to be inconclusive, given that there could also be other explanations for low uptake, such as lack of customer awareness of the programmes, or programme restrictions or uncertainty deterring uptake.

5.386 In summary, we consider that the programmes all have some form of restriction on eligibility which may significantly impact customers' ability to switch in practice; are all dependent on the cloud providers' own discretion to a degree; and are all reliant on existing terms and conditions (including application to the UK) not changing in the future. There are also other limits to the scope of the programmes in relation to service or data type eligibility restrictions and requirements for customers to be aware of the programmes and submit switching applications in advance.

5.387 In particular, we consider the following to be relevant to our assessment:

[1337] Transcript of hearing with AWS [✂].

(a) The free switching programmes are for customers that are switching cloud providers, not multi-cloud use. We are assessing the relevance of egress fees on customers' switching and multi-cloud decisions.

(b) The free switching programmes have been voluntarily and only recently introduced in the UK by AWS, Microsoft and Google, in response to the EU Data Act. They currently apply in the UK but that 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. Although AWS and Google have 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,[1338] we have the following concerns:

(i) First, as the programmes have been voluntarily introduced only recently, we consider it would be relatively easy for cloud providers to change or remove their applicability in the UK.

(ii) Second, whilst risk of reputational damage may provide an incentive for cloud providers to retain the programmes, we do not consider the programmes to have been in place long enough for them to be embedded in customers' awareness and hence they do not constitute a sufficiently strong prevention against the narrowing or removal of the programmes.

(c) AWS', Microsoft's and Google's free switching programmes all have some form of restriction on eligibility (for example, Microsoft's programme does not cover partial switching; AWS' and Google's programme terms preclude multi-cloud during switching; all programmes are limited to a 60-day time period and do not apply to ongoing multi-cloud use) and the cloud providers have a substantial degree of discretion over the application of the programmes. In particular, we note that:

(i) The 60-day time limit may mean the free switching programmes are ineffective in practice. Customer evidence suggests switching usually takes longer than 60 days,[1339] and any extension of the cloud providers' free switching programmes beyond 60 days is entirely at the respective cloud providers' discretion.

(ii) Some customer evidence also suggested that the process of switching involves periods of multi-cloud running,[1340] which would not appear to

[1338] Google's response to Egress fees working paper, Annex response (b); Transcript of hearing with [✂].
[1339] Responses to the CMA's information requests [✂]; Notes of meetings with [✂]; Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraphs 3.6.5, 3.6.7, 3.6.8; and page 43.
[1340] [✂] response to Ofcom's information request [✂]; Note of meeting with [✂]; Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 3.6.8.

be eligible under the published programme terms on AWS' and Google's websites.

5.388 For the above reasons, we consider that the free switching programmes introduced by AWS, Microsoft and Google have limited impact on our competition assessment. Where relevant, we have considered them in our proposed remedies assessment.

Customer views on the impact of egress fees

5.389 In order to assess whether egress fees are having an impact on customer behaviour, or are likely to do so, we asked 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.

5.390 In line with the qualitative nature of the evidence we gathered, we have given a narrative summary of the key points that we consider emerge from the evidence. Further detail is presented in the Appendix O.

5.391 We also commissioned qualitative customer research from Jigsaw Research. The evidence is set out in the Jigsaw report. This research was intended to capture a wider range and a different set of customers from those we contacted through direct channels. A narrative summary of the key points is also presented below.

Switching

Customer evidence

5.392 Of all the customers we asked questions to that have switched or have considered switching, only a few spontaneously identified egress fees as a challenge.[1341]

5.393 We note that many customers who had never considered switching had not done so because they were satisfied with or had only recently moved to their current cloud provider.[1342] A few customers that had never considered switching explicitly identified egress fees as a factor that could disincentivise switching.[1343]

5.394 A few customers that had considered switching 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.[1344]

[1341] Responses to the CMA's information requests [✂]; Note of meeting with [✂].
[1342] Responses to the CMA's information requests [✂].
[1343] Responses to the CMA's information requests [✂].
[1344] Notes of meetings with [✂].

5.395 Some 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.[1345]

Evidence from the Jigsaw report

5.396 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.[1346]

5.397 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. 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.[1347]

5.398 For some customers who have switched, 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.[1348]

5.399 We note that, when assessing customers' views, we do not consider it necessary for egress fees to be a significant concern for a majority of customers to be considered a meaningful barrier to multi-cloud or switching. The significance of egress fees for a substantial minority of customers is sufficient.

Multi-cloud

Customer evidence

5.400 Several customers we contacted identified themselves as using an integrated form of multi-cloud architecture, with different levels of integration.[1349]

5.401 A few customers with integrated multi-cloud architectures said that they did not consider egress fees to be a material challenge in setting up a multi-cloud architecture.[1350]

[1345] [✂] response to the CMA's information request [✂]; Notes of meetings with [✂].
[1346] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraphs 5.2.2 and 5.2.3.
[1347] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 5.2.5.
[1348] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 5.3.2.
[1349] Responses to the CMA's information requests [✂].
[1350] Responses to the CMA's information requests [✂].

5.402 Many of the customers with integrated multi-cloud architectures said that egress fees have been a challenge to multi-cloud architectures and/or they took egress fees into consideration when making their decision to set up a multi-cloud architecture.[1351] Reasons for this 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.

5.403 Of the handful of customers who use a siloed multi-cloud architecture,[1352] 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. At the time of submission one customer said that egress fees were not a material disincentive to multi-cloud for its use cases, but it now has a use case involving synchronisation of large amounts of data for which egress fees do provide a disincentive to multi-cloud.[1353] Since then it has developed a use case which 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.[1354]

5.404 For 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.[1355]

(b) A handful of 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.[1356]

(c) A few of the customers we contacted said they used a single cloud provider.[1357] We asked these customers why they used a single public cloud and if there were any potential challenges that they may 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.

[1351] Responses to the CMA's information requests [✂].
[1352] Responses to the CMA's information requests [✂].
[1353] [✂] response to the CMA's information request [✂].
[1354] [✂] submission to the CMA [✂].
[1355] [✂] response to the CMA's information request [✂].
[1356] Responses to the CMA's information requests [✂]; Notes of meetings with [✂].
[1357] Responses to the CMA's information requests [✂].

Evidence from the Jigsaw Report

5.405 In relation to egress fees, the Jigsaw report is consistent with our customer evidence with respect to switching and multi-cloud: 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 considered as the main or even one of the main barriers to a multi-cloud approach.[1358]

5.406 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 might 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.

5.407 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.[1359] Customers also noted that the increasing importance of AI, alongside the expected larger data volumes, will bring egress fees more into focus in the future. We briefly explore this topic in the section 'Impact of the emergence of AI', below.

5.408 Finally, 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 the 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.[1360] This perception was particularly strong among customers with high data volumes.[1361]

Other customer evidence

5.409 We have also received customer evidence on:

(a) egress fees' contribution to the predictability of cloud spend; and

(b) the impact of the emergence of AI.

[1358] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 5.2.2.
[1359] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 5.2.5.
[1360] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 5.3.6.
[1361] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 5.1.7.

5.410 We briefly address each below.[1362]

Egress fees' contribution to the predictability of cloud spend

5.411 We have considered how the predictability of cloud spend for customers might be relevant for our assessment of whether egress fees are a barrier to switching and/or multi-cloud. We consider the predictability of wider cloud spend and how egress fees may contribute to any difficulties in predicting that spend. We consider that:

(a) If wider cloud spend is hard to predict and it is uncertain whether a switch or a multi-cloud architecture would lead to net benefits, customers are more likely to choose not to use alternative cloud providers, even though the offer may be better for that customer.

(b) Insofar as egress fees contribute to this wider uncertainty, a customer may be more reluctant to switch or multi-cloud if they will not be able to accurately forecast the egress fees they would incur from doing so.

5.412 Overall, there is some evidence that both general cloud spend and egress fees in particular lack predictability. This means that customers may be more reluctant to switch and/or multi-cloud, as they will not be able to accurately forecast egress fees and general spend increases from doing so which makes the net benefits uncertain.

5.413 However, as we do not have evidence on the impact of the lack of predictability on customers' decisions, we have not taken this into account in our assessment of whether egress fees are a barrier to switching and/or multi-cloud.

Impact of the emergence of AI

5.414 The emergence of generative AI may have an impact on how much influence egress fees have on customers' switching and multi-cloud decisions. AI models may mean customers' data storage volumes increase, eg if they have large training data sets or the outputs of the AI model are stored in the cloud. Egress fees may be a material factor in any decision to move this data, either as part of a switch between clouds or as part of a multi-cloud architecture.

5.415 We have received some submissions from customers and a cloud provider that the rise of generative AI means that, in the future, egress fees may become more relevant.[1363]

[1362] For further detail, see Appendix O.
[1363] 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.

5.416 However, as we do not have sufficient evidence at present specifically on the impact of AI on customers' switching and multi-cloud decisions – and therefore on the relevance of egress fees in those decisions – we have not taken this into account in our assessment of whether egress fees are a barrier to switching and/or multi-cloud.[1364]

Internal documents

5.417 Our review of internal documents provides further indications that egress fees are likely to be a concern and potentially a barrier for some customers and can be a factor that can influence customer behaviour and/or decisions. For example:

(a) [✂];[1365]

(b) [✂].[1366] 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.[1367]

5.418 Our review also indicates that AWS, Microsoft and Google set egress fee prices not only based on costs and margins but also considering other factors, including competitors' pricing and the pricing of their other networking services. This indicates 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 focussing on this for individual services (eg egress) in isolation. We have not seen internal documents showing that these cloud providers consider future investment or innovation in making their pricing decisions.

Our assessment of the relevance of egress fees to customers' switching and multi-cloud decisions

5.419 We consider that egress fees influence 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 made difficult for customers as a result of higher costs of making a switch.[1368] Egress fees can therefore disincentivise and/or hinder customers who might be considering to, or attempting to, switch and/or multi-cloud. We summarise our reasoning below.

[1364] We further explore the potential impact of AI on switching and multi-cloud in general in the section on the impact of AI on cloud services in the competitive landscape chapter.
[1365] [✂] response to the CMA's information request [✂].
[1366] [✂] submission to the CMA [✂].
[1367] We have set out the evidence from cloud providers' internal documents on their commercial strategies in relation to egress fees in Appendix P.
[1368] CC3 (Revised), paragraphs 297 and 316.

5.420 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 responses to our customer engagement explicitly identified egress fees as a challenge when considering, or a factor disincentivising, switching. 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 sentiments were mirrored in the Jigsaw report, where some customers expressed concerns about egress fees for switching and/or multi-cloud.

5.421 When assessing customers' views, we consider that egress fees:

(a) do not have to be the main concern, but can be one of a number of concerns and still be meaningful, as long as they influence customers' decision-making, and

(b) do not have to be a relevant factor for a majority of customers to be considered a meaningful barrier to switching and/or multi-cloud. Their influence for a substantial minority of customers is sufficient to suggest that egress fees influence customer decision-making.

5.422 Our analysis of hypothetical scenarios indicates that customers may incur substantial costs from switching and multi-cloud. Specifically:

(a) Our estimates of hypothetical 'one-off' switching costs indicate that egress fees paid to move all of a UK customer's data from one cloud provider to another range between [0-5]% to [5-10]% of a customer's annual spend. The majority of AWS', Microsoft's and Google's UK customers would have had to pay [✂] of their total annual spend in 2022 for a total switch. If spreading costs over a three-year contract, a customer might expect to pay [0-5%] to [0-5%] of their total yearly cloud spend to switch.

(b) About [10-20]% of customers would have had to pay egress fees of more than [10-20]% of their total annual spend in a single year. If spreading costs over a three-year contract, these customers might expect to pay more than [0-5]% of their total yearly cloud spend to switch. [0-5]% - [5-10]% of customers would have to pay egress fees in excess of [5-10]% per year over a three-year contract to switch. These costs can also be materially higher for some customers.

(c) This cost could influence the decision on whether or not to switch cloud provider. Some customers would not have the incentive to switch to a rival cloud provider even if that cloud provider offered a price that was at least [0-5]% (or about [0-5]% over a three-year period) better or equivalently better quality or functionality, thereby reducing the level of competition in the market.

(d) Our estimates of hypothetical multi-cloud costs show that any non-trivial amount of data transferred out of a customer's cloud is likely to lead to non-negligible egress fee costs as a percentage of total monthly cloud spend. For example, for customers of any size, egressing 20% of their data would cost [0-5]% to [10-20]% of their total monthly cloud spend. Smaller customers are disproportionately affected. 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. We do not consider that Google's submission on multi-cloud costs,[1369] for reasons discussed in detail in Appendix M, shows that egress fees are not a barrier to multi-cloud.

5.423 Internal documents indicate that egress fees may be a concern or potential barrier for some customers and may be a factor that can influence customer behaviour and/or decisions. 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. Our review also indicates that there are strategic considerations in the setting of egress fee prices.

5.424 While cloud providers have submitted that falling effective prices for egress are indicative of competition in the market, we do not consider this is enough to conclude that the presence of egress fees cannot constitute a barrier to switching and multi-cloud. Additionally, 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 too.

5.425 We consider that analyses of the observed prevalence of egress fees are not wholly representative of customers' eligibility to pay them. 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.

5.426 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.

5.427 We have found that AWS', Microsoft's and Google's free switching programmes have limited and uncertain scope, especially for larger customers, based on evidence we have reviewed. These programmes generally do not cover multi-cloud use, are limited to a 60-day switching period and do not cover all products.

[1369] Google's submission to the CMA [✂].

We therefore consider that these programmes do not materially affect the conclusions of our analysis of either switching costs or multi-cloud costs. In addition, to the extent customers would incur periods of multi-cloud running in the process of switching, egress fees may represent an even higher proportion of their annual cost of cloud.

5.428 We focus much of our analysis on AWS, Microsoft 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 may 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

5.429 Cloud providers have 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.

5.430 We have 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.[1370]

Role of egress fees in funding investment

5.431 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 have examined the evidence on investment and innovation.

[1370] 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.

Cloud providers' views on investment and innovation

5.432 AWS, Microsoft, Google, IBM and Oracle submitted that egress fees recover costs.[1371] AWS, Microsoft 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 said that regulation which prevents its ability to make these investments would lead to reduced innovation and poorer outcomes for UK businesses.[1372]

(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.[1373]

(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.[1374] 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.[1375]

(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.[1376] Google submitted that a market-wide intervention in egress

[1371] 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.
[1372] AWS' response to the CMA's updated issues statement and working papers, paragraph 32.
[1373] AWS' response to the CMA's Egress fees working paper, 31 July 2024, paragraph 11.
[1374] Microsoft's response to the Competitive landscape, Committed spend agreements and Egress fees working papers, footnote 80.
[1375] Microsoft's response to the Competitive landscape, Committed spend agreements and Egress fees working papers, paragraphs 97-99.
[1376] Google response to the Egress fees working paper, Annex response (h).

fees puts the customer benefits of investment in networking products and infrastructure at risk and competitive market at risk.[1377]

(e) Google submitted that its networking prices reflect the quality and innovation of its networking products, as well as its cost and ongoing investment.[1378] Google also noted that removal of egress fees risks decreased innovation and product quality.[1379]

5.433 AWS, Microsoft 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.[1380] AWS also cited the development of [✂].[1381]

(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 submitted as an example that its investments will allow it to build low-latency infrastructure at higher costs and technical performance. Microsoft also 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 AWS, Microsoft and Google and often in the 'commons' (notably, open source).[1382]

(c) Microsoft also submitted that many cloud providers, including AWS, Microsoft and Google, have invested billions of pounds in building out extensive and sophisticated private networking infrastructure spanning the globe for the benefit of their customers and to supplement the public network.[1383]

(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 if offers innovative networking products such as Cross-Cloud Interconnect and a range of quality and service level options, including

[1377] Google's response to Egress fees working paper, Annex response (h).
[1378] Google's response to Egress fees working paper, paragraph 3(c).
[1379] Google response to the Egress fees working paper, paragraph 57.
[1380] AWS' submission to the CMA [✂].
[1381] AWS' submission to the CMA [✂].
[1382] Microsoft's response to the Competitive landscape, Committed spend agreements and Egress fees working papers, paragraphs 12 and 45 and footnote 115.
[1383] Microsoft's response to the Issues statement dated 17 October 2023, paragraph 36.

networking product add-ons such as additional security and Standard Tier and Premium Tier networking products.[1384]

5.434 Wasabi 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.[1385]

Analysis of the role of egress fees in funding investment

5.435 AWS, Microsoft and Google have submitted that data transfer fees, including egress fees, enable and/or incentivise their investment in networking services and innovation.[1386] We have 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.

5.436 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.[1387]

Analysis of potential benefits

5.437 Based on our analysis of AWS', Microsoft's and Google's submissions,[1388] we see that these cloud providers have made substantial investments in their global network infrastructure. We recognise that 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.

5.438 In terms of innovation, the examples cited by AWS, Microsoft and Google that appear to potentially be innovation are: AWS' development of an automated traffic optimisation and management service for network traffic between AWS and the internet;[1389] Microsoft's investments allowing it to build low-latency infrastructure at higher costs and technical performance and feature updates for its top 20 UK

[1384] Google's response to the Egress fees working paper, paragraphs 19 and 20.
[1385] Wasabi's response to the CMA's working papers, pages 1 to 2.
[1386] 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)
[1387] CC3 (Revised), paragraph 174.
[1388] 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; [✂] submission to Ofcom [✂].
[1389] AWS' submission to the CMA [✂].

cloud services;[1390] and Google's networking products such as Cross-Cloud Interconnect.[1391]

5.439 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.[1392] 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.[1393]

(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.[1394]

(c) Google submitted that [✂]. 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).[1395] We note that these are different products to the egress data transfer via the internet which is the focus of our analysis.

5.440 It is difficult to assess innovation and we have 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 AWS', Microsoft's and 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.

1390 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.
1391 Google's response to the Egress fees working paper, paragraph 20.
1392 AWS' response to the CMA's information request [✂].
1393 AWS' response to the CMA's information request [✂].
1394 Microsoft's response to the CMA's information request [✂].
1395 Google's response to the CMA's information request [✂].

Extent to which investment and innovation is dependent on egress fees

5.441 Of the submissions that AWS, Microsoft and Google have cited as evidence of investment in egress services or data transfer services more broadly, we note that:

(a) For AWS' submissions,[1396] 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.[1397]

(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.[1398]

5.442 The submissions that we have set out above as potential evidence of innovation are also largely for network investment and data transfer services overall. Only AWS' [✂] seems likely to directly benefit egress data transfer services as it [✂] (although this would likely also apply to ingress data transfers and potentially other services).

5.443 Given that the investments and innovation that AWS, Microsoft and Google have submitted largely all relate to the provision of their broader networking services (and/or cloud services overall), we consider this investment to be funded by these services overall.[1399] Whilst investment in providing a 'high quality' network will benefit egress data transfer services, our analysis does not indicate that egress data transfer services are the main user or beneficiary of this. For example, AWS', Microsoft's and Google's internal data transfer volumes are [✂] their egress volumes (excluding CDN egress, as this is a type of egress with different pricing and use case).[1400] [✂].[1401] AWS' CDN egress volumes are [✂] its non-CDN egress volumes.[1402]

[1396] 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 [✂].
[1397] 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.
[1398] Google's response to the Egress fees working paper, paragraph 20.
[1399] 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 [✂].
[1400] CMA analysis of responses to the CMA's information requests [✂].
[1401] Microsoft's response to the CMA's information request [✂].
[1402] AWS' response to the CMA's information request [✂]; AWS' response to Ofcom's information request [✂].

5.444 This is supported by our analysis of cloud providers' costs, as set out in Appendix Q, which finds that 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. AWS, Microsoft and Google each allocate a share of costs associated with network and wider cloud infrastructure to data transfer services.[1403] 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.

5.445 As noted above in our review of AWS', Microsoft's and Google's internal documents, we also did not see evidence that cloud providers considered future investment or innovation in making their data transfer pricing decisions.

5.446 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.

5.447 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 AWS, Microsoft and Google and so not the only source of profits.[1404]

Role of egress fees in passing on cost savings

5.448 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 role of the egress fees in passing on cost savings

[1403] Google only allocates costs to its 'Networking' segment overall, which is [✂]; Google's response to the CMA's information request [✂].
[1404] We have assessed overall cloud business profitability in Chapter 3, market outcomes section.

to customers, in particular the extent to which there is evidence of cost savings being passed on.

Cloud providers' views on pass-on cost savings

5.449 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 [✂] and Microsoft and Google both submitted that egress fee prices have decreased.

5.450 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.[1405]

AWS

5.451 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. 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.[1406] AWS also submitted that this data showed that from 2019 to 2023 [✂]. [1407] AWS also noted that its egress margin [✂], consistent with a pass-on rate exceeding [✂].[1408]

5.452 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.[1409]

5.453 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 cost. AWS said that these cost savings have been passed on to customers, for example through the expansion of AWS' free tier for egress, [✂] for customers [✂] and the introduction of free data transfer for customers who switch.[1410]

5.454 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

[1405] Oracle's response to the Updated Issues Statement and working papers, pages 1-2.
[1406] AWS' response to the CMA's Egress fees working paper issued 23 May 2024, paragraph 5.
[1407] AWS' submission to the CMA [✂].
[1408] AWS' submission to the CMA [✂].
[1409] AWS' submission to the CMA [✂]; AWS' response to the CMA's updated issues statement and working papers, paragraph 33.
[1410] AWS' submission to the CMA [✂].

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.[1411]

5.455 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 [✂].[1412]

Microsoft

5.456 Microsoft submitted that the list prices for its egress products (Bandwidth and ExpressRoute[1413]) have fallen over time. Microsoft submitted the following as evidence of this:[1414]

(a) Microsoft's introduction of internet egress routed via public ISP network as a cheaper egress option to routing via its Premium Network.

(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.

5.457 Microsoft also submitted that on average, real prices for both metered and unmetered egress plans have fallen since 2018.[1415]

Google

5.458 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

[1411] AWS' submission to the CMA [✂].
[1412] AWS' response to the CMA's information request [✂].
[1413] '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.
[1414] Microsoft's submission to the CMA [✂].
[1415] Microsoft's submission to the CMA [✂].

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.[1416]

5.459 Google submitted that [✂].[1417]

Analysis of pass-on cost savings

5.460 In this section, we examine the extent to which there may be a benefit from cost savings being passed on to customers.

5.461 AWS has 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.

5.462 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 have 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

5.463 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.

5.464 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 [✂].[1418]

5.465 AWS has 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

[1416] Google's response to the CMA's Egress fees working paper, paragraph 15.
[1417] Google's submission to Ofcom [✂].
[1418] AWS' submission to the CMA [✂]. AWS' response to the CMA's updated issues statement and working papers, paragraph 33; AWS' submission to the CMA [✂].

note that, having analysed the underlying data from the AWS submission, a [✂] of the costs that AWS include in their analysis of the movement of average egress costs per unit compared to average prices [✂].[1419] These costs [✂] rather than reflecting cost efficiencies or savings related to egress or data transfer services overall.

5.466 AWS has also submitted that its investments in its network [✂].[1420] 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. [✂].[1421]

5.467 Google's example of [✂].[1422] This does not appear to be cost savings arising from investment.

5.468 Similarly, 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.

5.469 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

5.470 AWS, Microsoft and Google have submitted that egress fee prices have decreased. We have identified some issues with the data AWS and Microsoft submitted on egress fee prices,[1423] 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.

5.471 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.[1424]

5.472 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

[1419] CMA analysis of AWS' submission to the CMA [✂]; AWS' response to the CMA's information request [✂].
[1420] AWS' response to the CMA's information request [✂].
[1421] See Appendix Q for further details. Source: CMA analysis of responses to the CMA's information requests [✂].
[1422] Google's submission to Ofcom [✂].
[1423] 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.
[1424] 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.

decrease in AWS', Microsoft's and/or Google's egress fee prices has been the passing on of cost savings resulting from investment.

5.473 We also note that AWS, Microsoft 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 in the last 5 years AWS and Microsoft have each only changed the free tier available to customers once and list prices for AWS non-CDN egress and Microsoft Premium tier egress have remained unchanged, whilst 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.[1425] Microsoft cites one instance of decreasing the price of its ExpressRoute service, in 2021 and [✂].[1426]

5.474 Additionally, we note that AWS cites [✂] for customers [✂] 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 find the free switching programmes introduced by AWS, Microsoft 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, [✂] 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 [✂] are not necessarily permanent price changes for customers.

5.475 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 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

5.476 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.

[1425] Microsoft's submission to the CMA [✂]; Google's response to the CMA's information request [✂]; AWS' response to the CMA's information request [✂]; Announcement of pricing change for egress traffic.
[1426] Microsoft's submission to the CMA [✂]; Microsoft's response to the CMA's information request [✂].

Cloud providers' views

5.477 AWS, Microsoft, Google and another cloud provider submitted that removing egress fees could result in inefficient and excessive egress usage by customers.[1427]

5.478 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.[1428] [✂]. 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.[1429]

5.479 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.[1430]

5.480 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. [✂]. IBM noted that customers' need for an efficient IT infrastructure is only relevant to multi-cloud use, as this is less relevant when switching cloud provider entirely.[1431]

5.481 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.[1432]

(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.[1433]

[1427] 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; [✂] submission to the CMA [✂].
[1428] AWS' response to the CMA's Egress fees working paper dated 23 May 2024, paragraph 18.
[1429] AWS' submission to the CMA [✂].
[1430] [✂] submission to the CMA [✂].
[1431] IBM's response to Egress fees working paper, page 3.
[1432] Google response to the Egress fees working paper, Annex response (h).
[1433] AWS' submission to the CMA [✂].

(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.[1434]

(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.[1435]

Our analysis of inefficient egress risk

5.482 In this section we consider first the extent to which inefficient egress usage by customers is a potential risk to be deterred. We then consider whether there are benefits arising from egress fees deterring inefficient egress usage.

Extent to which inefficient egress usage by customers is a potential risk

5.483 We recognise that each cloud providers' network capacity will generally be a finite resource and that inefficient or excessive egress usage could take up more of this finite network capacity.

5.484 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.

5.485 Cloud providers' submissions are 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.[1436]

[1434] AWS' response to the CMA's Egress fees working paper dated 23 May 2024, paragraph 19.
[1435] Microsoft's response to the competitive landscape, committed spend agreements and egress fees working papers, paragraph 95.
[1436] 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.

5.486 AWS submitted the example [✂].[1437] This does provide 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 customer. For example, [✂] may have been incentivised to have data traffic growth as a measure of business performance.

5.487 We also note that this is an example where [✂].[1438] 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 [✂] when charged at a fixed rate. Additionally, whilst AWS submitted that the 'example shows that, when AWS [✂],[1439] AWS does not actually state that there were any [✂] as a result of CloudFront's approach, just that it was [✂].

5.488 Microsoft and Google did not submit any evidence of inefficient or excessive egress usage.

5.489 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 [✂].[1440] [✂].[1441]

5.490 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 AWS, Microsoft and Google provide. 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.

5.491 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.[1442]

[1437] AWS' submission to the CMA [✂].
[1438] AWS' submission to the CMA [✂].
[1439] AWS' submission to the CMA [✂].
[1440] Microsoft's submission to the CMA [✂].
[1441] AWS' response to the CMA's information request [✂].
[1442] 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

5.492 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 Alliance[1443] 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.[1444] 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.

5.493 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 at paragraph 5.481(b), 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.[1445] We consider 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.

5.494 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

5.495 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 have considered whether there is a customer benefit, by incentivising efficient egress usage that allows

[1443] 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 27 February 2024.
[1444] Civo abolishes all egress fees to advance cloud simplicity strategy (accessed 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 22 October 2024.
[1445] AWS' submission to the CMA [✂].

providers to avoid unnecessary network investment and/or reducing the risk of reduced network quality (ie network availability and resiliency).

5.496 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.[1446] 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 inefficient egress such impacts would materialise, and as set out above, we consider the risk of inefficient egress to be low for switching and unclear for multi-cloud use with mitigating factors.

5.497 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.

5.498 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.

5.499 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',[1447] shows that volume alerts can be used to deter usage. A Microsoft internal document notes that:

> [✂][1448]

5.500 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

[1446] 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.
[1447] AWS' submission to the CMA [✂].
[1448] Microsoft's response to the CMA's information request [✂].

not the sole cause of any resulting benefits. The deterrent effect seems likely to provide material benefits to 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.

5.501 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 egress fees themselves are inducing firms to deter inefficient egress usage, in response to actual or expected actions by rivals.[1449]

Our assessment of the potential benefits of egress fees

5.502 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 many other services. However, we do not find any investment and innovation cited to be dependent or solely driven by the charging of egress fees.

5.503 In relation to AWS', Microsoft's and Google's submissions that egress fee prices have decreased and/or that cost savings are passed on to customers, we consider 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 seen limited evidence that inefficient or excessive egress is a risk, in the absence of egress fees.

5.504 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.

Provisional conclusions

5.505 We have provisionally found that the presence and relevance of egress fees to customers' decisions on switching and multi-cloud means that there is a weakened

[1449] CC3 (Revised), paragraph 174.

customer response to differences in price, service quality and/or innovation between cloud providers.

5.506 This means that customers may have less 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. As a result, egress fees 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 upon entering the market for cloud services.

5.507 While cloud providers have introduced a range of egress-free switching programmes, these have limited and uncertain scope and do not materially affect the conclusions of our analysis of either switching costs or multi-cloud costs.

5.508 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 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.

5.509 Our provisional 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, and that they also reduce the incentives of suppliers to compete for the business of their competitors.

Provisional conclusions on barriers to switching and multi-cloud

5.510 Our provisional view is that cloud providers do not have clear incentives to enable customers to switch and use multiple clouds.

5.511 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. For this reason, we have examined technical barriers and egress fees in more detail.

5.512 The conclusions of our assessment of technical barriers and egress fees are set out in the relevant sections above. In summary:

(a) Technical barriers raise the cost to customers of switching and the ability to integrate use of multi-cloud; and

(b) Egress fees are relevant to customers' decisions to switch and/or multi-cloud and can therefore reduce the ability of and/or incentives for, customers to switch to other cloud providers and/or run a multi-cloud architecture.

6. Licensing

- This chapter assesses whether Microsoft has partially foreclosed its rivals in cloud services through its software licensing practices.

- We have examined whether Microsoft has the ability to foreclose its rivals, whether it has the incentive to do so and whether its conduct has an adverse effect on its rivals. We considered: (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 conduct; and (iv) the impact on rivals' competitive offerings arising from Microsoft's conduct.

- We have provisionally found that Microsoft has significant market power in relation to each of the software products. This is because Microsoft has a moderate to high market share, customers are unable or unwilling to switch away from these products and/or there are limited alternatives.

- 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. We have provisionally found that Windows Server and SQL Server account for a material proportion of AWS' and Google's costs in providing different bundles of cloud services to customers and are important inputs.

- We also 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 we therefore consider that, for these customers, the inputs do not shape downstream competition beyond their significance in the cost base. However, for the subset of 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.

- We have provisionally found 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, VDI workloads currently make up a relatively small proportion of public cloud workloads and the evidence is mixed with regard to whether VDI will increase in importance over time. As such, our provisional view is that these three software products do not constitute important inputs on their own but rather contribute to the overall importance of the Microsoft software products as inputs to cloud services.

- We have examined the price and non-price differences between the software that Microsoft offers Azure customers and that which it supplies to its rivals in cloud services. We have provisionally found that the wholesale price paid by AWS and Google for Windows Server and SQL Server is higher than Microsoft's retail price to its cloud customers that have Windows Server and SQL Server Licences that qualify for the Azure Hybrid Benefit. This is evidence of the significance of Microsoft's conduct and the potential for it to disadvantage its rivals in cloud services.

- In relation to Windows 10/11 and Microsoft 365, Microsoft does not make this software 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 Desktop and Microsoft 365 on Google or AWS.

- We have assessed the impact of Microsoft's conduct on rivals' competitive offerings. 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. In addition, we found that usage of Windows Server and SQL Server on Azure is significantly higher than on AWS and GCP, although this does not demonstrate a causal link between Microsoft's conduct and relative usage. But taking this together with our market power findings, we consider that there is a higher likelihood that the differences in relative usage are at least partially driven by Microsoft's conduct.

- We have provisionally found that 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 the Microsoft's conduct relating to these software products contributes to an effect on AWS' and Google's competitive offerings.

- In summary, our provisional conclusion is that all three conditions of partial foreclosure are satisfied and that Microsoft has partially foreclosed AWS and Google in cloud services through its software licensing practices.

Background

6.1 Until recently, customers purchased licences for software and operating systems to be installed and used on their premises. 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.

6.2 Ofcom received submissions regarding the software licensing practices of some cloud providers, in particular Microsoft.[1450] The submissions raised concerns that

[1450] CMA, Issues statement, paragraph 32.

software providers – and in particular Microsoft – had made their software more expensive, had fewer features and/or worked less effectively when run on a rival's cloud infrastructure, and that this had disadvantaged rival cloud providers.[1451]

6.3 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 above in the Competitive landscape chapter there are indicators that Microsoft has significant market power in cloud services. As such, there is potential for Microsoft's licensing practices to have an adverse effect on competition.

6.4 Although we received some submissions relating to Oracle's software licensing practices, we have decided not to investigate these because:

(a) Oracle's market share of cloud is relatively small ([✂]% for IaaS and [✂]% 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 respondents used Oracle as their sole cloud provider, and their main use was as a secondary cloud, for example, for supporting legacy systems.[1452] 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;

(b) we have received far fewer submissions raising concerns about Oracle's licensing practices; and

(c) the provision of cloud infrastructure services is complex, and the CMA's resources are limited. As such, we prioritised the use of those limited resources to the areas where there is the potential for greater harm to arise.

6.5 Ofcom identified five software products as potentially relevant to the consideration of Microsoft's software licensing practices: Microsoft's Windows Server (which includes Active Directory functionality), SQL Server, Windows 10/11, Visual Studio and Microsoft 365/Office.

6.6 We have also focused on these five products because licensing concerns were raised in relation to these. We also understand that customers who have workloads involving these software products, which have migrated to the public

---

[1451] CMA, Issues statement, paragraphs 32 and 33.
[1452] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 3.4.8.

cloud, are more likely to 'lift and shift' those workloads compared to others, and continue using these software products in the cloud.[1453]

Description of the software licensing practices

6.7 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

6.8 While our analysis has focused on Microsoft's current licensing practices, it is useful to explain how these have developed.

Pre-2019

6.9 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.[1454,1455] 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.[1456,1457]

2019 changes

6.10 Microsoft modified its licensing terms in 2019:[1458] it created a new category of 'Listed Providers' of cloud infrastructure services. These were Microsoft, Alibaba, Amazon, and Google. Customers of Listed Providers could now no longer use

---

1453 [✂] submission to the CMA [✂]; [✂] submission to the CMA [✂] response to Ofcom's information request [✂]. Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), page 29. Lift and shift applications workloads.
1454 [✂] submission to the CMA [✂]. [✂].
1455 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 18 November 2024. A cloud provider submitted that customers were also able to BYOL specific Microsoft products to the shared hardware of non-Listed Providers that were part of Microsoft's Qualified Multitenant Hoster Program.
1456 [✂] submission to the CMA [✂].
1457 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.
1458 [✂] submission to the CMA [✂]; Updated Microsoft licensing terms for dedicated hosted cloud services, accessed 18 November 2024.

their pre-existing licences for Microsoft software on a BYOL basis on dedicated hardware of Listed Providers.[1459,1460]

6.11 Microsoft submitted that it updated its licensing terms to correct a ‘licensing loophole’.[1461] 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).[1462,1463]

6.12 Microsoft said that its 2019 changes to its licensing terms weren’t 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.[1464]

6.13 We have also 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.[1465]

(b) Google said that the list was relatively arbitrary,[1466] and that the Listed Providers were Microsoft’s largest competitors and likely among the most capable of competitively constraining Microsoft.[1467]

6.14 Microsoft submitted that [✂].[1468] Microsoft added that other players do not come close to the scale of Listed Providers and [✂].[1469] In this regard, Microsoft said that, based on scale, [✂].[1470]

[1459] [✂] submission to the CMA [✂]; [✂] Updated Microsoft licensing terms for dedicated hosted cloud services, accessed 18 November 2024. [✂]. [✂] submission to the CMA [✂].
[1460] 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 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. [✂] submissions to the CMA [✂].
[1461] Microsoft’s submission to the CMA [✂].
[1462] Microsoft’s submission to the CMA [✂].
[1463] ‘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’s submission to the CMA [✂].
[1464] Microsoft’s submission to the CMA [✂].
[1465] CISPE’s response to issues statement [✂].
[1466] Note of meeting with Google [✂].
[1467] Google’s response to the CMA’s information request [✂].
[1468] Microsoft’s submission to the CMA [✂].
[1469] Microsoft’s submission to the CMA [✂].
[1470] Note of meeting with Microsoft [✂].

6.15 Microsoft also said it had [✂] 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.[1471]

2022 changes

6.16 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 Microsoft's licensed software on their cloud infrastructure, and also that they could not use some versions of Microsoft's products.[1472]

6.17 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.[1473] 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.[1474]

6.18 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.[1475]

6.19 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.[1476] 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.[1477]

6.20 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

[1471] Microsoft's submission to the CMA [✂].
[1472] [✂] submission to the CMA [✂].
[1473] Microsoft's submission to the CMA [✂]. 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 18 November 2024.
[1474] See, Microsoft responds to European Cloud Provider feedback with new programs and principles - EU Policy Blog. Accessed 18 November 2024.
[1475] See, New licensing benefits make bringing workloads and licenses to partners' clouds easier; New options for partner hosted cloud, accessed 8 May 2024.
[1476] Microsoft's submission to the CMA [✂].
[1477] Microsoft's submission to the CMA [✂].

Microsoft licences from independent managed service providers if those providers host their services on Listed Providers' clouds.[1478]

Using Microsoft software products on public cloud

6.21 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.

The providers

6.22 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

6.23 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.[1479] 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'.[1480]

Cloud Solution Provider program

6.24 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).[1481] Over time, and subject to partner feedback, Microsoft has enabled Microsoft 365 Apps sold by cloud solution

---

[1478] See, New licensing benefits make bringing workloads and licenses to partners' clouds easier, accessed 18 November 2024; [✂] submission to the CMA [✂].
[1479] Microsoft's response to the CMA's information request [✂].
[1480] Microsoft's response to the CMA's information request [✂].

[1481] Microsoft's response to the CMA's information request [✂].

providers (CSPs) via this program to also be run in non-Listed Provider clouds.[1482] 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.[1483]

6.25 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.[1484] 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 way the licences for the Microsoft software products must belong to the end client.[1485]

Alternative cloud providers or independent managed service providers

6.26 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).[1486] The alternative cloud provider acquires the licences to use the Microsoft software in its own cloud services through a Services Provider Licensing Agreement (SPLA**)**.[1487]

6.27 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 providers' cloud.[1488] As noted above, from 1 October 2025, independent managed service providers will not be able to host these licences on Listed Providers' clouds.[1489]

The routes to obtaining the right to use the Microsoft software

6.28 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.

[1482] Microsoft's response to the CMA's information request [✂].
[1483] Microsoft's response to the CMA's information request [✂].
[1484] Microsoft's response to the CMA's information request [✂].
[1485] Microsoft's response to the CMA's information request [✂].
[1486] [✂] response to Ofcom's information request [✂]; [✂] submission to the CMA [✂].
[1487] Microsoft's response to Ofcom's information request [✂]; Microsoft's submission to the CMA [✂]; Microsoft's response to the CMA's information request [✂].
[1488] [✂] submission to the CMA [✂].
[1489] [✂] submission to the CMA [✂]; New licensing benefits make bringing workloads and licenses to partners' clouds easier, accessed 18 November 2024.

6.29 The main differences between the two routes are explained below.

BYOL

6.30 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).

6.31 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 to use those on-premises Microsoft product licences on non-Listed Providers' public or private clouds (ie they can make use of BYOL).[1490]

6.32 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.[1491] 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.[1492]

SPLA

6.33 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.[1493] 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.[1494] 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.[1495]

6.34 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) depends 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.

[1490] Microsoft's response to Ofcom's information request [✂]; Microsoft's submission to the CMA [✂].
[1491] Microsoft's submission to the CMA [✂]; Microsoft's response to Ofcom's information request [✂].
[1492] Microsoft's submission to the CMA [✂]. See: Microsoft licensing programs and License mobility through software assurance, accessed 18 November 2024.
[1493] Microsoft's submission to the CMA [✂]; Microsoft's response to Ofcom's information request [✂].
[1494] Microsoft's submission to the CMA [✂].
[1495] Microsoft's submission to the CMA [✂].

Differences between using Microsoft software products on Azure compared to on non-Azure clouds via SPLA

6.35 We received a number of submissions setting out issues and concerns relating to Microsoft's software licensing practices, including both price and non-price factors.

6.36 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 route, the cloud provider will need to purchase a licence via its SPLA and the customer will need to purchase licence included services.

6.37 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.[1496] 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'.[1497]

6.38 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.[1498]

6.39 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.[1499] 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. [1500] 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'.[1501]

6.40 The submissions relating to the non-price factors are wide-ranging and include (but are not limited to):

[1496] Listed Providers, accessed 18 November 2024; Submissions to the CMA [✂].
[1497] [✂] submission to the CMA [✂].
[1498] Azure Hybrid Benefit - Hybrid Cost Calculator, accessed 18 November 2024; See Explore Azure Hybrid Benefit for Windows VMs and Azure Hybrid Benefit - Azure SQL Database & SQL Managed Instance, accessed 18 November 2024.
[1499] [✂] response to Ofcom's information request [✂].
[1500] Microsoft's submission to the CMA [✂]; Responses to the CMA's information requests [✂].
[1501] Azure Hybrid Benefit for Windows Server, accessed 18 November 2024.

(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 Desktop Multisession).

Conceptual framework

6.41 We have examined whether Microsoft is able to foreclose its rivals in cloud services by means of raising their costs, restricting access to essential inputs (or degrading the quality of inputs) or making a significant proportion of customer demand less contestable.

6.42 As a provider of both software products and the public cloud services on which customers can use them, Microsoft is vertically integrated. Despite their potential to enhance efficiency and consumer welfare, vertical relationships can also lead to an AEC in a market, particularly by allowing the firms to foreclose rivals' access to inputs and customers and/or otherwise have a dampening effect on competition.[1502]

6.43 For a vertically related firm, foreclosure may be achieved by practices that restrict access to essential inputs or raise rivals' costs, or limit rivals' ability to acquire sufficient customers to benefit from economies of scale, learning effects and/or network effects.[1503]

6.44 Foreclosure of access to key inputs (input foreclosure) may lead to a reduced competitive constraint on a vertically related firm. When deciding whether to supply its competitors downstream with key inputs, a vertically integrated firm may take into account how these sales would affect the profits of its own downstream division. If it has significant market power in the upstream market, the firm may have an incentive to refuse access to the input or to raise its price, and consequently increase the costs of competing downstream firms. By being subjected to higher input prices, downstream competitors may be unable to compete effectively. As a result of such foreclosure effects a vertically integrated 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.

[1502] CC3 (Revised), Guidelines for market investigations: Their role, procedures, assessment and remedies, paragraphs 265 to 267.
[1503] CC3 (Revised), paragraph 268.

6.45 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.[1504]

6.46 The CMA will look at a variety of evidence in conducting its assessment of the overall impact of the vertical relationships on competition. Analysis of profitability and financial data can also help provide an insight into whether the foreclosure might be a profitable strategy.[1505]

6.47 Our concern is that Microsoft's licensing practices, as described above, may partially foreclose its 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; and that those licensing practices therefore—either on their own or in combination with other features of the market—prevent, restrict or distort competition.

6.48 This conduct could potentially harm competition in two ways:

(a) Any practices that make software licences more expensive when used with rival clouds compared to Microsoft's Azure service may raise rivals' costs of supplying cloud services. Microsoft's rivals, acting as profit-maximising firms, may be induced to pass a proportion of these higher costs on to their customers, thereby weakening the competition faced by Microsoft.

(b) Any practices that make a significant proportion of customer demand less contestable to rivals may, over the longer term, weaken rivals' ability to benefit from scale advantages in supplying cloud services, such as economies of scale, learning effects and/or network effects.[1506]

6.49 Microsoft submitted that various steps need to be satisfied to establish partial input foreclosure. It said that the critical 'must-pass' step is to establish that the upstream input is 'important' to downstream competition, what Microsoft calls the 'materiality test'.[1507] Microsoft noted that '[i]mportance can mean a couple of different things but only one type of "importance" is (...) relevant here: the cost element in raising rivals' costs with respect to Microsoft's software.'[1508]

6.50 We have examined whether Microsoft has the ability to foreclose its rivals, whether it has the incentive to do so and whether its conduct has an adverse effect on competition. Our statutory task in the context of this market investigation is to

1504 CC3 (Revised), paragraph 273.
1505 CC3 (Revised), paragraph 276.
1506 CC3 (Revised), paragraph 268.
1507 Microsoft's response to the licensing working paper dated 10 July 2024, paragraph 2.15
1508 Microsoft's response to the licensing working paper, dated 10 July 2024, paragraph 2.15.

decide whether any feature, or combination of features, prevents, restricts or distorts competition pursuant to the Enterprise Act 2002 and not whether there has been an infringement of the Competition Act 1998. Therefore, our provisional findings do not reach any view as to whether there has been a breach of the Competition Act 1998. We have set out evidence relevant to our assessment in the following sections:

(a) Microsoft's market power in related software markets;[1509]

(b) the importance of Microsoft software as inputs to cloud services;

(c) Microsoft's conduct; and

(d) The impact on rivals' competitive offerings arising from Microsoft's conduct.

6.51 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 provisional conclusions for each of these sections; we then provisionally conclude on ability, incentive and effect in the round at the end of the section.

6.52 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. As such, our analysis focused on AWS and Google.

Microsoft's market power in related software markets

Introduction

6.53 For a vertically integrated firm to foreclose its rivals, it must have significant market power in one or more markets along the supply chain.[1510]

6.54 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.

6.55 This section sets out:

[1509] 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 the Competition Act 1998.
[1510] CC3 (Revised), paragraph 274.

(a) the frameworks we used to define the relevant markets and to assess the extent of any market power held by Microsoft in the supply of software; and

(b) evidence and analysis relevant to the assessment of market definition and market power in relation to each of the five relevant software markets.

Framework

6.56 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, structured as follows:

(a) considerations we have applied within the frameworks for market definition and market power; and

(b) cross-cutting considerations relevant to the sources of evidence used in our assessment of market definition and market power.

Market definition

6.57 The principles of our approach to market definition, 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. One additional consideration that is relevant to our assessment of market definition in software markets is the concept known as the 'cellophane fallacy'.[1511]

Market power

6.58 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 prices that are above competitive levels, or output or quality levels that are below competitive levels.

6.59 The firm or firms holding market power will have the ability and incentive to influence market outcomes and other important aspects of competition[1512] 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.

[1511] See Appendix R.
[1512] CC3 (Revised), paragraph 178.

6.60 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.[1513]

6.61 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.[1514]

6.62 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.

6.63 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.

6.64 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.

6.65 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.[1515]

Evidence used in each assessment

6.66 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 views from customers and providers. Below we also consider some cross-cutting considerations relevant to interpreting these sources of evidence.

Market shares

6.67 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

[1513] CC3 (Revised), paragraph 185.
[1514] CC3 (Revised), paragraph 178.
[1515] 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.

less incentive to compete vigorously with its rivals (particularly if there are barriers to entry).[1516]

6.68 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.[1517]

6.69 Market shares depend on market definition and therefore can be subject to the binary fallacy,[1518] as well as the cellophane fallacy. Market shares should therefore be interpreted in the context of those factors. We have noted in our assessment whether market shares may overstate or understate market power.

Product characteristics

6.70 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.

6.71 Nevertheless, we recognise that product characteristics should be interpreted cautiously as they do not map directly to substitutability. Two products may have different characteristics and still be substitutable, or two products may ostensibly satisfy the same broad requirement and yet not serve as particularly good substitutes from a customer's perspective.[1519] We have taken these factors into account when considering product characteristics and alongside other types of evidence.

Customer and provider views

6.72 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. We also asked Microsoft and other providers about competitive

---

[1516] CC3 (Revised), paragraph 187.
[1517] CC3 (Revised), paragraph 192.
[1518] The binary fallacy refers to a situation whereby an overemphasis on market definition can lead to the incorrect assumption that competitor products outside the relevant market exert no constraint on those within the market.
[1519] 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 may be where high switching costs or a weak customer response may mean that substitutability may be limited even in the presence of similar product characteristics.

constraints they face in the relevant markets. Understanding the market from the perspective of providers can help with interpretation of other forms of evidence (such as customer responses).

6.73 For each of the relevant software products, we asked Microsoft's 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.[1520] Due to the possibility that the cellophane fallacy may limit the usefulness of this evidence in relation to market definition, we have considered this evidence carefully and used it as part of the market power assessments.

6.74 It is important to note that we asked customers to think about an increase in prices from the currently prevailing prices. If those prevailing prices have already been set at a level that is above the price level under fully effective competition, then responses may signal a greater willingness to switch away than would be the case under competitive prices, meaning responses may understate the overall extent of market power.[1521]

6.75 As set out above, the customer evidence we have collected is qualitative and so we have given a narrative summary of the key points that we consider emerge from the evidence.

6.76 As noted above, we asked customers about their use of the products on the public cloud. Most customers answered with this in mind, but some may also or only use the product on non-public cloud. Where we refer to evidence from customers that use the product on non-public cloud, we have indicated this below and explained why there is a link to the public cloud.

(a) Depending on the particular product, examples of evidence that may indicate this type of link include customer reasons for choosing the Microsoft product being related to history, experience, past installed base; or if a customer finds it difficult to switch away from the Microsoft product when moving to cloud. We also have evidence on whether customers reported previously using the relevant Microsoft software products on-premises, before moving to the public cloud.

(b) It may be that the competitive conditions on the public cloud and on-premises are similar. For example, where there is a very similar list of alternatives, and

---

[1520] 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.
[1521] This would create some ambiguity in how to interpret responses, particularly if customers demonstrate a high willingness to substitute away from Microsoft products. However, where customers demonstrate a low willingness to substitute away from Microsoft products this is less likely to affect the interpretation of the responses because the fact that they are understated would in that case be consistent with the same interpretation.

we have reason to believe the alternatives serve a similar level of substitutability to the relevant Microsoft products on both deployment types.

6.77 We note in particular that the use case on the public cloud for Office/Microsoft 365, Windows Desktop and Visual Studio is when the products are installed on a 'virtual desktop' as opposed to a physical desktop. A Virtual Desktop Infrastructure (VDI) allows a customer to run applications on remote virtual machines, which are accessed over the internet by end users using 'thin client' software.[1522] This could, for example, allow customers to run computationally-intensive applications on virtual servers instead of relatively more expensive dedicated workstations, and eliminate the need to transfer large amounts of data between cloud and client virtual machines.[1523] In summarising customer evidence, we make the distinction between customers who use these products on the public cloud (for which the primary use case is 'on VDI') and customers who use these products through non-public cloud.[1524]

Product background

6.78 In this section, we describe each of the relevant Microsoft software products and outline their use cases on the public cloud.

Microsoft Windows Server

6.79 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.[1525]

6.80 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 Linux and UNIX OSs. Windows Server is proprietary, and Linux distributions are open-source.[1526] 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.[1527]

[1522] What is Thin Client, accessed 18 October 2024.
[1523] Set up an auto scaling virtual desktop infrastructure (VDI) by using NICE EnginFrame and NICE DCV Session Manager, accessed 18 October 2024.
[1524] Where we have information about customers use of VDI on the public cloud, we use this to inform our categorisation of customer evidence. Further detail about how we consider the customer evidence for Windows Desktop, Visual Studio and Office/Microsoft 365 can be found in Appendix R.
[1525] Panek C (2019), Windows Server Administration Fundamentals, Wiley.
[1526] Dalheimer, MK and Welsh, M (2005), Running Linux, 5th Edition, O'Reilly.
[1527] 'Virtual Machines (VMs) for Linux and Windows', accessed 18 November 2024.

6.81 There are various ways in which customers use Windows Server on a virtual machine 'in the cloud'.[1528]

6.82 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.

6.83 Windows Server integrates well with Windows Desktop. For example, one cloud provider submitted that Windows Desktop has superior interoperability (including in speed) with Windows Server. [1529]

Microsoft SQL Server

6.84 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.

6.85 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 PaaS service). Proprietary RDBMS can also be provided in a managed service (such as Microsoft's Azure SQL[1530]).

6.86 While SQL Server can also be run on Linux server OSs, Microsoft submitted that it is designed to integrate with Windows Server[1531] and a cloud provider submitted that customers are most likely to run this workload on Windows Server.[1532] We also understand that SQL Server works well with other Microsoft products and technologies,[1533] and Microsoft has also submitted that it is designed to integrate

[1528] 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 18 November 2024; [✂] submission to the CMA [✂]; responses to the CMA's information requests [✂].
[1529] [✂] submission to the CMA [✂].
[1530] 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 16 October 2024).
[1531] Microsoft's response to the CMA's information request [✂].
[1532] [✂] response to the CMA's information request [✂].
[1533] For example, see PostgreSQL vs. SQL Server: What's the difference? accessed 18 November 2023. This is discussed further by some customers in SQL Server market power assessment, customer responses.

well with other Microsoft products that customers use together with SQL Server, such as Windows Server, Azure and Microsoft Dynamics.[1534] 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).[1535]

Microsoft Windows 10/11

6.87 Desktop OS software is designed to run a personal computer's hardware and provides a platform for the use of application software.[1536]

6.88 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 Unbuntu. 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 virtual desktop infrastructure (VDI) cloud service, such as AWS' WorkSpaces, can allow its staff to access a Windows 10/11-based virtual desktop;[1537] 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.[1538]

Microsoft's productivity suites

6.89 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.[1539] This package is available on a one-time purchase basis for one PC or Mac.

[1534] Microsoft's response to the CMA's information request [✂].
[1535] Software requirements for Database Servers for SharePoint Server Subscription Edition, accessed 3 October 2024.
[1536] [✂] response to the CMA's information request [✂].
[1537] 'What is Amazon WorkSpaces?', AWS online documentation, accessed 18 November 2024.
[1538] Submissions to the CMA [✂]; 'Windows 10 or Windows 11 Enterprise multi-session remote desktops', accessed 15 April 2024.
[1539] 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 18 September 2024).

(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.[1540] 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.[1541] The additional apps vary dependent on specification. The popular Microsoft 365 for Enterprise packages (E3/5) include Windows Desktop (and the ability to deploy this in the cloud on Azure), security capabilities and advanced identity and access management (including Entra ID).

6.90 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.[1542]

6.91 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,[1543] and there is some limited Microsoft 365 functionality available on Amazon Workspaces.[1544]

Microsoft Visual Studio

6.92 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 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.

[1540] See Compare Office 365 Enterprise Pricing and Plans, accessed 18 November 2024.
[1541] See Compare Microsoft 365 Enterprise Plans, accessed 18 November 2024.
[1542] See Compare Office 365 Enterprise Pricing and Plans, accessed 11 October 2024.
[1543] Note, this functionality is only available for those using Microsoft 365 in Azure.
[1544] 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 23 May 2024. This is discussed in more detail below.

6.93 Visual Studio is no longer supported for Mac OS,[1545] and is not supported for Linux.[1546] Therefore Mac OS or Linux customers would need to use Visual Studio Code,[1547] or Visual Studio on Windows Desktop or choose an alternative IDE. Visual Studio is also commonly used to develop business applications which run on Windows Server and Windows Desktop.

6.94 We understand that customers either use Visual Studio:

(a) on-premises;[1548] 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.[1549]

How the Microsoft products are purchased

6.95 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).

6.96 Many of the customers whose responses we consider below reported having an EA with Microsoft.[1550] 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.[1551] 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 [✂].[1552]

6.97 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.[1553]

6.98 For some customers, this purchasing factor may influence their decisions around consumption of the Microsoft products. We consider the role of EAs further below.

[1545] In August 2023, Microsoft announced that it would deprecate Visual Studio for Mac, from 31 August 2024 'Visual Studio for Mac Retirement Announcement', accessed 18 November 2024.
[1546] Visual Studio 2022 Downloads, accessed 18 November 2024.
[1547] Visual Studio 2022 Downloads, accessed 18 November 2024.
[1548] [✂] response to the CMA's information request [✂].
[1549] [✂] response to the CMA's information request [✂].
[1550] Responses to the CMA's information requests [✂].
[1551] Enterprise Agreement | Microsoft Volume Licensing, accessed 18 September 2024.
[1552] Microsoft's response to the CMA's information request [✂].
[1553] Microsoft's response to the CMA's information request [✂].

Stakeholder responses to our Licensing working paper

6.99 In response to our Licensing working paper, 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]".'[1554]

(b) AWS said 'its customers [have continued] concerns over use of Microsoft's numerous critical software products using the IT infrastructure of their choice'.[1555]

(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'.[1556]

6.100 An academic disagreed with our approach to market definition.[1557] We note this submission was not supported by accompanying evidence. We have followed our guidance in defining markets and case law, which is to start with the narrowest plausible market and then consider whether the market should be widened with reference to demand and supply side substitution.[1558] 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.[1559,1560]

6.101 Some stakeholders commented on the relevance of on-premises workloads.

[1554] Google's response to the Licensing working paper, 24 July 20204, paragraph 5.
[1555] AWS' response to the Licensing working paper, 1 July 2024, page 23-24.
[1556] Social Market Foundation's response to the Licensing working paper, page 22.
[1557] 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 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, 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 1).
[1558] See CC3 (Revised), paragraph 130.
[1559] See CC3 (Revised), paragraph 133.
[1560] In response to the Licensing working paper, Microsoft did not provide a particular view on our assessments of market definition. Microsoft response to the CMA's Licensing working paper, 10 July 2024.

(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.[1561]

(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 (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.'[1562] It submitted that in its view, shares on premises provide a strong indication of the competitive impact that Microsoft's licensing practices will have as the migration from on-premises to cloud continues.[1563]

(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 [✂] 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.[1564]

(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]'.[1565]

6.102 Microsoft noted that Amazon and Google are strong in other areas of their business. It said: 'Amazon and Google have competitive advantages (industry standard proprietary Linux OS from first-mover, Amazon, which it does not license; proprietary ad data from Google, which it does not license) but Microsoft concentrates its energies on competition on the merits – beating them in the market.'[1566]

6.103 We note that it would be difficult to argue Amazon has market power in the server OS market because: (i) it does not earn revenue on its server OS product; (ii) the

[1561] Microsoft response to the CMA's Licensing working paper, 10 July 2024, paragraph 5.1.
[1562] Google's submission to the CMA [✂].
[1563] Google's submission to the CMA [✂].
[1564] [✂] submission to the CMA [✂].
[1565] AWS' submission to the CMA [✂].
[1566] Microsoft's submission to the CMA [✂]. Microsoft also made a similar point in Microsoft's response to the CMA's Competitive landscapes working paper dated 1 July 2024, paragraph 23-31.

product is not strongly differentiated; (iii) it is mainly used on the cloud;[1567] and (iv) its market share is likely to be small. Google has a small market share in cloud, so any potential consumer harm from any attempted foreclosure strategy is likely to be small. Google's involvement in the online advertising market is out of scope of this market investigation, and is currently the focus of a CMA CA98 case.[1568]

Market definition and market power assessments

6.104 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.

Geographic market

6.105 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).[1569] The geographic market may be local, regional, national or wider.[1570]

6.106 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.[1571]

[1567] The original version of Amazon Linux was only made available for use within Amazon's EC2 cloud service (see 'Amazon Linux AMI FAQs, accessed 18 November 2024). The two newer versions, Amazon Linux 2 and Amazon Linux 2023, can be downloaded and used outside of Amazon's cloud services (see 'Linux from AWS', accessed 18 November 2024).
[1568] Refer: Investigation into suspected anti-competitive conduct by Google in ad tech.
[1569] CC3 (Revised), paragraph 145.
[1570] CC3 (Revised), paragraph 146.
[1571] CC3 (Revised), paragraph 147.

6.107 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.

6.108 Our provisional view is that there is a global geographic market for all the relevant products.[1572]

Microsoft Windows Server

6.109 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.

Product market definition

Provider submissions

6.110 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.[1573]

(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.[1574]

[1572] 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.
[1573] Microsoft's response to the CMA's information request [✂].
[1574] AWS' response to the CMA's information request [✂].

(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).[1575]

6.111 Views from providers suggest that the relevant market should not be expanded to include desktop OSs.

Customer submissions

6.112 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 we contacted identified other server OSs.[1576] No customers suggested that a desktop OS would serve as an alternative. This suggests that customers do not view them as substitutes on the demand side.

6.113 We asked the same customers (that use Windows Server) which other server OSs they used, if any. All customers submitted that they currently used both Windows Server and various Linux distributions (ie versions), suggesting that customers may value differentiated functionality of Linux.[1577] We considered whether each server OS could be used for different purposes, and that evidence is explained below.

Provisional view on the product market for Windows Server

6.114 Our provisional 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

6.115 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.

[1575] IBM's response to the CMA's information request [✂].
[1576] Responses to the CMA's information requests [✂].
[1577] Responses to the CMA's information requests [✂].

6.116 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 6.1: Microsoft's market share in server OS, variety of measures

Measure	*Deployment type*	*Source*	*Description*	*Microsoft's market share*
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 license 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]% [1578]

6.117 Microsoft has a substantial share of server OS 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 couple of years[1579] and the rate of 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).

6.118 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.

[1578] Source: CMA analysis of Microsoft's response to the CMA's information request [✂].
[1579] See Azure share of Windows Server compute analysis in Appendix R.

Product characteristics

6.119 We considered evidence from customers, software providers and academic research to understand the product characteristics of Windows Server. Individual distributions of Linux or UNIX may be free or paid for, as shown in the table below.

Table 6.2: Types of open-source server OSs (paid and free)

	Linux	*UNIX*
Paid for	Ubuntu (Ubuntu Pro for more than 5 machines)[1580] SUSE Linux Enterprise Server[1581] Red Hat Enterprise Linux[1582] Amazon Linux (included within Amazon EC2 and AWS charges)[1583]	IBM AIX[1584] Oracle Solaris[1585]
Free	Debian[1586] CentOS Linux[1587]	

Source: company websites (see footnotes 131-138) plus CMA Analysis

6.120 In order to interpret whether lack of product substitutability by functionality may act as a source of market power, we considered customers' reasons for choosing the Microsoft products to understand whether customers value the differences, ie whether these differences are relevant drivers of choice. We asked customers to explain the reasons they chose Windows Server rather than the alternatives they listed. Some customers provided reasons that relate to functionality of Windows Server. Alternatives that lack this functionality may therefore be less substitutable.[1588]

(a) Many customers we contacted said other software or applications require it or integrate with it.[1589]

(b) One customer said security and technical support was an additional reason for choosing Windows Server.[1590]

6.121 There is also evidence of product characteristics that may make other OSs weak substitutes, such as associations with different application building languages and

[1580] Ubuntu Pro: security, compliance and support, accessed 18 November 2024.
[1581] SUSE Linux Enterprise Server, accessed 18 November 2024.
[1582] Red Hat Enterprise Linux Server, accessed 18 November 2024.
[1583] Amazon Linux 2 FAQs, accessed 18 November 2024.
[1584] IBM Power AIX, accessed 12 December 2024.
[1585] Free with Oracle cloud product (The Oracle Solaris 11.4 operating system, accessed 18 November 2024) and paid for on non-Oracle hardware (Oracle Solaris, accessed 18 November 2024).
[1586] About Debian, accessed 18 November 2024.
[1587] About CentOS, accessed 18 November 2024.
[1588] Other reasons are discussed below section 'Customers' views' and Appendix R.
[1589] Responses to the CMA's information requests [✂].
[1590] To note this should have slightly less weight as the customer uses Windows Server on non-public cloud only. [✂] response to the CMA's information request [✂].

frameworks.[1591] Some customers told us that applications designed for one server OS will not run on another.[1592]

6.122 [✂].[1593] 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.[1594]

6.123 The above evidence suggests that there are some indications that Windows Server is used for different purposes to the next-closest products, Linux server OSs, for example running applications that require Windows Server. Therefore, this suggests Windows Server is differentiated and alternative server OSs may therefore be less close substitutes.

Customer views

6.124 Evidence from customers showed that:

(a) Most customers said they would be unlikely to move away from Windows Server in response to a 5% price rise. 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.

(b) In terms of alternatives to Windows Server, most customers generally listed server OSs from the Linux family, and some said there were none.

(c) Reasons customers gave for choosing Windows Server 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'. [1595]

Provider submissions

6.125 Key points from submissions received from providers are summarised below:

(a) Providers consider other server OS as competitors to Windows Server.

(b) One provider said there are some use cases for which Linux cannot be a substitute for Windows Server.

[1591] 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 18 November 2024).
[1592] Responses to the CMA's information requests [✂].
[1593] [✂] submission to the CMA [✂].
[1594] AWS' response to the CMA's information request [✂].
[1595] Further customer evidence is presented in Appendix R.

(c) Providers had mixed views on whether there are barriers to entry and expansion.[1596]

Provisional conclusions on Windows Server

6.126 We have provisionally found that Microsoft has a significant degree of market power in relation to Windows Server because Microsoft appears to have a high share of the server OS market and most customers would be unlikely to switch away from Windows Server in response to a small but significant price increase. The substitutability from Windows Server to other forms or server OS is limited.

6.127 This view is strengthened by the links between Windows Server and other Microsoft software outlined by customers.

6.128 Our view would not change, even if we had defined a wider or narrower market. We conducted this assessment on a market for server OS. There was scope for Windows Server to be in its own market because there were some limits to substitution between Windows Server and other server OS products, and we took these into account in our assessment. Microsoft would hold the entire share of a market defined on this basis. If the market had been widened to include desktop OSs, this would not make a difference to our view on market power, as desktop OSs serve a different use case and customers did not view them as substitutes. Customers would be unable or unwilling to switch away from Windows Server regardless of the frame of reference.

Microsoft SQL Server

6.129 For the purposes of this investigation, we consider that RDBMS is narrowest plausible candidate market which SQL Server sits within.

6.130 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.

Product market definition

Provider submissions

6.131 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.

[1596] Further provider evidence is presented in Appendix R.

(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.[1597]

(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').[1598]

(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.[1599] 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 [✂].[1600]

Customer submissions

6.132 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 set out a number of alternatives, mostly listing alternative RDBMS solutions.[1601]

6.133 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 types of DBMS) as alternatives for SQL Server, considering their organisational use case for SQL Server.[1602]

6.134 Some customers said that there would be no substitute for their use of SQL Server.[1603]

[1597] Microsoft's response to the CMA's information request [✂].
[1598] Oracle's response to the CMA's information request [✂].
[1599] 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, an 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 [✂].
[1600] IBM's response to the CMA's information request [✂].
[1601] The most frequently listed alternatives were 'Oracle' or Oracle Database, MySQL and PostgreSQL, which are RDBMS solutions.
[1602] Some customers did not answer this question in their response and instead listed the different types of DBMS solutions their organisation uses.
[1603] Responses to the CMA's information requests [✂].

6.135 **RDBMS:** Most customers reported that they would consider other RDBMS as an alternative to SQL Server.[1604] Some of these customers highlighted that use of alternatives may depend on use case,[1605] business or application requirements,[1606] or that there may be considerable inertia when moving away from SQL Server.[1607]

6.136 **NRDBMS:** Most customers responded that they would not consider NRDBMS as alternatives to SQL Server.[1608] One customer highlighted that NRDBMS tend to be more specialised in how they work, and the use cases for which they are suitable.[1609] One customer outlined that if the workload required a RDBMS, then it would not use an NDRBMS, however explained that if the workload did not, it would consider alternatives.[1610] Another customer nuanced its answer explaining that NRDBMS would typically not be a good choice to replace SQL Server due to the need to refactor the application and its lack of technical suitability in most use cases where an existing RDBMS database has been deployed.[1611]

6.137 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.[1612] One customer further explained that, although it would consider alternatives, that it would be difficult to replace a RDBMS with a NRDBMS.[1613]

6.138 **Data analytics services:** Some customers responded that they did not consider data analytics services as alternatives to their use of SQL Server.[1614] One customer highlighted that even though SQL Server has analytics capabilities this is not its primary use case and that they use further analytics tools.[1615] Another customer explained that its use of SQL Server as a data analytics tool is minor compared to its use as an RDBMS.[1616]

6.139 A few customers said that they did consider data analytics services as alternatives to their use of SQL Server.[1617] In these instances customers explained that, as

[1604] This included one customer who only uses the product on non-public cloud. Responses to the CMA's information requests [✂].
[1605] [✂] response to the CMA's information request [✂].
[1606] Responses to the CMA's information requests [✂].
[1607] [✂] response to the CMA's information request [✂].
[1608] Responses to the CMA's information requests [✂].
[1609] [✂] response to the CMA's information request [✂].
[1610] [✂] response to the CMA's information request [✂].
[1611] [✂] response to the CMA's information request [✂].
[1612] Responses to the CMA's information requests [✂].
[1613] [✂] response to the CMA's information request [✂].
[1614] Responses to the CMA's information requests [✂].
[1615] [✂] response to the CMA's information request [✂].
[1616] [✂] response to the CMA's information request [✂].
[1617] Responses to the CMA's information requests [✂].

above, their use of alternatives would depend on the use case or workload[1618], or that it would be difficult to replace a RDBMS with a data analytics service.[1619]

6.140 **Other types of RDBMS:** Most customers responded 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.[1620] One customer mentioned that any transition from SQL Server would be to an alternative RDBMS system available within the organisation and not to a new technology.[1621]

6.141 A few customers responded that they did consider data analytics services as alternatives to their use of SQL Server.[1622] One customer explained that, as above, its use of alternatives would depend on the use case.[1623] Another customer said it was looking into Graph Database Technologies[1624] as an alternative to structured/relational databases.[1625]

6.142 Some customers detailed their use of different database management systems in addition to their use of SQL Server. The use of multiple solutions across a range of customers suggests that customers view different database solutions as somewhat suited for specific use cases and/or applications, and therefore may be complements rather than substitutes in some circumstances.

Provisional views on the product market for SQL Server

6.143 Our provisional view is that alternative DBMS (which are not RDBMS) are not effective demand-side substitutes for RDBMS. This is because most customers would not consider other DBMS solutions (outside RDBMS) 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

6.144 The evidence illustrates that, although there are several alternative suppliers for RDBMS, Microsoft has the largest share considering RDBMS in 2023, and

---

1618 Responses to the CMA's information requests [✂].
1619 [✂] response to the CMA's information request [✂].
1620 Responses to the CMA's information requests [✂].
1621 [✂] response to the CMA's information request [✂].
1622 This included one customer who only uses the product on non-public cloud. Responses to the CMA's information requests [✂].
1623 [✂] response to the CMA's information request [✂].
1624 A graph database is a specialized, single-purpose platform used to create and manipulate data of an associative and contextual nature. The graph itself contains nodes, edges, and properties that come together to allow users to represent and store data in a way that relational databases aren't equipped to do (What is A Graph Database? A Beginner's Guide, accessed 18 November 2024).
1625 This customer only uses the product on non-public cloud. [✂] response to the CMA's information request [✂].

considering all measures its share has remained at between [20-30]% and [30-40]% over the period 2019-2023. [1626]

6.145 We note that for 2020-2021 Oracle had the largest share in RDBMS. Shares have remained relatively stable over time, and Microsoft has slowly gained share over the past few years.

Product characteristics

6.146 RDBMSs provide a system to store and retrieve data stored in separate tables.

6.147 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.[1627] This may in turn depend on the customer's workload or specific requirements.

Customer submissions

6.148 Evidence from customers showed:

(a) Most customers we asked 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) In terms of alternatives to SQL Server, customers generally listed a variety of alternative products including Databricks, Oracle RDBMS, MySQL and others.

(c) 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. [1628]

[1626] As outlined in Appendix D.
[1627] For example, see Oracle vs. SQL Server: Head-to-Head Comparison, accessed 18 November 2024, Comparing Database Management Systems: MySQL, PostgreSQL, MSSQL Server, MongoDB, Elasticsearch, and others, accessed 18 November 2024.
[1628] Further customer evidence is presented in Appendix R.

6.149 The evidence from customers shows that few customers were willing and able to switch away from SQL Server and that even though there are alternatives available, most customers who use SQL Server on the public cloud would not switch to these.

Provider submissions

6.150 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 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.[1629]

Provisional conclusions on SQL Server

6.151 We have provisionally found that Microsoft has a moderate but significant market share and, although a number of alternatives exist in the market for RDBMS, customers are generally unwilling to switch to alternative products in response to a small but significant price increase.

6.152 While the evidence is mixed, considering the evidence in the round our provisional 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.

6.153 This view would not be different, even if we had defined a wider or narrower market. We conducted this assessment on a market for RDBMS. If the market had been widened to include alternative DBMS, we would come to the same conclusion with respect to Microsoft's market power because customers did not view these additional alternatives as substitutes for their use of SQL Server.

[1629] Further provider evidence is presented in Appendix R.

Customer evidence suggests customers are unwilling to switch away from SQL Server regardless of the frame of reference.

Microsoft Windows 10/11

6.154 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

6.155 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).[1630] Microsoft also said mobile OSs could be seen as a substitute for desktop OSs, eg by a developer of a web browser because web browsing can be done on both types of OSs.[1631,1632]

6.156 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.[1633]

Customer submissions

6.157 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 identified other desktop OSs (Linux, MacOS and ChromeOS).[1634] MacOS was the most frequently mentioned alternative. Some customers said there were no alternatives.[1635]

6.158 No customers suggested that a server OS would be an alternative. One customer suggested Android – we understand this was for a specific use case for a segment of staff. This suggests that customers do not generally view server and mobile OSs as substitutes for desktop OSs on the demand side.

[1630] Microsoft's response to the CMA's information request [✂].
[1631] Microsoft's response to the CMA's information request [✂].
[1632] We note that Microsoft viewed its PC OS as distinct from its other OSs in the European Commission Digital Markets Act designation decision.
[1633] IBM's response to the CMA's information request [✂].
[1634] This included a few customers who only uses the product on non-public cloud. Responses to the CMA's information requests [✂].
[1635] This included a customer who only uses the product on non-public cloud. Responses to the CMA's information requests [✂].

6.159 Customer responses indicated that customers would not be able to switch from a desktop OS to a server or mobile OS easily, suggesting there is limited demand side substitution from desktop OSs to server or mobile OSs. In addition, the functionality and intended use for each of these types of OSs is very different to a desktop OS. For example. Microsoft said desktop OSs are used to manage the hardware of a PC device and allows applications to run on it, and server OSs manage the hardware of a server device and allows applications to run on it.[1636]

Provisional view on the product market for Windows 10/11

6.160 Our provisional 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.

Market power

Market shares

6.161 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

6.162 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.[1637]

(b) A Microsoft internal document described points of difference between Windows 11 and MacOS, such as [✂].[1638]

[1636] Microsoft's response to the CMA's information request [✂].
[1637] Lakka, S et al. (2012), 'Competitive dynamics in the operating systems market: Modelling and policy implications', Technological Forecasting & Social Change, pp88-105.
[1638] Microsoft's response to the CMA's information request [✂].

(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.[1639]

(d) A Microsoft internal document described points of difference between Windows 11 and ChromeOS, such as [✂]. It did not compare Windows 11 to any Linux distributions.[1640]

(e) Some customers told us using MacOS or ChromeOS would require specific types of hardware.[1641]

(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 proprietary software.[1642] 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.[1643]

(g) A Microsoft internal document did not compare Windows 11 to any Linux distributions.[1644] One interpretation of this could be that Microsoft considers Windows 11 to be substantially differentiated from them.

6.163 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.

6.164 We asked customers to explain the reasons they chose Windows 10/11 rather than the alternatives they listed. Some customers responded that they value the additional functionality of Windows 10/11. Alternatives that lack this functionality may therefore be less substitutable

(a) Many customers 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.[1645]

1639 Miller, MR (2019), My Google Chromebook, 4th Edition, Que Publishing, page 12.
1640 Microsoft's response to the CMA's information request [✂].
1641 Responses to the CMA's information requests [✂].
1642 Dalheimer, MK and Welsh, M (2005), Running Linux, 5th Edition, O'Reilly.
1643 [✂] response to the CMA's information request [✂].
1644 Microsoft's response to the CMA's information request [✂].
1645 This included a few customers who use the product on the non-public cloud. Responses to the CMA's information requests [✂].

(b) Several customers said a reason was usability/ user familiarity.[1646]

(c) One customer said a reason was using the virtualised and hosted desktop capability,[1647] one said the features suit its intended use,[1648] and another said a reason was security capabilities.[1649]

6.165 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 submissions

6.166 Evidence from customers showed that:

(a) Most customers 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 listed Linux and MacOS, and some said there were no alternatives.

(c) Reasons 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. [1650]

Provider submissions

6.167 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

---

1646 This included a few customers who use the product on the non-public cloud. Responses to the CMA's information requests [✂].
1647 [✂] response to the CMA's information request [✂].
1648 [✂] response to the CMA's information request [✂].
1649 This customer uses the product on the non-public cloud. [✂] response to the CMA's information request [✂].
1650 Further customer evidence is presented in Appendix R.

desktop OS was whether there were a significant number of third party applications available to run on it.[1651]

Provisional conclusions on Windows 10/11

6.168 We have provisionally found that 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.

6.169 This view would not be different, even if we had defined a wider or narrower market. We conducted this assessment on a market for desktop OS. There was scope for Windows 10/11 to be in its own market because there were some limits to substitution between Windows 10/11 and other desktop OS products, and we took these into account in our assessment. Microsoft would hold the entire share of a market defined on this basis. If the market had been widened to include server or mobile OSs this would not make a difference to our view on market power, because they serve different use cases and customers did not view them as substitutes (which contrasts to Microsoft which provided some examples of use cases where they could be substitutes). In addition, customer evidence suggests customers would be unable or unwilling to switch away regardless of the frame of reference.

Microsoft's productivity suites

6.170 Microsoft has various packages of products which provide some productivity functionality. For the purposes of this investigation, we consider solutions only for enterprise consumers.

6.171 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.

6.172 For the purposes of this investigation, we have treated productivity suites[1652] for enterprise as the narrowest product market which the Microsoft productivity products sit within.

6.173 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'

[1651] Further evidence is presented in Appendix R.
[1652] 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.

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 submissions

6.174 Almost all of the customers we contacted had some usage of Microsoft 365 for Enterprise.[1653] A few customers reported use of Office.[1654]

6.175 We asked customers for alternatives to the Microsoft products that they could use for the same purpose.[1655] Customers almost exclusively listed only alternative productivity suites, rather than individual products.

(a) Almost all customers we contacted responded listing Google Workspace as an alternative which they could use for the same purpose as Microsoft 365.[1656] Most of these customers listed no other alternatives for Microsoft 365.[1657]

(b) A few customers listed a component of Google Workspace (Google Docs) as an alternative.[1658] In addition, one customer mentioned Microsoft Office on-premises (desktop installed apps) as an alternative.[1659] Some customers also listed open-source productivity suites as alternatives to Microsoft 365.[1660]

(c) One customer also listed substitutes for component elements of Microsoft 365 including substitutes for security services and eDiscovery services.[1661]

(d) One customer also listed an alternative productivity application (Click Up).[1662,1663]

1653 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 [✂].
1654 This included one customer who only uses the product on non-public cloud. Responses to the CMA's information requests [✂].
1655 CMA's RFI issued 06 December 2023, question 4.
1656 This included a few customers who only use the product on non-public cloud. Responses to the CMA's information requests [✂].
1657 This included a few customers who only use the product on non-public cloud. Responses to the CMA's information requests [✂].
1658 This included a few customers who only use the product on non-public cloud. Responses to the CMA's information requests [✂].
1659 [✂] response to the CMA's information request [✂].
1660 This included a few customers who only use the product on non-public cloud. Responses to the CMA's information requests [✂].
1661 [✂] response to the CMA's information request [✂].
1662 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.
1663 ClickUp is an application which combines multiple functionalities such as documents, project management and communication into one application. [✂] response to the CMA's information request [✂].

(e) No customers listed a complete mix and match solution which included individual applications which would fully cover Microsoft 365 functionality.

6.176 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.

Provisional view on the product market for productivity suites

6.177 Our provisional view is that 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

6.178 We have considered three measures of share of supply for productivity suites.[1664] Considering measures which present Microsoft's share, for the time periods available, Microsoft has held an [80-90]% share of supply in each year in the global market for productivity suites.

Product characteristics

6.179 We considered evidence from customers and software providers to understand the product characteristics of Microsoft 365.

6.180 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.

6.181 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.

6.182 We asked customers the reasons for choosing Microsoft 365 over alternatives they had listed.[1665] Some customers provided reasons that relate to functionality or

[1664] See Appendix R.
[1665] CMA's RFI issued 06 December 2023, question 4.

capabilities of Microsoft 365. Alternatives that lack this functionality may therefore be less substitutable.

(a) Many customers mentioned they value the large variety of Microsoft software applications/functionality.[1666]

(b) Some customers mentioned they value collaboration functionality (for example Microsoft Teams).[1667]

(c) Some customers mentioned valuing security capabilities.[1668]

(d) One customer mentioned directly valuing device management capabilities,[1669] however several more mentioned valuing integration, or having to re-integrate systems as a reason for choosing Microsoft 365.[1670]

6.183 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 submissions

6.184 Evidence received from customers[1671] 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

6.185 Key points from submissions received from providers are summarised below:

(a) [✂].

1666 This included a customer who only use the product on non-public cloud. Responses to the CMA's information requests [✂].
1667 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 18 November 2024. This included a customer who only uses the product on non-public cloud. Responses to the CMA's information requests [✂].
1668 This included a customer who only uses the product on non-public cloud. Responses to the CMA's information requests [✂].
1669 This customer only uses the product on non-public cloud. [✂] response to the CMA's information request [✂].
1670 This included a few customers who only use the product on non-public cloud. Responses to the CMA's information requests [✂].
1671 Further customer evidence is presented in Appendix R.

(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.[1672]

(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.[1673]

Provisional conclusions on productivity suites

6.186 We have provisionally found that Microsoft has a significant degree of market power in relation to its productivity suites. This is because 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.

6.187 This conclusion would not be different, even if we had defined a wider or narrower market. We conducted this assessment on a market for productivity suites for enterprise. We considered the competitive constraint exerted by alternatives to individual applications within productivity suites in our assessment. There was scope to consider a narrower market for only Microsoft packages because customer evidence suggests that customers are unwilling or unable to switch away from these. Microsoft would hold the entire share of a market defined on this basis. If the market had been widened to a market for productivity software, this would not make a difference to our view on market power because customers did not view them as substitutes. In addition, customer evidence suggests customers would be unable or unwilling to switch away regardless of frame of reference.

Microsoft Visual Studio

6.188 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.

[1672] Further provider evidence is presented in Appendix R.
[1673] Responses to the CMA's information requests [✂].

Product market definition

Provider and stakeholder submissions

6.189 Microsoft explained that Visual Studio can be used for building applications to run on non-Windows environments (in addition to Windows environments).[1674]

6.190 In response to our Licensing working paper, only one stakeholder commented on this market definition and agreed with our view.[1675]

Customer submissions

6.191 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.[1676] 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.[1677] Some of these are specialised eg Xcode is specialised in Apple development

(c) Other customers listed examples of software that are slightly different to IDEs eg VIM (a lightweight text editor) and GitHub (a developer platform).[1678]

(d) A few customers said there were no alternatives.[1679]

6.192 Overall, this suggests there could be a wide range of alternative products to Visual Studio.

6.193 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

[1674] Microsoft's response to the CMA's information request [✂].
[1675] 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.
[1676] This included one customer who only use the product on non-public cloud. Responses to the CMA's information requests [✂].
[1677] Responses to the CMA's information requests [✂].
[1678] This included a few customers who only use the product on non-public cloud. Responses to the CMA's information requests [✂].
[1679] This included one customer who only used the product on non-public cloud. Responses to the CMA's information requests [✂].

extent they would consider an IDE tailored to other types of software development, eg Java, to be a substitute for Visual Studio.

(a) Most customers said they wouldn't, or would be unlikely to, consider another IDE to be a substitute for Visual Studio.[1680] 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) A few customers said they would consider other IDEs to be a substitute for Visual Studio, for example Eclipse and IntelliJ were mentioned as alternatives.[1681] One customer said Visual Studio is still required for some purposes.[1682]

(c) A few other customers didn't express a strong opinion, explaining they tend to use the IDE most suited to each task or let the developer choose their preferred tool.[1683]

6.194 This suggests that customers have mixed views on whether an IDE tailored for non-Windows development would be a good substitute for Visual Studio.

6.195 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.

(a) Some customers said this was possible:[1684] some spoke positively of it; one said it was common;[1685] one said that it gives them access to more solutions;[1686] one was less positive, saying there would be fewer features.[1687]

(b) Some indicated it may be possible to an extent, but they don't do this;[1688] one of these explained that the efficiency of the developer would be reduced.[1689] Another customer said they cannot use another IDE for development of Windows applications built in C#/VB.net which we understand are

[1680] This included a few customers who only use the product on non-public cloud. Responses to the CMA's information requests [✂].
[1681] This included a few customers who only use the product on non-public cloud. Responses to the CMA's information requests [✂].
[1682] This customer only uses the product on non-public cloud. [✂] response to the CMA's information request [✂].
[1683] Responses to the CMA's information requests [✂].
[1684] This included one customer who only uses the product on non-public cloud. Responses to the CMA's information requests [✂].
[1685] [✂] response to the CMA's information request [✂].
[1686] [✂] response to the CMA's information request [✂].
[1687] [✂] response to the CMA's information request [✂].
[1688] This included a few customers who only use the product on non-public cloud. Responses to the CMA's information requests [✂].
[1689] [✂] response to the CMA's information request [✂].

programming languages developed by Microsoft and closely associated with Microsoft technology.[1690, 1691]

6.196 This suggests that other IDEs can be used for Windows development, but Visual Studio may be the most well suited for this.

6.197 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.

Provisional views on the product market for Visual Studio

6.198 Our provisional view is that the relevant product market is the market for IDEs.

Market power

Market shares

6.199 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 related[1692] 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 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

6.200 We considered evidence 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

[1690] This customer only uses the product on non-public cloud. [✂] response to the CMA's information request [✂].
[1691] 'The history and legacy of Visual Basic - Ryan Lucas', accessed 18 November 2024.
[1692] We explain the caveats associated with each of these in the Appendix R.

perform the same basic functionality, enabling software developers to build applications, web pages or services.

6.201 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.

6.202 We asked customers to explain the reasons they chose Visual Studio rather than the alternatives they listed. Many customers responded that they value the additional functionality of Visual Studio.[1693] 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),[1694] and integrations with other Microsoft products.[1695]

6.203 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 submissions

6.204 Evidence received from customers showed (a detailed discussion of the customer evidence is presented in Appendix R):

(a) Most customers 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 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) Most customers said there would be significant or some switching costs when switching from Visual Studio to an alternative IDE including: retraining,

[1693] This included a few customers who only use the product on non-public cloud. Responses to the CMA's information requests [✂].
[1694] This included one customer who only uses the product on non-public cloud. Responses to the CMA's information requests [✂].
[1695] This included one customer who only uses the product on non-public cloud. Responses to the CMA's information requests [✂].

impact on developer efficiency, staff recruitment issues and others. Some customers said there would be no or low switching costs.

Provider submissions

6.205 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. [1696]

Provisional conclusions on Visual Studio

6.206 We have found that, 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).

6.207 We have provisionally found that 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.

6.208 This provisional conclusion would not be different, even if we had defined the market more narrowly. We conducted this assessment on a market for IDEs. There was scope for Visual Studio to be in a market for IDEs for Windows development because this use case created some limits to substitution between Visual Studio and other IDEs, and we took these into account in our assessment. Microsoft would have a higher market share on this basis than we have found for a market for IDEs. In addition, customer evidence suggests customers would be unable or unwilling to switch away regardless of frame of reference.

[1696] Further provider evidence is presented in Appendix R.

Assessment of any cumulative effect of Microsoft's market power across markets

6.209 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.

6.210 This section builds on those provisional 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).

6.211 We asked customers whether, if at all, the way Microsoft software products are i) sold, for example, some products may be included together in packages or suites, or ii) purchased, for example, under framework agreements with Microsoft or other third parties, influenced their consumption decision around the Microsoft software products at the point of purchase.

6.212 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 said their purchasing decisions were impacted by the way the Microsoft products are sold or purchased.[1697] 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 Server, SQL Server and Visual Studio were not impacted because capability is more discrete.[1698]

(b) Some customers said that their purchasing decisions were not impacted but did not explain why.[1699]

(c) Others said their purchasing decisions were not impacted because they do not have an Enterprise Agreements.[1700]

6.213 We asked a similar question regarding whether the same factors act as barriers to switching away from the Microsoft products.

6.214 Most customers said these factors could act as a barrier to switching,[1701] however most of these considered other factors to be more important.[1702]

1697 Responses to the CMA's information requests [✂].
1698 [✂] response to the CMA's information request [✂].
1699 Responses to the CMA's information requests [✂].
1700 Responses to the CMA's information requests [✂].
1701 Responses to the CMA's information requests [✂].
1702 Responses to the CMA's information requests [✂].

6.215 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 said there were technical benefits to using Microsoft products together, for example customers expected a high level of interoperability, [1703] skills and staff training can be transferred between tools,[1704] and industry-wide adoption of Microsoft makes document sharing easier.[1705]

(b) Most customers thought there were technical limitations when using one of Microsoft's software products without one another (or with alternatives).[1706] Examples of limitations included extra work and infrastructure being required to integrate non-Microsoft software,[1707] functionality of some Microsoft products may be reduced when using non-Microsoft products with it (eg using Zoom with Outlook),[1708] or using M365 without Windows 10/11,[1709] security vulnerabilities and ineffective collaboration.[1710]

(c) Some customers said there were no or minimal technical limitations to using Microsoft products without one another or with alternatives.[1711] For example, one customer said it expected most non-Microsoft software would integrate with Active Directory/ Entra ID.[1712]

Provider submissions

6.216 In response to the Licensing working paper, some parties we contacted commented on the cumulative effect of market power across markets.

6.217 An academic said these cumulative effects can have benefits for consumers: 'The CMA is thus claiming that complementarities in consumption - or what one could

[1703] Responses to the CMA's information requests [✂].
[1704] Responses to the CMA's information requests [✂].
[1705] [✂] response to the CMA's information request [✂].
[1706] Responses to the CMA's information requests [✂].
[1707] Responses to the CMA's information requests [✂].
[1708] [✂] response to the CMA's information request [✂].
[1709] [✂] response to the CMA's information request [✂].
[1710] [✂] response to the CMA's information request [✂].
[1711] Responses to the CMA's information requests [✂].
[1712] [✂] response to the CMA's information request [✂].

call economies of scope in consumption - which are clearly benefits for consumers - may give rise to only greater CMA concern'.[1713]

6.218 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.[1714]

6.219 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).[1715]

6.220 The SMF said that there may be an issue with the way Microsoft sells its products: '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.'[1716]

Enterprise Agreements

6.221 In this section we consider additional evidence relating to Enterprise Agreements, and whether this affects customer choice of Microsoft software products. Customers can choose which products and at what volume these are included in the agreement.

Customer submissions

6.222 The Jigsaw report draws out various factors influencing customer purchasing decisions around choice of cloud, one of which is Enterprise Agreements. [1717] The report highlights that these agreements are perceived by some as good value compared to the cost of other software licences and services, however that some

[1713] 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 52.
[1714] Google's submission to the CMA [✂].
[1715] [✂] response to the CMA's information request [✂].
[1716] SMF response to the Licensing working paper, page 30.
[1717] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), page 29.

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.[1718]

6.223 In response to the working paper, SMF highlighted the role of contracts and agreements. It said: ‘Because its software is embedded across most UK offices, one participant 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’.[1719]

6.224 We gathered customer evidence to assess the importance Enterprise Agreements in customer purchasing decisions.

6.225 Many customers[1720] that responded to the CMA’s RFI had an Enterprise Agreement. Most of these lasted five years,[1721] some of these lasted three years.[1722] Most customers had at least two of the relevant Microsoft software products included as part of their Enterprise Agreement.[1723]

6.226 We asked customers about the benefits and disadvantages of Enterprise Agreements:

(a) Benefits: Customers that we contacted said that price certainty,[1724] flexibility[1725] and discounts compared to list prices[1726] were the key benefits of Enterprise Agreements.

(b) Disadvantages: Customers said that the disadvantages of Enterprise Agreements included restricting the ability to discontinue use or fall below minimum purchase requirements[1727] and a high commitment level over the term.[1728]

6.227 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.

[1718] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), page 29.
[1719] SMF response to the Licensing working paper, page 28.
[1720] Responses to the CMA’s information requests [✂].
[1721] Responses to the CMA’s information requests [✂].
[1722] Responses to the CMA’s information requests [✂].
[1723] Responses to the CMA’s information requests [✂].
[1724] Responses to the CMA’s information requests [✂].
[1725] Responses to the CMA’s information requests [✂].
[1726] Responses to the CMA’s information requests [✂].
[1727] Responses to the CMA’s information requests [✂].
[1728] Responses to the CMA’s information requests [✂].

(a) All customers responded that their Enterprise Agreement did not influence their consumption at the point of purchase or wasn't a major factor in their decision.[1729]

(b) Most customers outlined that their Enterprise Agreement would not act as a barrier to switching to alternative software products.[1730] 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.[1731] One customer outlined that its EA does act as a barrier to moving to alternative products.[1732]

6.228 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 the flexibility which EAs provide.

6.229 While these agreements may contribute to Microsoft's cumulative effect of market power across markets, the evidence suggests that these agreements are not a key source of market power, and that, absent these agreements customers would likely make the same purchasing decisions.

Provider submissions

6.230 In response to the Licensing working paper, Google submitted that Enterprise Agreements create sticky demand, saying: '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.'[1733]

Our assessment

6.231 We note that selling products together can have benefits for customers, such as reducing search costs, and that technical quality and integration can be a parameter of competition. 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,

[1729] Responses to the CMA's information requests [✂].
[1730] Responses to the CMA's information requests [✂].
[1731] Responses to the CMA's information requests [✂].
[1732] [✂] response to the CMA's information requests [✂].
[1733] Google response to the licensing working paper, paragraph 23.

and that our evidence relating to market power with respect to individual products would already capture these additional cumulative effects.

6.232 Most customers 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 customers said there were technical benefits to using Microsoft products together. Most customers thought there were technical limitations when using one of the Microsoft's software products without one another.

6.233 With respect to Enterprise Agreements, we recognise there are perceived benefits to customers relating to discounts for products when purchased together, but that there may also be drawbacks such as commitment over a time period or to a specific set of products.

6.234 In addition to the cumulative effects due to links between the five Microsoft software products we have focussed on, and other products from Microsoft's range (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 EAs (which is not already accounted for in the individual market power assessments for each product).

Provisional view on cumulative effect of market power

6.235 Our provisional view is that Microsoft has significant market power in relation to each of the relevant software products and any cumulative effects not already accounted for in the individual assessments, would support our view on the individual products.

Provisional conclusions on Microsoft's market power in related software markets

6.236 We have provisionally found that Microsoft has a significant degree of market power in relation to Windows Server, SQL Server, Windows 10/11, Visual Studio and its productivity suites.

6.237 This is because at least some of the following factors apply to each of these products: customer evidence suggests that customers are unwilling or unable to switch away from the product in response to a 5% price rise; there are limited competitive alternatives; the product is differentiated; there are barriers to customers switching; and Microsoft has a moderate to high share of each of the relevant markets.

6.238 The evidence set out above shows whether Microsoft has the ability and incentive to foreclose rivals: we have provisionally found that customers are unlikely to switch away from the relevant Microsoft software products, which therefore is consistent with Microsoft having the ability and incentive to foreclose rivals.

The importance of Microsoft software inputs

6.239 In this section, we assess the importance of the Microsoft software inputs in shaping both the prices and quality of downstream cloud offerings. We structure our assessment as follows:

(a) First, 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.

(b) Second, we consider the significance of Microsoft software inputs to AWS and Google other than the proportion of the cost base they account for. This section focuses on Windows Server, SQL Server, Windows Desktop, Microsoft's productivity suites, and Visual Studio as all products are subject to non-price related licensing conduct. In particular, we consider:

(i) Our analysis of the distribution of cloud use on Azure by spend on each service;

(ii) Evidence of the importance of the Microsoft software products to VDI;

(iii) Evidence on the relative size of cloud-based VDI and its potential to grow in the future; and

(iv) Qualitative evidence on the importance of other non-price factors including Extended Security Updates (ESUs).

Significance in the cost base of rival providers

6.240 In the context of price-based foreclosure, our assessment of the importance of the inputs focuses on their significance in the cost base of rival providers. The greater the proportion of total costs that a particular input accounts for, the more likely it is that changes in the price of that input will affect changes in the price of (or else the margins earned on) the downstream product.

6.241 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.

6.242 The price-based conduct we have considered relates to the wholesale 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. Accordingly, as submitted by Microsoft, the relevant numerator for this analysis is the Microsoft software IP input costs.[1734]

6.243 Before setting out different analyses of the Microsoft software inputs' significance in the cost base, we make the following two observations that are of cross-cutting relevance to the different analyses:

(a) In some industries, when calculating the input cost as a proportion of rivals' costs of supplying their customers, we would use the unit cost of the downstream product as the denominator.[1735] 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, the relevant denominator in any single case depends on the composition of the bundle that a customer is purchasing. Any given input may therefore represent a larger proportion of the cost base for some customers, and a smaller proportion for others. We consider below a number of different indicators of the importance of the input to the cost base and discuss the issue of the relevant denominator alongside those estimates.

(b) Additionally, while we would ideally analyse the significance of the Microsoft software input relative to the cost base of the downstream product, both Microsoft's and our own analysis uses downstream revenues as the denominator as a proxy as only revenue data was available. In particular,

---

[1734] Microsoft's response to the Licensing working paper, paragraph 1.3.

[1735] For example, one cloud provider submitted that the Windows Server licence typically comprises around [40-50]% of the on-demand list price for [✂] 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. [✂] response to Ofcom's information request [✂].

Microsoft's analysis considers the input cost share of revenues from all cloud services. Our analysis considers the input cost share of (i) revenues from all cloud services, and (ii) revenues from Windows Server VMs only.

Microsoft's analysis of AWS' and Google's licensing costs

6.244 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.[1736]

6.245 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.

6.246 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'.[1737]

6.247 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.

6.248 In addition, 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.

[1736] 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 the appendix for a discussion of the SQL Server Editions and their role in determining estimates of the significance of the input.
[1737] 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.

6.249 Our own analysis (set out below) addresses three of these issues and suggests that the proportion of revenues that the licensing input accounts for is higher than that submitted by Microsoft.

6.250 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. While a measure of cost that incorporates all cloud services might be appropriate if customers made a single indivisible choice about all of their cloud spend, we have seen evidence that customer decision-making happens at a less aggregated level.

6.251 For example, we have set out in chapter 3 evidence received from cloud providers suggests that most customers multi-cloud to some degree; this also covers evidence from cloud providers that suggests that digital native customers, larger enterprises, and enterprises with complex regulatory requirements are more likely to multi-cloud.

6.252 Additionally, our data analysis suggests that the prevalence of multi-cloud tends to increase with customers' total spend on cloud services, with 50% of customers in the highest revenue bracket using more than one cloud (see Appendix I). Customers with a higher propensity to multi-cloud are less likely to decide which cloud to deploy all their workloads in a single decision.

6.253 Even customers that do not multi-cloud may choose to deploy some workloads on one cloud and later add others in an independent decision. Microsoft has submitted that it competes for each workload.[1738] A cloud provider submitted that it discounts at the deal level, ie across the bundle of services that a customer includes in a single purchasing decision.[1739]

6.254 Where customers make choices about a subset of their cloud spend, it would be more appropriate to assess the significance of the licensing input costs relative to the cost of providing smaller subsets of services. By taking the widest possible denominator that includes all services, Microsoft's analysis risks understating the significance of licensing spend in the cost base.

6.255 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 below.

[1738] Microsoft's response to the Technical barriers working paper, paragraph 23.
[1739] [✂] submission to the CMA [✂].

6.256 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.

6.257 Fourth, Microsoft's analysis does not acknowledge that many customers use both Windows Server and SQL Server:

(c) Out of all Azure customers that use Windows Server, [10-20]% also use SQL Server on a PAYG basis.

(d) Out of all Azure customers that use SQL Server on a PAYG basis, [90-100]% also use Windows Server. [1740]

(e) Customers that us both Windows Server and SQL Server PAYG account for [70-80]% of total Azure revenues.

6.258 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.

6.259 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. In this respect, we make a number of observations:

(a) Analysis submitted by a cloud provider 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.[1741]

(b) Additionally, Microsoft's CEO stated in July 2023 that 'it remains early when it comes to the long-term cloud opportunity',[1742] and told investors in early 2024

[1740] 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).
[1741] [✂] response to the CMA's information request [✂].
[1742] Microsoft's Annual Report 2023.

that 'Azure again took share as organisations bring their workloads to the cloud'.[1743]

(c) As set out above, Windows Server had a [70-80]% share of supply in the Server OS market deployed in non-cloud environments in 2022.[1744] One provider submitted that traditional enterprise customers are more likely to have a historical reliance on Microsoft software based on Microsoft's historical dominance.[1745]

6.260 In light of this evidence, the significance of Microsoft software inputs may be reflected not only by their current significance in the cloud cost base, but also their current significance in on-premises use, which is an indicator of potential significance when competing for customers migrating to the cloud. To the extent that on-premises usage of the relevant Microsoft software is more significant than current usage levels on cloud, analysis based on current costs may understate the significance of those inputs for competition for cloud services.

6.261 Our own analysis, which we set out below, addresses three of the issues set out above.

Our analysis of SPLA costs

6.262 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 licensing spend would account for if Microsoft paid the same wholesale prices that it charges AWS or Google for Windows Server and SQL Server.

6.263 We consider this analysis helps to account for some of the issues relating to Microsoft's analysis above.

6.264 First, our analysis is not restricted to looking at licensing costs as a proportion of spend on all cloud services. Customers with a higher propensity to multi-cloud are less likely to decide which cloud to deploy all their workloads on in one go. Even customers that do not multi-cloud may choose to deploy some workloads on one cloud and later add others in a completely unrelated decision. As set out above, customers may vary in terms of how aggregated their decisions on cloud services are. To account for this apparent variation in the level of aggregation at which customers make these decisions, we have compared the Windows Server licensing input costs to:

[1743] Microsoft's Q1 2024 earnings call.
[1744] 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. One party also submitted a bespoke analysis [✂], which suggests that Windows Server workloads account for 60-70% of all spending on on-premises workloads running in the UK. [✂] response to the CMA's information request [✂].
[1745] In its 2004 Microsoft (Windows Media Player) infringement decision, the European Commission found that Microsoft's share of the work group server OS market was 60-75%. [✂] response to the CMA's information request [✂].

(a) total customer spend on Azure services, and

(b) customer spend on Windows Server VMs on Azure.

6.265 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.

6.266 In each case, (a) and (b) represent narrow and broad purchasing decisions, respectively. By comparing the licensing input costs to these two denominators, our analysis provides an upper and lower bound estimate of the Microsoft software inputs' significance when competing for cloud customers purchasing different bundles of services.

6.267 To elaborate, Windows Server is used as an operating system on VMs. SQL Server is a software product that is installed on VMs, and predominantly Windows Server VMs. For example, [90-100]% of SQL Server usage on Azure occurred on Windows Server VMs in 2022. As such, a Windows Server VM is the narrowest set of downstream products that Windows Server can serve as an input for, whereas a Windows Server VM with SQL Server installed is the narrowest set of downstream products that SQL Server can serve as an input for. On the other hand, the full set of Azure services that a customer uses represents the broadest bundle that each of Windows Server and SQL Server can serve as an input for.

6.268 Based on the evidence outlined above, we consider that neither of these upper and lower bound estimates are likely to represent the significance of the licensing input costs in competing for any single customer. Rather, we consider that values within the ranges that we present are likely to be more reliable as indicators. For customer groups with a higher propensity to multi-cloud, such as higher spend customers, the significance is likely to be closer to the upper bound, whereas for customers with a lower propensity to multi-cloud, such as lower spend customers, the significance is likely to be closer to the lower bound.[1746]

6.269 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, 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 wholesale prices that it charges to AWS

[1746] Our data analysis suggests that the prevalence of multi-cloud tends to increase with customers' total spend on cloud services, with 50% of customers in the highest revenue bracket using more than one cloud (see Appendix I).

and Google. We then compared these hypothetical input costs to the spend denominators described above.

6.270 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.

6.271 However, our upper bound estimates of the significance of the input (ie those that compare the licensing costs to spend on Windows Server VMs and Windows Server VMs plus SQL Server) are less likely to be overestimates. This is because 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.

6.272 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).

6.273 Our analysis does not overcome the revenue denominator issue as we do not have access to data on Microsoft's cost base with respect to each customer. It also does not address the forward-looking issue around the potential for the importance of Microsoft software workloads to increase in the future. As such, it may still understate the significance of the licensing input in the cost base because of either of these factors.

6.274 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 below, SQL Database accounted for [0-10]% of total Azure revenues in 2022. 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).

6.275 As noted 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 wholesale prices that it charges AWS and Google. To do this, we multiplied each Azure customer's total 2022 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.[1747] We then took these hypothetical licensing input costs as a proportion of each customer's:

[1747] 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.

(a) total spend on Azure services, and

(b) spend on the narrowest set of downstream products that each software product can serve as an input for on Azure.

6.276 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.

6.277 There are three editions of SQL Server that are available to license via the SPLAs. Enterprise Edition is significantly more expensive than Standard Edition, and Standard Edition is significantly more expensive than Web Edition. The usage data for SQL Server does not break down usage into usage levels for each edition. Given the differences in the price of each edition, the proportion of usage of each edition that we assign to customers substantially impacts our estimates of the SQL Server licensing input costs.

6.278 Our analysis assumes the following three flat ratios of Enterprise to Standard edition: 50:50, 40:60, and 25:75. The evidence that we set out in Appendix T suggests that the 50:50 ratio may be more accurate for customers in the highest revenue brackets, whereas the 40:60 and 25:75 ratios may be more accurate for customers in the lower brackets.

6.279 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 heterogenous 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. As outlined in Appendix T, larger customers are more likely to use more expensive versions of SQL Server than smaller customers. This means that the mean SQL Server input cost proportions may be less representative of the majority of the customers within these brackets.[1748]

6.280 As explained in full in Appendix T, we have excluded customers with <$10k spend from this analysis. These customers account for a [✂] minority of total Azure spend. Therefore, our analysis still covers the [✂] majority of Azure spend.

6.281 Google's SPLA with Microsoft [✂]. We present the results using Google's prices below and present the full results in Appendix T.

6.282 In the figures and tables below, ACR stands for Azure Consumed Revenue and denotes total customer spend on Azure services; VM denotes spend on Windows

[1748] For Windows Server, the mean and median results are very similar.

Server VMs; and VM+SQL denotes spend on Windows Server VMs and SQL Server IP on Azure.

6.283 Figure 6.1 and Table 6.3 below shows Microsoft's Windows Server licensing input costs if it paid the same wholesale 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.

Figure 6.1: [✂]

Table 6.3: 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]	[50-60]
5M-10M	[5-10]	[50-60]
10M-20M	[5-10]	[60-70]
>20M	**[5-10]**	**[70-80]**

Source: CMA analysis of Microsoft's response to the CMA's information request [✂]; Microsoft's response to Ofcom's information request [✂].

6.284 As discussed above, we consider that the significance of the input is best represented by 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 [70-80]% of the relevant spend denominator for customers in each bracket.

6.285 Figure 6.2 and Table 6.4 below shows Microsoft's SQL Server licensing input costs if Microsoft paid Google's SPLA wholesale 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.

Figure 6.2: [✂]

Table 6.4: 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

	50:50 Enterprise to Standard		*40:60 Enterprise to Standard*		*25:75 Enterprise to Standard*	
Revenue bracket ($)	ACR (%)	VM (%)	ACR (%)	VM (%)	ACR (%)	VM (%)
10k-1M	[20-30]	[60-70]	[20-30]	[60-70]	[10-20]	[40-50]
1M-5M	[5-10]	[20-30]	[5-10]	[20-30]	[0-5]	[10-20]
5M-10M	[0-5]	[20-30]	[0-5]	[10-20]	[0-5]	[10-20]
10M-20M	[0-5]	[10-20]	[0-5]	[10-20]	[0-5]	[10-20]
>20M	**[0-5]**	**[10-20]**	**[0-5]**	**[10-20]**	**[0-5]**	**[10-20]**

Source: CMA analysis of Microsoft's response to the CMA's information request [✂]; Microsoft's response to Ofcom's information request [✂].

6.286 Our analysis suggests that SQL Server accounts for at least [0-5]% and as much as [20-30]% of the relevant spend denominator for across the five highest revenue brackets.

6.287 Figure 6.3 and Table 6.5 below shows Microsoft's SQL Server input costs if it paid Google's SPLA wholesale 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.

Figure 6.3: [✂]

Table 6.5: 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

	50:50 Enterprise to Standard		*40:60 Enterprise to Standard*		*25:75 Enterprise to Standard*	
Revenue bracket ($)	ACR (%)	VM (%)	ACR (%)	VM (%)	ACR (%)	VM (%)
10k-1M	[10-20]	[50-60]	[10-20]	[40-50]	[10-20]	[30-40]
1M-5M	[0-5]	[10-20]	[0-5]	[10-20]	[0-5]	[10-20]
5M-10M	[0-5]	[10-20]	[0-5]	[10-20]	[0-5]	[10-20]
10M-20M	[0-5]	[10-20]	[0-5]	[10-20]	[0-5]	[10-20]
>20M	**[0-5]**	**[20-30]**	**[0-5]**	**[10-20]**	**[0-5]**	**[10-20]**

Source: CMA analysis of Microsoft's response to the CMA's information request [✂]; Microsoft's response to Ofcom's information request [✂].

6.288 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.

6.289 Figure 6.4 and Table 6.6 below shows Microsoft's combined Windows Server and SQL Server licensing input costs if it paid Google's SPLA wholesale prices. These costs are expressed as a share of customer spend across all Azure services and spend on Windows Server VMs and SQL Server IP on Azure on average for customers in each revenue bracket that use both products.

6.290 As explained above, [✂]% of SQL Server PAYG customers also use Windows Server. Therefore, we consider that the significance of the licensing input to competition for SQL Server customers is best represented by the combined costs across both products.

Figure 6.4: [✂]

Table 6.6: 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

	50:50 Enterprise to Standard		*40:60 Enterprise to Standard*		*25:75 Enterprise to Standard*	

Revenue bracket ($)	ACR (%)	VM (%)	ACR (%)	VM (%)	ACR (%)	VM (%)
10k-1M	[30-40]	[100-200]	[30-40]	[100-200]	[30-40]	[90-100]
1M-5M	[10-20]	[70-80]	[10-20]	[70-80]	[10-20]	[60-70]
5M-10M	[10-20]	[70-80]	[10-20]	[70-80]	[10-20]	[70-80]
10M-20M	[10-20]	[70-80]	[10-20]	[70-80]	[10-20]	[70-80]
>20M	**[10-20]**	**[80-90]**	**[10-20]**	**[80-90]**	**[5-10]**	**[70-80]**

Source: CMA analysis of Microsoft's response to the CMA's information request [✂]; Microsoft's response to Ofcom's information request [✂].

6.291 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 [100-200]% of the relevant spend denominator across all revenue brackets.

6.292 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. 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.

6.293 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

6.294 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 [✂] Windows Server usage do not amount to material demand for the purposes of assessing whether AWS or Google are foreclosed.[1749]

6.295 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 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.[1750] Further, customers in

[1749] Microsoft's submission to the CMA. [✂].
[1750] Microsoft's submission to the CMA. [✂].

the [✂] revenue brackets [✂] spent [10-20]% and [5-10]% on Windows Server VMs, respectively.[1751]

6.296 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,[1752] which are more valuable for the retention or building of scale.[1753] 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.

6.297 As set out above, Microsoft also submitted that AWS and Google are not foreclosed in relation to customer sub-segments with [✂] Windows Server usage. In this regard, Microsoft considered three customer groups: the [✂] 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, 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.[1754,1755]

6.298 However, we note certain issues with this analysis and Microsoft's interpretation of the results, some of which we have also identified in relation to previous analyses:

(a) Microsoft's estimate of the overall average spend on Windows Server VM represents a lower bound estimate of 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) In addition, 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 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.

(c) We do not consider that [10-20]%, as a lower bound estimate, is an insignificant portion of total revenues;

1751 Microsoft's submission to the CMA. [✂].
1752 Microsoft's response to the Licensing working paper, paragraph 7.21.
1753 Microsoft's response to the Licensing working paper, paragraph 1.3, part d.
1754 Microsoft's submission to the CMA. [✂].
1755 Microsoft conducted this analysis on 2023 data.

(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.[1756]

6.299 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.

6.300 We consider that 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.

6.301 We agree that these customer groups represent a small portion of the market in comparison to other types 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.

6.302 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.

6.303 Using the data made available to us, we also 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.[1757]

6.304 Finally, we note that Microsoft 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

[1756] CMA analysis of Microsoft's response to the CMA's information request [✂].
[1757] CMA analysis of Microsoft's response to the CMA's information request [✂].

customers on average. Indeed, in our own analysis, we found that combined Windows Server VM spend 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

6.305 As outlined above, even in circumstances where an input constitutes a small part of the cost base in monetary terms, it may still be significant in terms of affecting the overall quality or attractiveness of suppliers' 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.

6.306 In this section, we set out evidence on:

(a) What customers use cloud for, according to the overall distribution of cloud spend including spend on services that make use of Microsoft software;

(b) The importance of Microsoft software to the quality of cloud-based VDI offerings;

(c) The current and future significance of VDI as a cloud workload; and

(d) Cross-cutting evidence on whether certain non-price factors impact on the quality of relevant cloud services.

Distribution of cloud use by spend

6.307 We have assessed the importance of the cloud services that make use of each software product relative to overall cloud spend. We have also further calculated the proportion of total revenues that spend on Microsoft software workloads accounts for. The greater the proportion of cloud spend that Microsoft software workloads account for, the greater scope there is for any price or non-price differences in the Microsoft input to impact downstream competition.

6.308 We used Azure UK revenue data for 2023. To the extent that Microsoft's licensing conduct causes customers to disproportionately deploy their Microsoft software workloads on Azure, this analysis overstates the importance of Microsoft software workloads relative to overall cloud revenues.

6.309 Figure 6.5 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 6.5: Distribution of Azure cloud use by spend on each service

[✂]

6.310 Our analysis suggests that:

(a) Virtual Machines is [✂] cloud service on Azure, accounting for [20-30]% of total UK revenues. Of that, spend on Windows Server and Windows Desktop VMs account for [40-50]% and [10-20]% respectively.

(b) Storage is the [✂] cloud service, accounting for [20-30]% of total Azure spend.

(c) SQL Database is the [✂] 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 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. [✂].[1758] Together, therefore, VDI revenues account for [0-5]% of total Azure revenues.

(g) All other services [✂] account for less-than [0-5]% of Azure UK revenues each.

6.311 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 [✂].

The importance of the Microsoft software products to VDI offerings

6.312 The primary use case for Microsoft's client-side products, ie Windows Desktop, 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.

6.313 We have provisionally found that Microsoft has significant market power with respect to all of these client-side products.

[1758] Note of meeting with [✂].

6.314 We have also seen evidence that shows that Microsoft software is important to the supply of cloud-based VDI:

(a) One cloud provider submitted that [60-70]% of all [✂] run Microsoft Office, and approximately [90-100]% of [✂] instances run Windows instead of Linux.[1759]

(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 [✂] VMs which showed [90-100]% are based on Windows.[1760]

Current importance of cloud-based VDI workloads

6.315 **Azure revenue analysis:** 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 [✂] on Azure. However, given the importance of the Microsoft software input to VDI (as evidenced just above) and the potential of the conduct to foreclose rivals (see below), we expect the proportion of market-wide revenues that VDI accounts for to be lower.

(a) As part of our initial evidence gathering, we asked some customers what proportion of their total cloud spend were 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%[1761] while the other half spent less than 10%.[1762]

(b) In order to gather further customer evidence on this question, we contacted additional customers some of which were active in industries which are more likely to use VDI. In response, most customers we contacted spent 5-10% of their total cloud spend on VDI.[1763] However some spent up to 20-30%,[1764] and some spent 0-5%.[1765]

6.316 **Internal documents**: We have also seen internal documents which detail the current size of the market for cloud based VDI.

[1759] [✂] submission to the CMA [✂].
[1760] [✂] response to the CMA's information request [✂].
[1761] Responses to the CMA's information requests [✂].
[1762] Responses to the CMA's information requests [✂].
[1763] Responses to the CMA's information requests [✂].
[1764] Responses to the CMA's information requests [✂].
[1765] Responses to the CMA's information requests [✂].

(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.[1766]

(b) Microsoft submitted an internal document which outlines that the global desktop as a service market was forecast to be worth [✂] USD in 2022.[1767]

(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 USD in 2023.[1768]

(d) 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 USD, and PaaS is reported as $142 billion USD.[1769]

6.317 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

6.318 We undertook an analysis of the relative usage from 2020-2023 of VMs that use Windows Server, Linux, and Windows Desktop as the OS on Azure.

6.319 As noted above, we understand that usage of Windows Desktop VMs represents the underlying compute resource used to run VDI instances on Azure. Therefore, to the extent that Windows Desktop 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.

6.320 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

[1766] Google's response to the CMA's information request [✂]. Cloud-Based VDI Market Scope & Dynamics Forecast 2032 (imarcgroup.com).
[1767] Microsoft's response to the CMA's information request [✂].
[1768] Microsoft's response to the CMA's information request [✂].
[1769] Microsoft's response to the CMA's information request [✂].

calculate shares of usage by each OS, as well as annual growth rates and CAGRs for each OS.

6.321 Figure 6.6 below shows shares of Azure VM usage by operating system over the Years 2020-2023.

Figure 6.6**: Shares of Azure VM usage by operating system 2020-2023**.

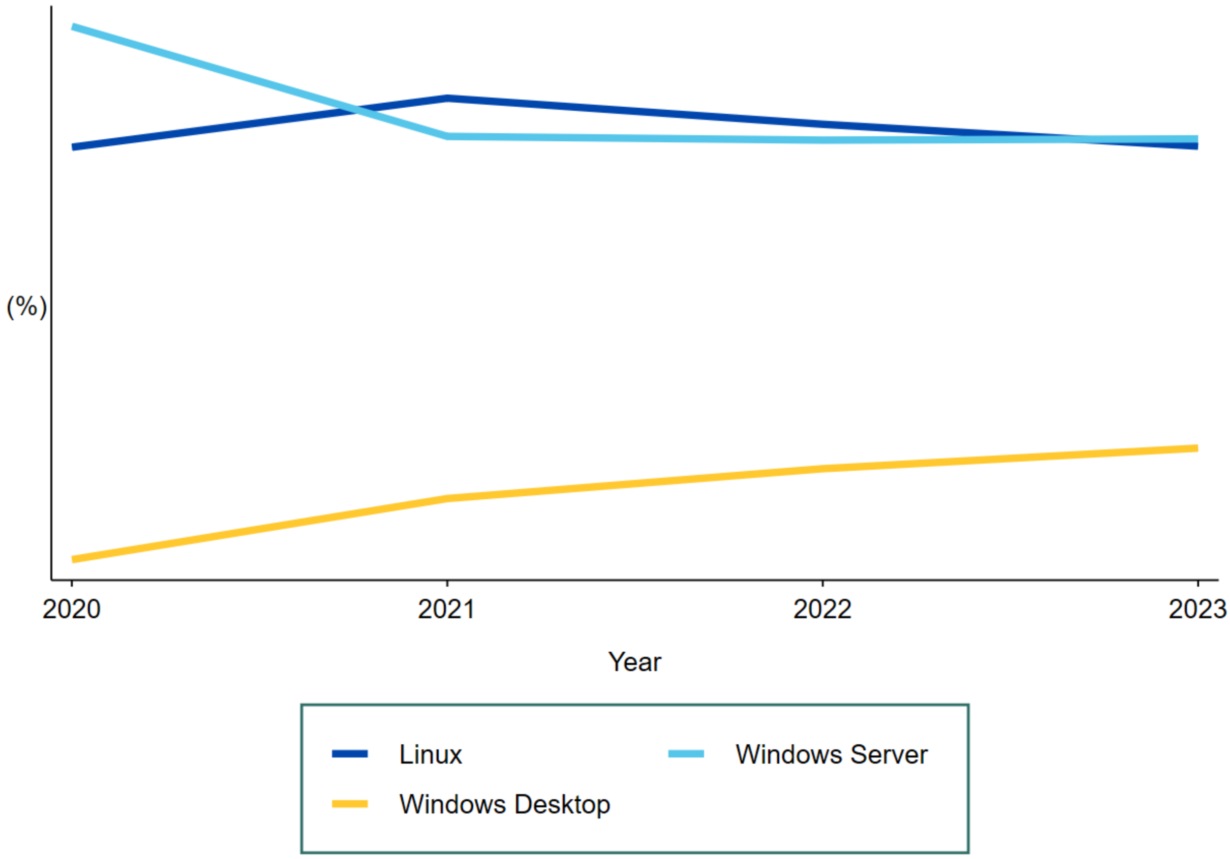

Source: CMA analysis of Microsoft's response to the CMA's information request [✂].

6.322 Our analysis suggests that the share of Windows Desktop VM usage has increased from [✂] in 2020 to [✂] in 2023. Since 2020, Windows Desktop VM usage has increased at a CAGR of [✂], compared with CAGRs for Windows Server and Linux of [✂] and [✂] respectively.[1770] Our analysis shows that VDI demand on Azure has increased at a CAGR of [✂]. 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.[1771]

[1770] CMA analysis of Microsoft's response to the CMA's information request [✂]. We note that Microsoft began offering VDI services in 2019.
[1771] CMA analysis of first-party revenue data and IDC and Synergy data. Refer to chapter on Market structure and concentration for further information.

6.323 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.[1772] As such, this analysis may understate the share of VM resource that is used to support VDI workloads on Azure.

6.324 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 Azure, we consider that this analysis provides a good indication of the general growth in demand for cloud-based VDI.

VDI provider views

6.325 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.[1773]

(b) A 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.[1774] 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.[1775]

(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.[1776]

1772 [✂] response to the CMA's information request [✂].
1773 Note of meeting with [✂].
1774 Note of meeting with [✂].
1775 Note of meeting with [✂].
1776 Note of meeting with [✂].

Cloud providers' views

6.326 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 5-10 years.[1777] 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.[1778] 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).[1779]

6.327 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.[1780]

(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. [1781] This growth is accelerated from [10-20]% CAGR (across all cloud deployment types) from 2017-2022.[1782]

[1777] 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 7 October 2024), 16.5% CAGR 2022-2032 (Source: Business research insights, Cloud Infrastructure Services Market Size, Share, Growth, and Industry Analysis, accessed 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 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 consider that despite this caveat, these estimates provide an indication of the future trajectory against which to compare VDI workloads.
[1778] Microsoft's response to the CMA's information request [✂].
[1779] Microsoft's response to the CMA's information request [✂].
[1780] Imrac Group, Cloud-Based VDI Market Scope & Dynamics, accessed 4 October 2024.
[1781] Microsoft's response to the CMA's information request [✂].
[1782] Microsoft's response to the CMA's information request [✂].

(c) The same Microsoft internal document dated January 2023 reported that [✂] of all enterprise PC users will rely on DaaS within three years.[1783]

Customer views

6.328 We asked customers, some of which were active in industries which are more likely to use VDI, 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,[1784] and a few customers expected VDI as a proportion of total cloud spend to increase.[1785]

Provisional views on the importance of Microsoft products to VDI offerings

6.329 We have provisionally found that VDI currently accounts for a small proportion of Azure UK public cloud revenues and likely a smaller proportion of market-wide revenues.

6.330 VDI may 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 at a faster rate than the broader cloud market. 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 cloud services.

6.331 Our provisional view is that, while there is mixed evidence as to whether VDI workloads will grow to be more important relative to other public cloud workloads over time, in combination with other Microsoft workloads, VDI workloads are a sufficiently important input into rivals' competitive offerings such that they contribute to Microsoft's ability to worsen rivals' competitive offerings at present.

Cross-cutting evidence on the importance of quality factors

6.332 As we set out above, in assessing the importance of the input, we consider not only the proportion of rivals' costs that the input accounts for, but also for example 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.

6.333 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

[1783] Microsoft's response to the CMA's information request [✂].
[1784] Responses to the CMA's information requests [✂].
[1785] Responses to the CMA's information requests [✂].

therefore experience any differences directly.[1786] 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.

6.334 We asked customers about the differences, if any, between using Microsoft software products[1787] 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 differences[1788]:

(a) Most customers did not identify differences in quality factors between public clouds;[1789]

(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;[1790] and

(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.[1791]

6.335 We asked customers about specific quality differences that Google, CISPE and/or AWS submitted 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 speak to price factors:[1792]

(a) Most customers we contacted felt that they were unable to answer the question because they had not compared these differences across clouds or

[1786] 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.
[1787] The question referred to Windows Server, SQL Server, Windows 10/11, Microsoft Office, Microsoft 365 and Active Directory/Azure AD specifically.
[1788] Where this is the case, AWS and Google are theoretically more likely to be able to compensate customers through offering price discounts.
[1789] Responses to the CMA's information requests [✂].
[1790] [✂] response to the CMA's information request [✂].
[1791] [✂] response to the CMA's information request [✂].
[1792] 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.

that any differences were not relevant because they have a small Microsoft estate on public cloud;[1793]

(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).[1794] For example, 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’.[1795]

(ii) Other factors that a few customers identified included the features and/or functionality available on software;[1796] minimum purchase requirement, eg duration of licences or number of virtual cores;[1797] and the basis by which you are charged for use of Microsoft products on public cloud.[1798]

6.336 Overall, the evidence suggests that most customers are currently not aware of quality differences in using Microsoft software products on Azure compared to rival clouds. As such, on the whole we do not consider that current quality differences indicate that the input is materially more important than set out above. 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, Microsoft software may be particularly important in terms of affecting the overall quality or attractiveness of AWS’ and Google’s competitive offerings.

Our assessment

6.337 We have considered the importance of the Microsoft software inputs, both in terms of their significance to the cost base of rival providers and their potential to shape downstream competition, for example by degrading the quality of rival providers’ offerings. 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.

6.338 Our view is that Microsoft’s analysis of AWS’ and Google’s SPLA costs understates the importance of the Microsoft software input.

[1793] Responses to the CMA’s information requests [✂].
[1794] Responses to the CMA’s information requests [✂].
[1795] [✂] response to the CMA’s information request [✂].
[1796] Responses to the CMA’s information requests [✂].
[1797] Responses to the CMA’s information requests [✂].
[1798] [✂] response to the CMA’s information request [✂].

6.339 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 [70-80]% 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 [100-200]% of the relevant spend denominator across all revenue brackets.

6.340 As our analysis compares the licensing input costs to revenues rather than costs, and it does not address the forward-looking issue around the potential for the importance of Microsoft software workloads to increase in the future, it may still understate the significance of the licensing input in the cost base because of either of these factors.

6.341 Based on the evidence outlined above, we consider that the upper and lower bound estimates in each revenue bracket are unlikely to represent the significance of each product in competing for any single customer. Rather, we consider that values within the range between them are likely to be more reliable as indicators. The significance in competing for larger customers is likely to be closer to the upper bound estimate (as such customers have a greater propensity to multi-cloud), whereas the significance in competing for smaller customers is likely to be closer to the lower bound estimate, on average.

6.342 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 the evidence shows that Windows Server is an important input on its own.

6.343 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.

6.344 Finally, we conducted a two-limbed assessment of the importance of the client-side Microsoft software products. Our assessment suggests that Windows Desktop 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.

6.345 We consider that Microsoft's client-side products do not constitute an important input on their own. However, we consider that the client-side products contribute to the overall importance of the Microsoft software inputs.

Provisional conclusions

6.346 We have provisionally found that the relevant Microsoft software products are sufficiently important as inputs to shape downstream competition. This is consistent with Microsoft having the ability to partially foreclose its rivals.

6.347 The evidence related to Windows Server is particularly strong and this software on its own is a sufficiently important input for rivals' clouds offerings such that it is consistent with Microsoft having the ability to partially foreclose its rivals.

Microsoft's conduct

6.348 Input foreclosure may arise from a firm supplying rivals with an input at either a higher cost or lower quality than it enjoys itself, or refusing to supply the input to rivals altogether.

6.349 In this section, we set out evidence on Microsoft's current conduct through considering 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 its downstream rivals. In doing so, we have considered the significance of any differences in price and/or quality in relation to customer decisions: the greater the differences in the input, the more likely it is that the conduct has an effect on competition.

6.350 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; and,

(c) the evidence related to the price and non-price differences for the client-side products Visual Studio, Microsoft 365, Office and Windows Desktop.

Our comparison of the price that Microsoft charges Azure customers for Windows Server and SQL Server and the price it charges AWS and Google

6.351 We have assessed whether, and if so the extent to which, Microsoft's conduct incurs a price disadvantage on AWS and Google in competing for public cloud customers. In particular, we compared the prices that Microsoft charges its own customers to use Windows Server and SQL Server with the 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. The higher the prices that Microsoft charges to AWS and Google compared with its own customers, the more difficult AWS and Google will find it to cover their other cloud costs while remaining competitive with Microsoft on price.

6.352 In order to assess the price disadvantage that AWS and Google could face compared 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.

6.353 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 wholesale prices that it charges AWS and Google.

6.354 As such, this analysis understates the 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.

6.355 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 wholesale prices that it charges AWS and Google. To do this, we multiplied each Azure customer's total 2022 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 hypothetical 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 wholesale prices that it charges AWS and Google.

6.356 As such, our analysis effectively compares Azure's customer-facing prices (net of discounts) and the wholesale 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.[1799]

6.357 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 additional input costs that AWS and Google would have to absorb to match Microsoft's competitive offering to customers that want to use each software product on public cloud (and cover their costs) all else equal.

6.358 For customers that make indivisible choices about all their cloud spend, the additional costs that AWS and Google must absorb to compete with Microsoft for customers that demand each product is best represented by the difference in prices as a proportion of total spend across all Azure services. Conversely, for customers that allocate their cloud spend on a workload-by-workload basis, the additional costs that AWS and Google must absorb to compete with Microsoft are best represented by the difference in prices as a proportion of spend on Windows Server VMs (and SQL Server IP in the case of SQL Server). In reality, customers are likely distributed between these two extremes. As such, the additional costs that AWS and Google must absorb to compete with Microsoft for any single customer that demands each product will fall within the range that we present for each revenue bracket.

6.359 However, given our understanding that larger customers are more likely to multi-cloud than smaller customers, we consider that the additional costs that AWS and Google must absorb to compete for larger customers will be closer to the upper bound than for smaller customers, on average.

6.360 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 conduct to foreclose AWS and Google may vary with respect to customers of different sizes.

6.361 We also present the median results to account for the fact that some revenue brackets encompass a broad range of heterogenous 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

[1799] 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.

Server). This means that the mean results may be less representative of the majority of the customers within these brackets.

6.362 The AHB allows customers with qualifying existing on-premises Windows Server and SQL Server licences and Software Assurance subscriptions to effectively migrate these licences onto Azure VMs without paying any additional licensing fees. Alternatively, customers that are not eligible for the AHB can purchase a licence-included VM and thereby pay for the Windows Server and/or SQL Server software IP on a PAYG basis. For Windows Server, due to the availability of data our analysis focuses on the difference in customer-facing and wholesale prices relating to AHB usage only. Appendix T sets out our full methodology including a discussion of the AHB IP spend data point.

6.363 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.[1800]

6.364 As discussed above, there are three editions of SQL Server that are available to license via AWS' and Google's SPLAs. The usage data for SQL Server does not break down into usage of each edition, and the proportion of usage that we assign to each edition substantially impacts our estimates of the SQL Server licensing input costs.

6.365 Our analysis assumes the following three flat ratios of Enterprise to Standard edition: 50:50, 40:60, and 25:75. Based on evidence set out in Appendix T, we consider that estimates based on the 50:50 ratio may be more accurate for customers in the highest revenue brackets, whereas estimates based on the 25:75 ratio may be more accurate for customers in the lower revenue brackets. Please see Appendix T for our full reasoning and methodology behind these assumptions and the analysis more generally.

6.366 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.

[1800] 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.

6.367 Google's SPLA with Microsoft includes [✂]. We present the results using Google's prices below and present the full results in Appendix T.

6.368 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.

6.369 In the figures presented below ACR stands for Azure Consumed Revenue and denotes total customer spend on Azure services, VM denotes spend on Windows Server VMs, and VM+SQL denotes spend on Windows Server VMs and SQL Server IP on Azure.

6.370 Figure 6.7 and Table 6.7 below shows the difference between the wholesale 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.

Figure 6.7: [✂]

Table 6.7: Average percentage difference between the wholesale prices that Google pays for Windows Server and Microsoft's customer-facing prices

Revenue bracket ($)	*Percentage difference (%)*
10k-1M	[4000-5000]
1M-5M	[3000-4000]
5M-10M	[1000-2000]
10M-20M	[1000-2000]
>20M	**[4000-5000]**

Source: CMA analysis of Microsoft's response to Ofcom's information request [✂].

6.371 Our analysis suggests that the wholesale prices that Microsoft charges Google for Windows Server are at least [1000-2000]% higher than its customer-facing prices relating to Windows Server AHB usage.

6.372 Figure 6.8 and Table 6.8 shows the median percentage difference between Microsoft's customer-facing prices and the wholesale prices that Google pays for Windows Server for customers in each revenue bracket.

Figure 6.8: [✂]

Table 6.8: Median percentage difference between the wholesale prices that Google pays for Windows Server and Microsoft's customer-facing prices

Revenue bracket ($)	*Percentage difference (%)*
10k-1M	[900-1000]
1M-5M	[1000-2000]
5M-10M	[1000-2000]
10M-20M	[1000-2000]
>20M	**[3000-4000]**

Source: CMA analysis of Microsoft's response to the CMA's information requests [✂]; Microsoft's response to Ofcom's information request [✂]. Median percentage difference between the wholesale prices that Google pays for Windows Server and Microsoft's customer-facing prices

6.373 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.

6.374 Figure 6.9 and Table 6.9 below shows the difference between the wholesale 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.

Figure 6.9: [✂]

Table 6.9: Average difference between the wholesale 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

Revenue bracket ($)	*ACR (%)*	*VM (%)*
10k-1M	[5-10]	[30-40]
1M-5M	[0-5]	[30-40]
5M-10M	[0-5]	[20-30]
10M-20M	[0-5]	[30-40]
>20M	**[0-5]**	**[30-40]**

Source: CMA analysis of Microsoft's response to the CMA's information request [✂]; Microsoft's response to Ofcom's information request [✂].

6.375 Our analysis suggests that the difference between Microsoft's customer facing prices and the wholesale prices it charges Google for Windows Server account for at least [0-5]% of customer spend across all Azure services and as much as [30-40]% of spend on Windows Server VMs across all revenue brackets.

6.376 As discussed above, the additional costs that Google must absorb 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.

6.377 Figure 6.10 and Table 6.10 below shows the median difference between the wholesale 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.

Figure 6.10: [✂]

Table 6.10: Median difference between the wholesale 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

Revenue bracket ($)	*ACR (%)*	*VM (%)*
10k-1M	[0-5]	[30-40]
1M-5M	[0-5]	[20-30]
5M-10M	[0-5]	[10-20]
10M-20M	[0-5]	[40-50]
>20M	**[0-5]**	**[20-30]**

Source: CMA analysis of Microsoft's response to the CMA's information request [✂]; Microsoft's response to Ofcom's information request [✂].

6.378 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.

6.379 Figure 6.11 and Table 6.11 below shows the percentage difference between the wholesale 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.

Figure 6.11 [✂]

Table 6.11: Average percentage difference between the wholesale prices that Google pays for SQL Server and Microsoft's customer-facing prices relating to SQL Server PAYG usage

Revenue bracket ($)	*50:50 Enterprise to Standard (%)*	*40:60 Enterprise to Standard (%)*	*25:75 Enterprise to Standard (%)*
10k-1M	[400-500]	[300-400]	[200-300]
1M-5M	[300-400]	[200-300]	[100-200]
5M-10M	[500-600]	[500-600]	[300-400]
10M-20M	[200-300]	[100-200]	[100-200]
>20M	**[40-50]**	**[30-40]**	**[0-5]**

Source: CMA analysis of Microsoft's response to the CMA's information request [✂]; Microsoft's response to Ofcom's information request [✂].

6.380 Our analysis suggests that the wholesale prices that Microsoft charges Google for SQL Server are at least [0-5]% higher than its customer-facing prices for customers in highest revenue bracket. The wholesale prices are at least [100-200]% higher than Microsoft's customer-facing prices for customers in all other revenue brackets.

6.381 Figure 6.12 and Table 6.12 below shows the median percentage difference between the wholesale prices that Google pays for Windows Server and Microsoft's customer-facing prices relating to SQL Server PAYG usage.

Figure 6.12: [✂]

Table 6.12: Median percentage difference between the wholesale prices that Google pays for SQL Server and Microsoft's customer-facing prices relating to SQL Server PAYG usage

Revenue bracket ($)	*50:50 Enterprise to Standard (%)*	*40:60 Enterprise to Standard (%)*	*25:75 Enterprise to Standard (%)*
10k-1M	[100-200]	[100-200]	[80-90]
1M-5M	[100-200]	[80-90]	[40-50]
5M-10M	[100-200]	[100-200]	[70-80]
10M-20M	[70-80]	[50-60]	[20-30]
>20M	[30-40]	[20-30]	[-5-0]

Source: CMA analysis of Microsoft's response to the CMA's information request [✂]; Microsoft's response to Ofcom's information request [✂].

6.382 Figure 6.13 and Table 6.13 below shows the difference between the wholesale 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.

Figure 6.13: [✂]

Table 6.13: Average difference between the wholesale 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

	50:50 Enterprise to Standard		*40:60 Enterprise to Standard*		*25:75 Enterprise to Standard*	
Revenue bracket ($)	ACR (%)	VM (%)	ACR (%)	VM (%)	ACR (%)	VM (%)
10k-1M	[10-20]	[40-50]	[10-20]	[30-40]	[5-10]	[20-30]
1M-5M	[0-5]	[10-20]	[0-5]	[5-10]	[0-5]	[0-5]
5M-10M	[0-5]	[10-20]	[0-5]	[5-10]	[0-5]	[0-5]
10M-20M	[0-5]	[5-10]	[0-5]	[5-10]	[0-5]	[0-5]
>20M	**[0-5]**	**[5-10]**	**[0-5]**	**[5-10]**	**[0-5]**	**[0-5]**

Source: CMA analysis of Microsoft's response to the CMA's information request [✂]; Microsoft's response to Ofcom's information request [✂].

6.383 Our analysis suggests that the difference between Microsoft's customer facing prices and the wholesale prices it charges Google for SQL Server account for at least [0-5]% of customer spend across all Azure services and as much as [40-50]% of spend on Windows Server VMs and SQL Server IP across the four highest revenue brackets.

6.384 As discussed above, the additional costs that Google must absorb 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.

6.385 Figure 6.14 and Table 6.14 below shows the median difference between the wholesale 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.

Figure 6.14: [✂]

Table 6.14: Median difference between the wholesale 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

	50:50 Enterprise to Standard		*40:60 Enterprise to Standard*		*25:75 Enterprise to Standard*	
Revenue bracket ($)	ACR (%)	VM (%)	ACR (%)	VM (%)	ACR (%)	VM (%)
10k-1M	[5-10]	[30-40]	[5-10]	[20-30]	[0-5]	[10-20]
1M-5M	[0-5]	[5-10]	[0-5]	[0-5]	[0-5]	[0-5]
5M-10M	[0-5]	[5-10]	[0-5]	[0-5]	[0-5]	[0-5]
10M-20M	[0-5]	[0-5]	[0-5]	[0-5]	[0-5]	[0-5]
>20M	**[0-5]**	**[0-5]**	**[0-5]**	**[0-5]**	**[-5-0]**	**[-5-0]**

Source: CMA analysis of Microsoft's response to the CMA's information request [✂]; Microsoft's response to Ofcom's information request [✂].

6.386 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.

6.387 Figure 6.15 and Table 6.15 below shows the combined difference in the wholesale 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.

Figure 6.15: [✂]

Table 6.15: Average combined difference between the wholesale 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

	50:50 Enterprise to Standard		*40:60 Enterprise to Standard*		*25:75 Enterprise to Standard*	
Revenue bracket ($)	ACR (%)	VM (%)	ACR (%)	VM (%)	ACR (%)	VM (%)
10k-1M	[10-20]	[50-60]	[10-20]	[40-50]	[10-20]	[30-40]
1M-5M	[5-10]	[30-40]	[5-10]	[30-40]	[0-5]	[20-30]
5M-10M	[5-10]	[30-40]	[5-10]	[30-40]	[0-5]	[20-30]
10M-20M	[5-10]	[40-50]	[5-10]	[30-40]	[5-10]	[30-40]
>20M	**[0-5]**	**[40-50]**	**[0-5]**	**[30-40]**	**[0-5]**	**[30-40]**

Source: CMA analysis of Microsoft's response to the CMA's information request [✂]; Microsoft's response to Ofcom's information request [✂].

6.388 Our analysis suggests that the average difference between the wholesale 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 [0-5]% of spend across all Azure services and as much as [50-60]% across the revenue brackets.

6.389 As discussed above, the additional costs that Google must absorb 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.

6.390 Figure 6.16 and Table 6.16 below shows the median combined difference in the wholesale 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.

Figure 6.16: [✂]

Table 6.16: Median combined difference between the wholesale 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

50:50 Enterprise to Standard	*40:60 Enterprise to Standard*	*25:75 Enterprise to Standard*

Revenue bracket ($)	ACR (%)	VM (%)	ACR (%)	VM (%)	ACR (%)	VM (%)
10k-1M	[10-20]	[30-40]	[5-10]	[30-40]	[5-10]	[20-30]
1M-5M	[5-10]	[20-30]	[0-5]	[20-30]	[0-5]	[10-20]
5M-10M	[0-5]	[30-40]	[0-5]	[30-40]	[0-5]	[20-30]
10M-20M	[5-10]	[40-50]	[0-5]	[40-50]	[0-5]	[40-50]
>20M	**[0-5]**	**[20-30]**	**[0-5]**	**[20-30]**	**[0-5]**	**[20-30]**

Source: CMA analysis of Microsoft's response to the CMA's information request [✂]; Microsoft's response to Ofcom's information request [✂].

6.391 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.

6.392 The magnitude of the percentage differences between Microsoft's customer-facing prices and the wholesale prices that it charges AWS[1801] 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.

6.393 The differences between Microsoft's customer-facing prices and the wholesale prices it charges AWS and Google as a proportion of customer spend provide an indication of the extent of the effect of Microsoft's conduct on competition. These results show that AWS and Google would have to absorb at least several percentage points' higher costs to compete for such customers, with these cost proportions increasing for customers that allocate their cloud spend in a less aggregated way.

6.394 While the difference in Microsoft's customer-facing prices and the wholesale 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 Server[1802] (and exclusively on GCP) such that it is appropriate to consider the combined impact of the conduct relating to Windows Server and SQL Server.

6.395 Additionally, as noted above, marketing material of Microsoft 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

[1801] See Appendix T for the results of our comparison of Microsoft's customer-facing prices and the wholesale prices it charges AWS for Windows Server and SQL Server

[1802] 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 [✂].

without impacting AWS' or Google's input costs, although our analysis has not assessed that.

6.396 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 conduct relative to customer spend across all cloud services.

A cloud provider's analysis of Windows Server licensing cost versus infrastructure price

6.397 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.[1803] 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.

6.398 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.[1804] [✂].[1805]

6.399 The analysis uses a similar numerator to the analysis 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).[1806] The cloud provider uses the price of infrastructure for that workload on Azure as the denominator.[1807]

6.400 The table below shows the average difference in licensing related costs as a proportion of Azure infrastructure prices.

Table 6.171: The cloud provider's analysis of differences in licensing costs paid by that cloud provider compared to customer facing price for Azure customers

[✂]

6.401 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 customers would pay to use Windows Server on Azure accounts for [✂] of Azure's infrastructure price under PAYG plans on general purpose instances.

[1803] [✂] submission to the CMA [✂].
[1804] [✂] submission to the CMA [✂].
[1805] [✂] submission to the CMA [✂].
[1806] [✂].
[1807] [✂] submission to the CMA [✂].

6.402 The cloud provider submitted that the results show that the difference in licence-related cost is greater than [✂].[1808] [1809]

6.403 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.[1810]

6.404 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 by the cloud provider (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) [✂].[1811]

(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.[1812] 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.

6.405 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 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

1808 [✂]. [✂] submissions to the CMA [✂].
1809 [✂]. Submissions to the CMA [✂].
1810 [✂] submission to the CMA [✂].
1811 [✂] response to the CMA's information request [✂].
1812 The analysis currently weights each instance equally.

interpretation of the results.[1813] Further, while the denominator used by the cloud provider is narrow, we consider that the results are consistent with the analysis we conducted, which uses a wider denominator.

Price and non-price differences for the client-side products

6.406 In this section we consider Microsoft's conduct in relation to the client-side products whether this may foreclose rivals by impacting the quality of AWS, Google's and other cloud providers' competitive offerings.

6.407 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. We consider the client-side products to be 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 Desktop.

6.408 Table 6.18 outlines the licensing practices with respect to the client-side products.

Table 6.18: Availability of the Microsoft products for use in the public cloud

	Office desktop applications	Microsoft 365	Windows Desktop	Visual Studio
Microsoft	Can BYOL with software assurance	Microsoft 365 applications available to run in Azure.[1814]	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.[1815] Cannot BYOL.[1816] Currently available under third party SPLA, [✂].[1817]	Not available through the SPLA.[1818] BYOL to Amazon Workspaces for a subset of Microsoft 365 packages only.[1819] Not available on Amazon Appstream.[1820]	Not available through the SPLA.[1821] Cannot BYOL to the public cloud.* Currently made available for resale through third parties to host on Listed Providers infrastructure [✂].[1822]	Available via the AWS SPLA for use with EC2.[1823] Limited ability to BYOL: Customers cannot BYOL the [✂] of Visual Studio to AWS Workspaces.[1824]
Google	Technically 'available' under the SPLA [✂].[1825]	Not available via BYOL or SPLA.[1827]	Not available through SPLA.[1829]	Available through the SPLA [✂].[1831]

1813 In fact, adjusting for relative usage may either reduce *or* increase the point estimates.
1814 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 [✂].
1815 [✂] response to the CMA's information request [✂].
1816 [✂] response to the CMA's information request [✂].
1817 [✂] submission to the CMA [✂].
1818 [✂] submission to the CMA [✂].
1819 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 15th 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 [✂].
1820 Note of meeting with [✂].
1821 Microsoft's response to the CMA's information request [✂].
1822 [✂] submission to the CMA [✂].
1823 Microsoft licensing on AWS - AWS Prescriptive Guidance, accessed 04 October 2024.
1824 [✂] submission to the CMA [✂].
1825 [✂]. [✂] submission to the CMA [✂].
1827 [✂] submission to the CMA [✂].
1829 [✂] submission to the CMA [✂]. Microsoft's response to the CMA's information request [✂].
1831 [✂] submission to the CMA [✂].

	BYOL unavailable to the public cloud for licenses post-2019.[1826] Currently available under third party SPLA to use on GCP, [✂].	Microsoft has commented that Google offers its own Workspaces solution, so it does not need to make this available to Google.[1828]	Cannot BYOL to the public cloud.* Currently made available for resale through third parties to host on Listed Providers infrastructure [✂].[1830]	Customers cannot BYOL.[1832]
Non-Listed Providers	Can BYOL with Software Assurance.	Microsoft 365 Applications available to run in Non-Listed Provider clouds. [1833]	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.[1834]

6.409 Table 6.18 shows that:

(a) Google and AWS cannot resell Microsoft 365 or Windows Desktop through their SPLAs.

(i) Google submitted that [✂].[1835]

(ii) [✂].[1836] [✂].

(b) For Windows Desktop, 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%.[1837]

(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.[1838]

6.410 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

[1826] [✂] response to the CMA's information request [✂].
[1828] [✂].
[1830] [✂] submission to the CMA [✂].
[1832] [✂] submission to the CMA [✂].
[1833] Microsoft's submission to the CMA [✂].
[1834] Bring Your Own Windows desktop licenses in WorkSpaces Personal - Amazon WorkSpaces, accessed 15 July 2024; Microsoft FAQ – AWS, accessed 15th July 2024.
[1835] [✂]. Google's submission to the CMA [✂].
[1836] [✂] response to the CMA's information request [✂].
[1837] [✂]. [✂] submission to the CMA [✂].
[1838] The VDA fee does not apply to customers wanting to run Windows Desktop on Azure or non-Listed Providers' private clouds. AWS' submission to the CMA [✂].

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.[1839,] 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.[1840] 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.[1841]

6.411 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 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.[1842] 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.[1843]

6.412 Microsoft submitted that customers can purchase subscriptions to Microsoft 365 and Windows desktops and run those applications in any non-Listed Provider cloud as well as Azure.[1844] It said that customers that want to use a Windows based Desktop as a Service ('**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.[1845] Customers that want to use a Windows-based DaaS solution from AWS can now bring the Microsoft 365 Apps to AWS WorkSpaces directly.[1846]

6.413 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

[1839] Microsoft's submission to the CMA [✂]. 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 [✂].
[1840] Microsoft's submission to the CMA [✂].
[1841] Microsoft's submission to the CMA [✂].
[1842] Microsoft's submission to the CMA [✂].
[1843] Microsoft's submission to the CMA [✂].
[1844] Microsoft's submission to the CMA [✂].
[1845] Microsoft's submission to the CMA [✂].
[1846] Microsoft's submission to the CMA [✂].

why a customer cannot bring the client-side versions of 365 to a VDI solution on third party infrastructure.

6.414 [✂]. But customers are still prohibited from bringing their own pre-purchased licences of that client-side software.[1847]

6.415 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.[1848] 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.[1849]

6.416 We have seen that [✂] and the standard SPLAs [✂].[1850]

Relevant price and non-price restrictions

6.417 Below, we consider the price and non-price restrictions for each of the client-side products.

Windows 10/11 multi-session

6.418 Customers are unable to use [✂] VDI providers services to run Windows Desktop in multisession mode. One cloud provider 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.[1851]

6.419 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 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.[1852]

Microsoft 365 and Teams restrictions

6.420 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 license. This means

[1847] Note of meeting with [✂].
[1848] Microsoft's submission to the CMA [✂].
[1849] Microsoft's response to the CMA's information request [✂].
[1850] [✂] response to the CMA's information request [✂].[✂] response to the Ofom's information request [✂].
[1851] [✂] submission to the CMA [✂].
[1852] [✂] submission to the CMA [✂].

that customers that already own a Windows Server licence must repurchase this [✂], 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),[1853] and therefore, cannot BYOL to the public cloud on this mode of deployment. [✂].[1854]

6.421 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 license which includes shared computer activation), whereas this was not the case if they want to use Azure Virtual Desktop, which can run Windows desktop OS.[1855] It has submitted that this means the end customers total cost of ownership is higher when using Microsoft's competitors systems than its own.

6.422 [✂].[1856]

Visual Studio subscriber software and restrictions

6.423 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.[1857] 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.[1858]

6.424 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.

6.425 While anecdotal, [✂] licensing costs with respect to Visual Studio on [✂] compared to on Azure. One example estimated that the cost of the development and testing licences on [✂] this cloud provider was greater than the monthly licensing cost for Visual Studio IDE licences alone.[1859] We understand this cloud provider's [✂] argument to be that the need to purchase additional development

[1853] AWS' submission to the CMA [✂].
[1854] Note of meeting with AWS [✂].
[1855] [✂] submission to the CMA [✂].
[1856] [✂] submission to the CMA [✂].
[1857] Visual Studio Subscriptions (Formerly MSDN Licenses) - Everything You Need To Know, accessed 12 November 2024; Available software downloads - Visual Studio Subscription, accessed 12 November 2024; [✂].
[1858] [✂] response to Ofcom's information request [✂].
[1859] [✂] response to the CMA's information request [✂].

and testing licenses on non-Azure clouds increases costs [✂] above any differences in the licensing cost of Visual Studio IDE on Azure compared to non-Azure clouds.

6.426 AWS submitted that the [✂] of Visual Studio continues to be restricted for customers accessing the software through BYOL on Amazon Workspaces.[1860]

Potential alternative solutions for Listed Providers

6.427 We understand that Windows Server can be used to create a Windows VDI environment, as well as Windows Desktop.[1861] 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.[1862] In addition, [✂] 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.[1863] Therefore, we understand that post 2025, customers who base their VDI network on Windows Server not Windows Desktop will be unable to use it to host Microsoft Office.

6.428 Further, while independent managed service providers can resell Windows Desktop 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.[1864]

Our assessment

6.429 We have presented our comparison of the price that Microsoft charges Azure customers for Windows Server and SQL Server with the wholesale prices it charges AWS and Google, and a cloud provider's analysis of Windows Server licensing cost compared to the infrastructure price. Both analyses address a similar question: how the input cost paid by cloud providers compares to the customer facing price of Microsoft software on Azure.

6.430 While the analyses use slightly different methodologies, they both find that the input costs to rivals are higher than Microsoft's customer facing price.

6.431 Our guidelines set out that, by being subjected to higher input costs, downstream competitors may be unable to compete effectively.[1865] The fact that Microsoft has

[1860] AWS' submission to the CMA [✂].
[1861] [✂] response to the CMA's information request [✂].
[1862] Please refer above for a more detailed explanation of the Windows server licensing conditions.
[1863] [✂] submission to the CMA [✂].
[1864] 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 4 October 2024.
[1865] CC3 (Revised), paragraph 270.

set rivals' 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 rivals.

6.432 We also note that the difference we calculated in customer-facing prices and 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.

6.433 We consider that the cloud provider's analysis 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.

6.434 We have also found that for Windows Desktop, Microsoft 365, Office and Visual Studio, the licensing practices mean that for customers with existing licences, they cannot bring these to Listed Provider'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 Desktop and Microsoft 365 on either AWS or GCP.

6.435 In addition, potential workarounds for cloud providers 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.

Provisional conclusions

6.436 We have provisionally found that Microsoft's conduct is consistent with actions taken as a result of an incentive to partially foreclose rivals.

6.437 We also have provisionally found that the magnitude of the differences between customer-facing prices on Azure and SPLA input costs, and Microsoft's conduct in relation to other price and non-price differences, is significant in the context of customer spend and their available options, and so may have a material effect on competition.

The impact on rivals' competitive offerings from Microsoft's conduct

6.438 In this section, 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 its rivals and competition. We consider:

(a) whether AWS/Google have passed through or tend to pass through input costs;

(b) the extent to which list prices for select instances differ across clouds;

(c) competitor margins and, in particular, whether the margins are low enough to be consistent with high input prices and likely pass-through of high costs;

(d) relative usage of Microsoft software products on Azure compared to AWS;

(e) customer evidence on the reasons for their choice of cloud; and

(f) examples of lost customers from AWS and Google.

Pass-through of input prices

6.439 If a firm is charged a higher input price (for example, an increase over time or a higher price than would be charged but for conduct such as a foreclosure strategy), the firm may react in several different ways. It may pass through the higher price, resulting in a weaker competitive offering overall. It may alternatively absorb some or all of the higher price by holding their price, quality, range, service and innovation levels constant and instead 'absorb' the higher input price by reducing their margins. These choices are not mutually exclusive: a firm can partially absorb and partially pass through a higher input price, spreading the impact of the higher price between its own margin and its competitive offering.[1866]

6.440 In general, we would expect that a profit-maximising firm that experiences higher input costs will tend to pass through at least some of these higher costs.[1867] We interpret the evidence we have set out below in light of this expectation.

6.441 One cloud provider submitted that it had passed through input cost increases. In particular, that cloud provider said that it had 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.[1868]

6.442 One cloud provider 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.[1869] The cloud provider 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.[1870]

[1866] The extent to which costs are passed on depends on a range of factors, including the responsiveness of the demand and supply conditions, and the degree of competition between businesses throughout the supply chain. RBB Economics -cost pass through: theory, measurement and potential policy implications, page 1.
[1867] The theoretical models and empirical evidence considered in the cost pass-through report prepared for the OFT generally indicated that there would be at least some pass through of input costs, though the exact degree of pass through depends on a range of factors. RBB Economics -cost pass through: theory, measurement and potential policy implications.
[1868] [✂] response to the CMA's information request [✂].
[1869] [✂] submission to the CMA [✂].
[1870] [✂] submission to the CMA [✂].

6.443 Evidence on the extent to which a different cloud provider has passed through high input costs on Microsoft software is more mixed:

(a) [✄] submitted it is forced to pass on some of the higher licensing costs to customers.[1871] [✄]. [✄] said that the difference in price relates to the cost of Windows Server licence acquired through the SPLA which [✄] passes through to its customers.[1872] [✄] also said absorbing the costs was not viable in a lot of cases, [✄]. [1873]

(b) [✄] submitted that Microsoft had been raising Windows Server licensing fees over the past several SPLA renewals and that [✄] had historically absorbed these cost increases and not raised prices for its customers.[1874]

(c) [✄] submitted that it had passed on cost benefits for existing services to customers in the form of lower prices for these services, including for [✄].[1875] [✄] said that, in this sense, there may be indirect pass on (as it may have decreased prices if it had lower input costs).[1876] We consider that the analysis [✄] submitted on passing through cost decreases may have limitations.[1877]

6.444 AWS submitted that it could not absorb Microsoft licensing costs such as buying an additional copy of a pre-existing licence for customers, and that doing so would be 'an economic absurdity'.[1878] In comparison, 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 their free Linux VM offering.[1879]

6.445 Microsoft also 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 of at least 40 servers and $250 for migrations of more than 80 Windows Servers.[1880]

6.446 The evidence submitted by the first cloud provider is consistent with high costs leading to it charging a higher price and having a less competitive offering. Evidence over time has also shown a consistent picture, with changes in input prices being associated with changes in customer-facing prices. In relation to the

[1871] [✄] submission to the CMA [✄].
[1872] [✄] submission to the CMA [✄].
[1873] Transcript of hearing with [✄].
[1874] [✄] response to Ofcom's information request [✄].
[1875] [✄] submission to the CMA [✄].
[1876] Transcript of hearing with [✄].
[1877] [✄].
[1878] AWS' submission to the CMA [✄].
[1879] Microsoft's submission to the CMA [✄].
[1880] Microsoft's submission to the CMA [✄].

second cloud provider, there is some evidence that the input prices that Microsoft charges to the cloud provider appear to be associated with relatively high customer-facing prices. However, the cloud provider has historically not passed through changes in input costs over time.

6.447 We consider this evidence on pass-through alongside other evidence when we set out our provisional conclusions below.

List prices

6.448 Microsoft submitted a comparison of list prices for selected instances involving the use of Microsoft software across AWS, Azure and GCP. Before setting out this analysis, we make the following observations which are relevant to its interpretation.

(a) First, while list prices could in principle provide an indication of how cloud providers' competitive offerings compare, a comparison of list prices will not take into account any pass-through of high input prices that is actioned through lower discounts instead of higher list prices. As set out in Chapter 7, CSDs alone covered a significant proportion of total annual cloud spend.[1881] A cloud provider also submitted that discounts are typically at the deal level,[1882] which would imply that prices on deals involving Microsoft software may understate price differences.

(b) Second, Microsoft's submission presents the differences in list prices for Windows VM PAYG prices. As such, these are the list prices customers would pay if they do not have a pre-existing Windows Server licence that qualifies for AHB. If they did, the customer would therefore pay a lower price on Azure.

(c) Third, Microsoft selected these instances [✂] and compared them to similar instances provided by AWS and Google. [1883] We do not know how the list prices of other instances compare across AWS, Azure and GCP.

6.449 We present Microsoft's submission below.

Table 6.192: Microsoft's submission on list prices for Windows Server VM PAYG prices across AWS, Azure and GCP

Instance family	*AWS instance series*	*Azure instance series*	*GCP instance series*	*AWS Windows list price ($ per vcore hour)*	*Azure Windows list price ($ per vcore hour)*	*GCP Windows list price ($ per vcore hour)*
General purpose	M4	Dv3	N1, N2	$0.104	$0.104	$0.109*

[1881] See the CSAs section 'Prevalence of CSDs'.
[1882] [✂] submission to the CMA [✂].
[1883] [✂] submission to the CMA [✂].

Compute optimised	C5	FSv2	C2	$0.097	$0.097	$0.113*
Memory optimised	R5	Ev3	M1, M2	$0.120	$0.124*	$0.239*
Storage optimised	I3	LSv2	Z3	$0.137	$0.137	Unknown**
GPU	G3	NCaST4_v3	G2	$0.135	$0.154***	$0.141*

** this is the model price for virtual machines within this instance series; the price varies for individual virtual machines*
*** the price for this virtual machine is unknown as it is yet to be generally available [at the time Microsoft submitted the evidence]*
**** this is the mean price for virtual machines within this instance as price varies by individual virtual machine for this series. In this particular case, it wasn't possible to estimate the mode as each series in this instance had different prices*
Source: Microsoft's submission to the CMA [✂].

6.450 common instance. For the same instances, AWS has equal or lower list prices than Azure's list prices.

6.451 We consider that the list prices analysis submitted by Microsoft provides some evidence on the differences in competitive offerings for customers that do not qualify for AHB on Azure. In particular, the results indicate that, for these instances, AWS' competitive offering is similar or better than Microsoft's, whereas Google's competitive offering is at least marginally worse.

6.452 However, we consider it to be appropriate to be cautious in interpreting these results, given the limitations we have set out above. As such, we have considered the evidence in the round with the other evidence we have on competitive offerings.

Margins and mark-ups

6.453 Microsoft submitted analyses of the mark-ups that AWS and Google charge and of the margin that Microsoft would hypothetically make if it were charged SPLA prices for its use of the relevant software. Microsoft submitted that the observation that AWS and Google earn a margin is evidence that there is no impediment to competition.[1884] Google has also submitted data on the total revenues and costs it faced on Microsoft workloads, which we have used to calculate margins. Each of these pieces of analysis are discussed below.

6.454 As an initial observation, we consider that a rival having a positive margin can be consistent with a foreclosure strategy affecting their competitive offering. By raising a rival's costs, that rival may be induced to increase its prices or otherwise worsen their competitive offering, which may in turn lead to that rival exerting a weaker competitive constraint on other suppliers in the market.

6.455 There is no reason to think that a firm's input costs will affect their prices only when their margins drop to zero. We think it would be unreasonable to assume that partial foreclosure can be prevented by rivals moving away from profit-

[1884] [✂] submission to the CMA [✂].

maximising behaviour and severing the link between their input costs and their prices. Nevertheless, we have considered the relevant analyses of margins and mark-ups as it could plausibly provide useful context to take into account alongside other evidence.

6.456 Microsoft submitted two analyses related to margins. In the first, Microsoft calculated the mark-up over the Windows Server software license IP cost as the difference between the premium for a Windows VM over a Linux VM of the same VM specification and AWS' and Google's SPLA costs per vcore hour. Based on this, Microsoft submitted that AWS and Google currently charge list prices with a markup over their SPLA costs of approximately [10-20]%.[1885]

6.457 In its second analysis, Microsoft calculated the hypothetical margin it would make if it faced the SPLA licensing cost it charges AWS and Google for Windows Server. 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. Microsoft submitted that this analysis takes a narrow view of the competitive process by focusing exclusively on Windows Server VMs, rather than considering broader packages. It said that considering broader packages would show even more favourable economics to AWS and Google. Microsoft added that the (global) analysis was conservative because it assumed that Windows Server VMs were 100% utilised. Microsoft found that even at [90-100]% utilisation Windows Server VMs would break even in 2023 globally and that in 2023 Q3 and 2023 Q4, utilisation was [90-100]% and [90-100]%, respectively.[1886]

6.458 We make the following observations relevant to the interpretation of these analyses:

(a) The analysis on the margins that AWS and Google make on Windows Server VMs implicitly assumes that AWS and Google receives the list price for their services. To the extent AWS and Google discount on their list prices, the margins calculated by Microsoft will tend to overstate their margins. We note that, for example, a cloud provider submitted that it discounts at the deal level.[1887]

(b) This analysis also includes all customers and all usage, regardless of whether they had a pre-existing licence for Windows Server that qualifies for AHB. AWS and Google would have to discount substantially to match Microsoft's offering on such workloads (as set out above). Therefore, the

1885 [✂] submission to the CMA [✂].
1886 [✂] submission to the CMA [✂].
1887 [✂] submission to the CMA [✂].

margin AWS and Google would make on those workloads would be even lower than the calculated margins. Given the margins are already relatively low (particularly on a global basis), even with both workloads with licences that qualify and do not qualify for AHB, this implies that Microsoft would struggle to offer the discounts it is necessary for AWS and Google to offer to match the AHB discount.

6.459 Google also submitted data on its total revenues and costs incurred directly from licence sales of Windows Server and SQL Server.[1888] We used this data to calculate the 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. In particular, between 2021 and 2023, Google made an estimated margin of:

(a) [-10- -5]% to [0-5]% on SQL Server Enterprise, which accounted for [10-20]% of usage in 2023;

(b) [20-30]% to [20-30]% on SQL Server Standard, which accounted for [40-50]% of usage in 2023; and

(c) A negative margin [✂] on SQL Server Web, which accounted for [30-40]% of usage in 2023.

6.460 We observe the following points about this analysis:

(a) Usage varies significantly by revenue band: higher-spend customers use SQL Server Enterprise more frequently, whereas lower-spend customers use Web more frequently;

(b) Google's analysis only captures the margins that Google makes on the customers that it wins and by its construction excludes the margins that Google would have made on the customers that it loses or does not compete for because it would not earn sufficient margin. This may tend to cause the analysis to overstate Google's margins.

(c) The observations made in relation to Microsoft's analysis above on the inclusion of all customers and all usage, regardless of whether they had a pre-existing licence for Windows Server that qualifies for AHB is also relevant to this analysis.

6.461 We note that the calculated margins and mark-ups do not appear to be particularly high in the context of software markets, particularly given that there are several

1888 [✂] response to the CMA's information request [✂]. 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.

factors which may cause the calculated margins to represent overestimates of the relevant margin. We discuss this evidence in the context of our provisional views below.

6.462 The evidence suggests that AWS and Google have a positive mark-up on list prices for Microsoft workloads, and that they make a margin on those workloads. However, we consider that positive and low margins are not inconsistent with Microsoft's licensing practices having an effect on AWS' and Google's competitive offerings. The low margins on Microsoft workloads may suggest AWS and Google are less likely to compete fiercely for Microsoft workloads, or for customers without sufficient non-Microsoft workloads to compensate for the low margins.

Relative usage

6.463 We compared Windows Server and SQL Server usage levels across Azure, AWS, and GCP. To the extent that usage of Windows Server and SQL Server is significantly higher on Azure than on AWS and GCP, this would be consistent with Microsoft's licensing conduct 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.

6.464 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 over 2020-2023; Google provided us with equivalent data for the years 2021-2023.[1889] We used this data to compare the relative average usage rates of Windows Server and SQL Server on each platform over time.

6.465 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.[1890] 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.[1891] 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.

6.466 This analysis cannot determine with certainty whether there is causal link between Microsoft's conduct and relative usage levels. Variation in usage levels across providers could in principle be explained by other factors, including for example a

[1889] 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.
[1890] Microsoft's time-series data includes usage of the Enterprise, Standard, Web, Express, Developer, and Power BI editions of SQL Server.
[1891] [✂] response to the CMA's information request [✂].

more general preference to use Microsoft software on Azure or the use of different software products that are substitutes for the Microsoft software on AWS and GCP. 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.

6.467 The figure below shows the relative average usage of Windows Server over the years 2020-2023 on Azure and AWS and the years 2021-2023 on GCP.

Figure 6.17**: Average Windows Server usage on Azure, AWS, and GCP in each Year**

[✂]

Source: CMA analysis of responses to the CMA's information requests [✂].

6.468 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.

6.469 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 6.18**: Total average usage of SQL Server on Azure compared with average usage of SQL Server on AWS and GCP as licensed through the SPLAs in each Year**

[✂]

Source: CMA analysis of responses to the CMA's information requests [✂].

6.470 Our analysis suggests that SQL Server usage on Azure is significantly higher than on AWS and GCP. In particular, 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.

6.471 However, as explained above, this graph overstates the difference in usage of SQL Server on Azure compared with AWS and GCP to some extent. We present our sensitivity analysis below.

6.472 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 usage across all editions and deployment types for comparison.

Figure 6.19: **SQL Server average usage in 2022**

[✂]

Source: CMA analysis of responses to the CMA's information requests [✂].

6.473 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.

6.474 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.

6.475 The results of our relative usage analysis suggest that customers disproportionately choose to deploy their Windows Server and SQL Server workloads on Azure. 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 a substantial effect on customer choice.

6.476 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 the licensing conduct 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.

Customers' choice of cloud providers

6.477 We asked customers about their perception of the difference between providers' competitiveness in terms of hosting workloads that use Microsoft software as an input, the factors that affect their choice of cloud provider and the role of licensing practices in their choices.

6.478 We have set out the evidence we gathered from randomly selected AWS, Microsoft and Google customers from (i) different revenue brackets[1892] for

[1892] In particular, the revenue brackets were < $10,000, $10,000 – $1Million, $1Million – $5Million, $5Million – $10Million, $10Million – $20Million, > $20Million.

licensing specific questions and (ii) large customers for general public cloud questions, including the importance of different factors in their choice of cloud.

6.479 In order to gather further evidence on customers' use of the Microsoft software on VDI, we randomly selected additional customers, some of which we identified were active in industries more likely to use VDI, and some of which were randomly selected. We have also considered customer views in the Jigsaw report.

Perceived differences in competitive offerings

6.480 We gathered evidence from customers on their perceptions of any differences in terms of price between using Microsoft software products[1893] on Azure compared to other public clouds:

(a) most customers said that there were price advantages from using Microsoft software products on Azure;[1894]

(b) of these customers, the majority indicated that the price advantage arose from receiving AHB on Azure compared to needing to repurchase Microsoft licences for use on other clouds.[1895,1896] 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.[1897]

6.481 Similarly, when asked about the selection factors considered when choosing a public cloud, a large customer submitted that Microsoft make the cost and ability to use software licences with competitors such as AWS difficult, particularly with respect to Windows Server and SQL Server. It specified that it is uncompetitive to use a competitor cloud for Microsoft software products use cases because it is expensive, complex and difficult.[1898]

6.482 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 did not identify quality differences in using Microsoft software products across different clouds.

[1893] The question referred to Windows Server, SQL Server, Windows 10/11, Microsoft Office, Microsoft 365 and Active Directory/Azure AD specifically.
[1894] Responses to the CMA's information requests [✂].
[1895] 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 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.
[1896] Responses to the CMA's information requests [✂].
[1897] [✂] response to the CMA's information request [✂].
[1898] [✂] response to the CMA's information request [✂].

Selection factors

6.483 Customer evidence was largely consistent with Microsoft's licensing practices having an effect on their choice of public cloud provider. In particular:

(a) We sought to understand the selection factors customers consider in their choice of public cloud. When prompted to consider a list of possibly relevant factors,[1899] most customers identified licensing as a factor in their choice of Azure.[1900] Similarly, many large customers rated the cost and ability to use software licences as important or very important in their choice of public cloud.[1901] In their free-text explanations, some customers explicitly mentioned the ability to bring on-premises licences to the public cloud was important.[1902]

(b) 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.[1903] 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.[1904]

(c) Many customers said that licensing terms affected their choice of Microsoft workloads.[1905]

6.484 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.[1906]

6.485 There was a broad consensus among the customers that existing skills and familiarity with the Microsoft ecosystem were factors in choosing Azure as their cloud provider.[1907] A few of these customers identified themselves as a 'Microsoft

---

[1899] 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.
[1900] Responses to the CMA's information requests [✂].
[1901] We also asked large customers to provide an indication of the importance of different factors (provided by the CMA) that they consider when choosing their main public cloud provider.
[1902] Responses to the CMA's information requests [✂].
[1903] CMA commissioned primary customer research conducted by Jigsaw, dated 23 May 2024, paragraph 7.3.2.
[1904] CMA commissioned primary customer research conducted by Jigsaw, dated 23 May 2024, paragraph 7.2.3.
[1905] Responses to the CMA's information requests [✂].
[1906] CMA commissioned primary customer research conducted by Jigsaw, dated 23 May 2024, paragraph 1.4.30-31.
[1907] 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

shop' or a 'Microsoft first' organisation.[1908]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'.[1909]

6.486 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.[1910] One of these customers said that licensing was not a persuasive factor in its decision making.[1911] Similarly, a few large customers rated the cost and ability to use software licences as unimportant or very unimportant in their choice of public cloud. Of these, two customers indicated that licences were not relevant for their applications[1912] (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.[1913]

(b) As set out above, most customers did not identify quality differences between using Microsoft software products across public clouds. They therefore said that their choice of cloud was not affected. For example, one customer said that it considered other factors, such as scalability and flexibility.[1914]

(c) Many customers said that licensing terms did not affect their choice of workload placement,[1915] with some of these specifying that other factors were more important in their choice of cloud.[1916]

6.487 In addition, we asked customers specifically about their choice of cloud provider for their VDI workloads. Almost all customers we contacted responded that they use VDI.[1917] Some of these customers host VDI on-premises.[1918] Considering the customers that host VDI on public cloud:

(a) many customers which we contacted host their VDI network on Azure;[1919]

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 [✂].

1908 Responses to the CMA's information requests [✂].

1909 CMA commissioned primary customer research conducted by Jigsaw, dated 23 May 2024, paragraph 7.3.1.

1910 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 [✂].

1911 [✂] response to the CMA's information request [✂].

1912 Responses to the CMA's information requests [✂].

1913 [✂] response to the CMA's information request [✂].

1914 [✂] response to the CMA's information request [✂].

1915 Responses to the CMA's information requests [✂].

1916 Responses to the CMA's information requests [✂].

1917 Responses to the CMA's information requests [✂].

1918 Responses to the CMA's information requests [✂].

1919 Responses to the CMA's information requests [✂].

(b) some customers host their VDI network on both Azure and AWS;[1920]

(c) a few customers host their VDI network both on Azure and on-premises;[1921] and

(d) one customer did not host any of its VDI network on Azure, hosting across both AWS and IBM.[1922]

6.488 We asked some customers why they chose the public cloud to host their VDI workloads.

(a) Many customers which we contacted mentioned costs[1923] and whether the VDI provider is an existing supplier[1924] as important factors.

(b) Some customers mentioned price and non-price differences relating to Microsoft products as reasons for choosing their public cloud provider:

(i) Some customers mentioned choosing Azure as they already use Microsoft software;[1925]

(ii) Some customers mentioned choosing Azure to enable Windows 10/11 multisession use (a capability not available on other public clouds);[1926]

(iii) Some customers mentioned licensing costs or lower costs for Microsoft products.[1927]

6.489 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;[1928]

6.490 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.[1929]

6.491 In order to gather further customer evidence, we asked some further customers (who were not selected based on where they deployed their VDI network) whether

---

1920 Responses to the CMA's information requests [✂].
1921 Responses to the CMA's information requests [✂].
1922 [✂] response to the CMA's information request [✂].
1923 Responses to the CMA's information requests [✂].
1924 Responses to the CMA's information requests [✂].
1925 Responses to the CMA's information requests [✂].
1926 Responses to the CMA's information requests [✂].
1927 Responses to the CMA's information requests [✂].
1928 Responses to the CMA's information requests [✂].
1929 Responses to the CMA's information requests [✂].

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;[1930]

(ii) One customer referred to the importance of being able to run the Windows 10/11 multisession feature;[1931]

(iii) One customer referred to the fact that it did not have to pay the VDA per-user licence on Azure.[1932]

6.492 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 reported benefits (or hypothetical benefits) to having some workloads located in the same public cloud as their VDI network.[1933]

(i) Workloads which some customers we contacted mentioned as preferable to host in the same public cloud as their VDI network included data,[1934] core line of business applications,[1935] key cloud-based workloads[1936] and some customers reported unspecified applications or workloads.[1937]

(ii) The primary reason for customers hosting workloads in the same public cloud as their VDI network was latency.[1938] 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.[1939]

(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

[1930] Responses to the CMA's information requests [✂].
[1931] [✂] response to the CMA's information request [✂].
[1932] [✂] response to the CMA's information request [✂].
[1933] Responses to the CMA's information requests [✂]. 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.
[1934] Responses to the CMA's information requests [✂].
[1935] [✂] response to the CMA's information request [✂].
[1936] [✂] response to the CMA's information request [✂].
[1937] Responses to the CMA's information requests [✂].
[1938] Responses to the CMA's information requests [✂].
[1939] [✂] response to the CMA's information request [✂].

be the customer's overall public cloud strategy which would determine its location of VDI workloads, not the other way around.[1940]

(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.[1941]

Stakeholder views

6.493 We asked a VDI provider about where customers chose to host their VDI workloads. It said that vast majority of its customers who host their VDI workloads on the public cloud did so on Azure.[1942]

6.494 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.[1943]

6.495 Another VDI provider which we contacted 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.[1944]

6.496 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.[1945] 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.[1946]

[1940] [✂] response to the CMA's information request [✂].
[1941] Responses to the CMA's information requests [✂].
[1942] Note of meeting with [✂].
[1943] Note of meeting with [✂].
[1944] Note of meeting with [✂].
[1945] Note of meeting with [✂].
[1946] Note of meeting with [✂].

Our assessment

6.497 Customer evidence shows that licensing terms may 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.

6.498 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, even 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 may still influence future decision making. In particular, for Azure customers considering switching, licensing terms may result in an additional friction to doing so.

6.499 However, the evidence suggests that most customers are not aware of and/or do not consider quality factors in their choice of public cloud. But we have also seen evidence that suggests that, for some customers, non-price factors 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.

6.500 Stakeholders' views are consistent with Microsoft's licensing practices having an effect on customer choice of cloud with respect to VDI workloads.

Customer examples from AWS and Google

6.501 AWS and Google submitted customer examples that they said show the effect of Microsoft's licensing practices.[1947] We set out an overview of these examples below.

(a) AWS submitted a number of anonymised customers that it said illustrated the harmful impact of Microsoft's licensing. 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.[1948] For example, AWS said that one of these customers, a [✂], was previously a customer of AWS but migrated its

[1947] Google's response to the CMA's information request [✂]. AWS' submission to the CMA [✂].
[1948] AWS' submission to the CMA [✂].

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.[1949]

(b) Google submitted [✂] which it said illustrate the prohibitive costs associated with Microsoft's licensing practices. One of these [✂] indicated that the customer would be willing to pay a [✂]% premium to use the cloud provider's cloud since it was their primary and preferred cloud. However, the customer found that AHB meant that the cloud provider was several times more expensive than Azure and it ultimately decided to migrate their application to Azure.[1950]

6.502 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. [1951] According to the Google's calculations, total cloud costs were several times more expensive on its cloud compared to on Azure, despite the IaaS generally being cheaper on the non-Azure cloud.[1952]

6.503 While these examples are to an extent illustrative of the scope for the licensing practices to give rise to an effect on the competitive offering of AWS and Google, we place very limited weight on them. This is because the examples are anecdotal and we are unable to assess how representative they are of customers' choices impacted by the licensing practices. We attach greater weight to the customer evidence we have gathered directly.

Submissions from cloud providers

6.504 AWS and Google have submitted that they have lost customers due to Microsoft's restrictions with respect to the 'client-side' products.

(a) AWS submitted that it had lost [✂] 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.[1953] This provider's total revenue for its VDI solutions was around [✂] USD in 2022.[1954]

[1949] AWS' submission to the CMA [✂].
[1950] Google's response to the CMA's information request [✂].
[1951] Google's response to the CMA's information request [✂].
[1952] Google's submission to the CMA [✂].
[1953] AWS' response to the CMA's information request [✂].
[1954] AWS' response to the CMA's information request [✂].

(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.[1955]

6.505 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.[1956] 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 [✂]. The same Google internal document also noted that Windows Server [✂] is significantly more expensive than Azure (via Azure Hybrid Benefit) and some VDI workloads such as Microsoft 365 apps are not eligible for resale [✂] and are not BYOL eligible, limiting Google's ability to enter this market as a first party VDI provider.[1957]

(b) The same Google internal document outlined that while it is working with partners [✂] 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.[1958]

(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.[1959] One internal document provided a sample list of lost VDI deals, where out of [✂] deals presented,

[1955] Google's response to the CMA's information request [✂]. In support of its response Google provided the internal documents summarised below.
[1956] We note that the majority of documents responsive to this question are summarised as part of the Importance of the input section above.
[1957] Google's response to the CMA's information request [✂].
[1958] The availability of third party SPLA for use on listed provider's public clouds is detailed above in section [✂]. Google's response to the CMA's information request [✂].
[1959] Google's response to the CMA's information request [✂].

[✄] were reported as being lost directly due to Microsoft's licensing restrictions:

(i) [✄].

(ii) [✄].[1960]

(d) Google has also submitted [✄] 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.[1961]

(i) One deal valued at [✄] USD was said to have been lost due to client deciding to move their VDI to Azure due to Microsoft licensing.[1962]

(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.'[1963]

6.506 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.

6.507 We asked AWS about the effect of customers being able to BYOL some Microsoft 365 licenses to Amazon Workspaces. AWS said that the ability to host Microsoft 365 on Workspaces has [✄]. 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, [✄].[1964]

6.508 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.[1965] 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 server products to be used to build cloud services, whereas client products are designed to be installed on individual PCs for individual use.[1966]

1960 Google's response to the CMA's information request [✄].
1961 Google's response to the CMA's information request [✄].
1962 Google's response to the CMA's information request [✄].
1963 Google's response to the CMA's information request [✄].
1964 Note of meeting with AWS [✄].
1965 Microsoft's submission to the CMA [✄].
1966 Microsoft's submission to the CMA [✄].

6.509 Microsoft has also submitted that Windows Server and SQL Server are the only software services potentially affecting downstream competition.[1967]

6.510 In addition, 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-prem licences that would benefit from AHB. [✂].[1968]

Our assessment

6.511 We have assessed the extent to which observed outcomes are consistent or inconsistent with Microsoft's licensing practices having an impact on competitors and competition. This allows us to assess whether AWS and/or Google have the ability to absorb or counteract the effect.

6.512 We have considered the impact on rivals' competitive offerings in light of the significance of Microsoft's conduct (as set out in the previous section): this showed that Microsoft currently sets a high input cost for AWS and Google to host Windows Server and SQL Server. This input cost is higher than the price that Microsoft charges its own customers to be able to use Windows Server and SQL Server on Azure, provided the customer has pre-existing licences for the product (and qualifies for AHB). This difference for Windows Server, and Windows Server and SQL Server combined, accounts for a significant portion of the amount that customers spend on Windows Server VMs and a smaller, though still material, portion of the total amount customers spend on Azure.

6.513 In considering whether AWS and Google have the ability to absorb or counteract the effect of Microsoft's conduct, we have considered how their competitive offering compares to Microsoft's.[1969] Evidence from both AWS and Google is consistent with high input costs resulting in them charging a higher price and having a less competitive offering. One of these, [✂], has also passed through changes in the input cost, though the other, [✂], has historically absorbed increases.

6.514 We have also 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.

6.515 We found that AWS' and Google's margins are relatively low for Windows Server and SQL Server workloads.

[1967] Microsoft's submission to the CMA [✂].
[1968] Microsoft's submission to the CMA [✂].
[1969] In doing so, we assume that if AWS or Google had the ability to absorb or counteract the conduct, they would have done so.

6.516 In addition, we found that average usage of Windows Server and SQL Server is higher on Azure than on AWS and GCP, suggesting customers disproportionately choose to deploy their workloads using Windows Server and SQL Server on Azure. This alone cannot determine whether there is a causal link between Microsoft's conduct and relative usage; the disparity may also be explained by other factors such as a more general preference to use Microsoft software on Azure or the use of different software products that are substitutes for Microsoft software on AWS and GCP. We have therefore considered these findings against the background that customers are unlikely to switch from Windows Server and SQL Server (as set out in our assessments of market power), which suggests that there is a higher likelihood that the differences in relative usage are at least partially driven by Microsoft's conduct.

6.517 We place limited weight on 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.

6.518 First, we consider 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.

6.519 Second, we consider that 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.

6.520 Finally, we note that Microsoft's analysis does not consider the potential counterfactual. In particular, it is not clear what the number of large customers that Microsoft would have in the absence of the licensing practices we have identified above.

6.521 We also considered customers' views on using Microsoft software on public cloud and found that customers perceive AWS and GCP to be more expensive than Azure, and some customers consider licensing factors in making their choice of cloud.

6.522 Finally, we consider that Microsoft's licensing practices with respect to Office, Microsoft 365, Windows Desktop and Visual Studio could 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. We note that competition for VDI workloads may be the lower bound for any effect on competition with respect to VDI. This is 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.

Provisional conclusions

6.523 We have provisionally found that Microsoft's licensing practices in relation to Windows Server, and Windows Server and SQL Server together, are harming AWS' and Google's competitive offerings. We also have provisionally found that the impact of Microsoft's licensing practices on VDI workloads contributes to the impact on AWS' and Google's competitive offerings.

Summary of our assessment and provisional conclusions

6.524 We have considered whether Microsoft's licensing practices may partially foreclose its rivals in the supply of cloud services, particularly in competing for customers who are purchasing cloud services that use certain Microsoft software as an input.

6.525 We structured our assessment around three limbs: whether Microsoft has the ability to foreclose its rivals; whether it has the incentive to do so; and whether its conduct has an adverse effect on the competitiveness of its rivals.

Ability

6.526 We assessed Microsoft's ability to harm rivals in cloud services considering its market power upstream (ie in software products)[1970] and the importance of those products as inputs for AWS and Google.[1971]

6.527 We have provisionally found that Microsoft has significant market power in relation to each of Windows Server, SQL Server, Windows 10/11, Visual Studio and its productivity suites. Customers generally said that they were unable or unwilling to switch away from the respective Microsoft software products which means they cannot mitigate against the effect of any licensing practices Microsoft chooses to employ.

6.528 In order to assess the significance of Windows Server and SQL Server to the cost base of rivals, we considered multiple indicators that reflect different possible bundles of cloud services customers purchase and considered the range of values between them. We have found that across all revenue brackets, Windows Server and SQL Server combined account for a material proportion of overall costs and,

[1970] 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.
[1971] Microsoft could only harm the competitiveness of its rivals if the input it supplies plays an important role in shaping downstream competition.

for values of our indicators that are in the middle of the range and further away from our upper and lower bound estimates (and therefore better reflect the significance of Windows Server and SQL Server for customers that make their decisions in a disaggregated way), are very high. We also consider that the evidence supports a provisional finding that Windows Server is an important input on its own.

6.529 Evidence on usage of VDI suggests that it represents a small portion of the overall spend on public cloud services. However, the evidence suggests that Windows 10/11 and the Microsoft productivity suites are a particularly important input to VDI and Visual Studio is an important input 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.

6.530 We also assessed the effect of quality differences between using the relevant Microsoft software products on Azure compared to AWS and Google. We have found that quality differences do not suggest that the input is significantly more important than we found based on the significance to the cost base of rivals.

6.531 Customers are not generally aware of quality differences between using Microsoft software products on Azure compared to rival clouds. However, some customers did identify access to security updates and the availability of Windows multisession capability for VDI workloads as important, and therefore for these customers Windows Server, SQL Server and/or Windows Desktop may be particularly important in terms of affecting the overall quality or attractiveness of AWS' and Google's competitive offerings. Most customers are not aware of quality differences because they have not compared these for Microsoft software products across clouds.

6.532 In summary, we have provisionally found that Microsoft has the ability to partially foreclose its rivals using the relevant Microsoft software products.

Incentive

6.533 We have considered Microsoft's incentive to harm rivals, including its conduct and pricing of its software inputs, as well as how its position in relevant markets may affect its incentives:

(a) Microsoft currently sets a high input price for AWS and Google to pay in order to host Windows Server and SQL Server. This price 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 to be able to use Windows Server on Azure, provided the customer has pre-existing licences for the product (and qualifies for AHB). For SQL Server, the input cost for AWS and Google is higher than Microsoft's PAYG customer-facing price.

(b) Customers are unlikely to switch away from the relevant Microsoft software products to other software on AWS and Google, particularly if they used those software products in a traditional IT environment and would incur switching costs. This would mean that the cost to Microsoft of raising rivals' costs are lower, meaning there is a greater incentive to partially foreclose.

(c) Microsoft, through Azure, has a high share of supply in both IaaS and PaaS, which are highly concentrated markets, and 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 the Google, and vice versa.

(d) IaaS and PaaS are high growth markets, and IaaS is characterised by strong economies of scale. As such, the potential for Microsoft to win customers in IaaS and PaaS and position itself to capture future growth and lower its infrastructure costs may act as an additional incentive to partially foreclose rivals.

6.534 In summary, we have provisionally found that Microsoft has the incentive to harm rivals through conduct relating to the relevant Microsoft software products.

Effect

6.535 We have considered whether our assessment is consistent with Microsoft's licensing practices being likely to have the effect of reducing competition in cloud services markets, in particular by materially disadvantaging its rivals.

6.536 We have evidence that is consistent with AWS' and Google's competitive offerings being worsened by Microsoft's licensing practices in relation to Windows Server and SQL Server:

(a) AWS and Google pass through input costs, and customers perceive AWS and Google to be more expensive. 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 qualify for AHB, AWS has a mostly similar list price for certain instances of Windows Server VMs as Azure. Google's list prices on these instances are mostly at least 5% higher than Microsoft's.

(b) AWS' and Google's margins are relatively low on Windows Server and SQL Server workloads. Low margins are not inconsistent with a situation where

AWS and Google are passing through high input costs to some extent while absorbing some in their margin. We do not consider that positive margins are inconsistent with partial foreclosure.

(c) Windows Server and SQL Server are used disproportionally on Azure compared to on AWS and GCP. Some customers said that their choice of cloud was influenced by licensing factors, and particularly price factors.

6.537 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.

6.538 In relation to Windows 10/11 and Microsoft 365, Microsoft does not make these available to AWS and Google through their respective SPLAs, and customers cannot make use of the BYOL route to deploy existing licences for this software on AWS' or Google's public cloud (with the exception of certain Microsoft 365 licences on Amazon Workspaces). [✂].

6.539 This means 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, we consider that Microsoft's conduct relating to these software products to contributes to an effect on rivals' competitive offerings.

6.540 AWS and Google together are the most significant competitors to Microsoft by share of supply, both in IaaS and PaaS markets. As such, we consider that the impact of Microsoft's licensing practices on AWS' and Google's competitive offerings is likely to affect a significant portion of cloud services.

6.541 In summary, we have provisionally found that Microsoft's conduct is harming competition in cloud services.

6.542 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. We have received very limited evidence to suggest that non-Listed Providers are benefiting in terms of growth or scale, and note that their smaller market shares suggest that they would capture a much lower proportion of customers than Microsoft.

Provisional conclusion

6.543 Our provisional conclusion is that all three conditions for partial foreclosure of AWS and Google are satisfied, and that Microsoft's conduct is harming competition in cloud services.

7. Committed spend agreements

- This chapter sets out our assessment of Committed Spend Agreements or Discounts (CSAs or CSDs). These 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.
- We assess whether CSAs may lead to harm to competition by acting as a potential commercial barrier to switching and multi-cloud. We focus on AWS and Microsoft based on their position as the two largest providers of cloud services by some distance.
- We assess the extent to which CSAs could lead to harming competition by either: (i) 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 (ii) reducing the incentive of rival suppliers to compete for each other's existing customers, leading to softening of competition between those suppliers.
- We have found that 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 has indicated that the existence of AWS' and Microsoft's CSAs, as currently structured, do not adversely affect rivals' ability or incentive to compete against AWS and Microsoft.
- In view of the above, we have provisionally found that 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, this does not mean that harm to competition might not materialise in the future. In particular, we have considered that the link between sticky and contestable demand in AWS' and Microsoft's CSAs would likely harm competition if: (i) the market matures such that the share of sticky demand increases significantly; and/or (ii) 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

7.1 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.[1972]

7.2 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.

7.3 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'.

7.4 We first set out our conceptual framework for conducting our assessment, and then move to our analysis.

7.5 We focus mainly on AWS and Microsoft. We do so based on their position as the two largest providers of cloud services by some distance. 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 and other smaller cloud providers.

Conceptual framework and analysis structure

7.6 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.

7.7 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 don't meet those conditions.

[1972] In this report, we use the terms CSA and CSD interchangeably to mean the same type of agreement.

7.8 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.

7.9 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;[1973] 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.[1974]

7.10 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

[1973] The scenario we have considered is AWS' and Microsoft's CSAs reducing the ability of smaller rival suppliers to profitably compete for the contestable demand.
[1974] The scenarios we have 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.

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.

7.11 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.

7.12 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.

7.13 As shown in Figure 7.1 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 7.1: Discount that the rival must apply to win units of contestable demand

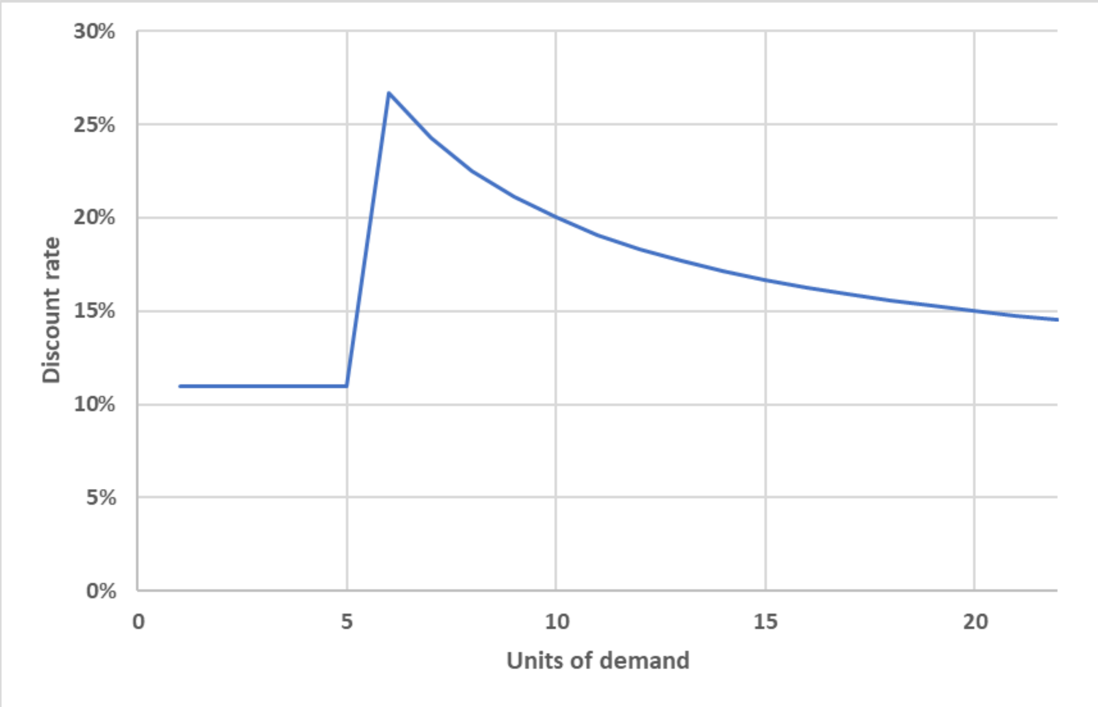

Source: CMA analysis.

7.14 As Figure 7.1 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.

7.15 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.

7.16 Figure 7.2 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 7.2: Total profit to rival for different units of contestable demand it wins

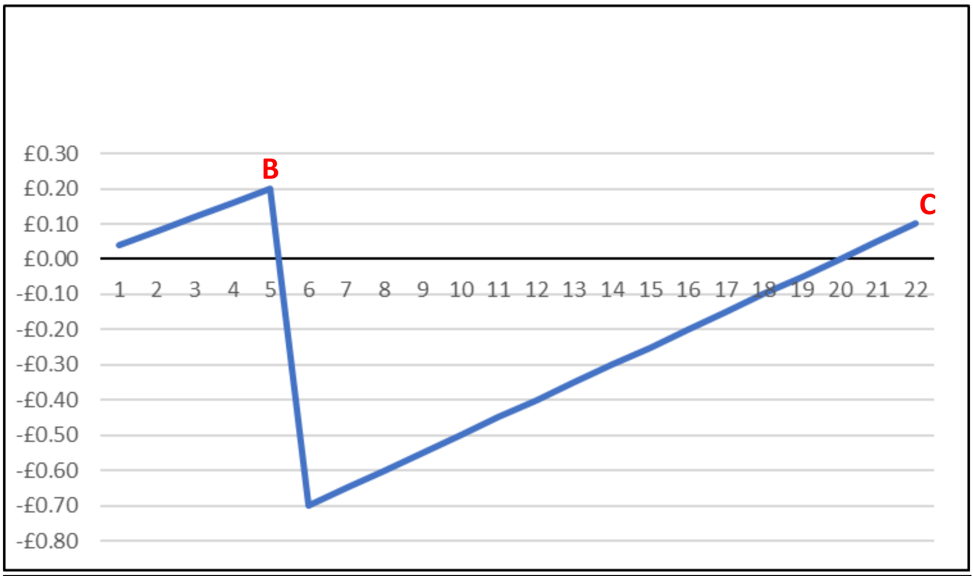

Source: CMA analysis.

7.17 We note the above is a stylised example and is not necessarily representative of the CSAs offered by AWS or Microsoft, nor does it intend to be reflective of the actual margins of rivals.

7.18 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. 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) First, we have considered the prevalence of CSAs. This provides an indication of the size of the segment that could be affected by these agreements.

(b) Second, we have considered 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) Third, we have considered 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.[1975]

(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) Finally, we have also considered whether there are any potential benefits that may arise from CSAs even if they may harm competition. For example, some

[1975] 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*1,030)/30 ~ 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*130)/30 ~ 0.09).

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.

7.19 We note that cloud providers offer several other discount programmes besides CSDs. These range from discounts on specific services based on booking cloud 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).[1976]

7.20 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.[1977] 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

7.21 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.

7.22 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

7.23 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.[1978]

(b) IBM submitted that CSAs can reinforce the technical barriers to multi-cloud and present a challenge from customers' standpoint. Nevertheless, IBM also

[1976] We consider cloud credits in Chapter 4.
[1977] We discuss cloud credits targeted to startups in Chapter 4.
[1978] Former UK Cloud employees' response to working papers and updated issues statement, page 3.

submitted that CSAs warrant a nuanced approach as they also allow customers to benefit from lower prices and stability/predictability of spend.[1979]

(c) 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.[1980]

(d) An ISV said that CSDs incentivise customers to buy third-party services in the larger cloud providers' marketplaces, thereby increasing these cloud providers' advantages.[1981]

7.24 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.[1982]

(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.[1983]

(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.[1984]

(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.[1985]

(e) OkTik (an ISV) submitted that it relies upon Microsoft's MACC,[1986] for sales of its product and could be 'heavily impacted' if MACCs were ever to be restricted in the UK.[1987]

Industry bodies

7.25 We have received submissions from ACT/The App Association, Computer and Communications Industry Association (CCIA), and The Startup Coalition that

[1979] IBM's response to the Committed spend agreements working paper, page 1.
[1980] Company A's response to the Issues Statement dated 17 October, paragraph 2.1.
[1981] [✂] submission to the CMA [✂].
[1982] AWS' response to the Issues Statement dated 17 October 2023, paragraphs 30-32.
[1983] Microsoft's response to the Issues Statement 17 October 2023, section 6.2.
[1984] Google's response to the Issues Statement dated 17 October 2023, paragraph 32.
[1985] Oracle's response to the Issues Statement dated 17 October 2023, page 5.
[1986] MACC stands for 'Microsoft Azure Consumption Commitment' and is Microsoft's main CSD programme.
[1987] OkTik's response to the Committed spend agreements working paper.

considered CSDs to be beneficial to customers and investment in cloud services.[1988]

7.26 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,[1989] AWS and Google are members of the CCIA,[1990] and Google is a supporter of the Startup Coalition.[1991]

Customers

7.27 We have also received submissions from a couple of customers who provided comments on 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'.[1992]

(b) A banking provider submitted that overall, 'regulatory intervention should be carefully evaluated to avoid jeopardizing the broader business and economic advantages associated with cloud adoption and the commercial deals that have been struck to date absent clearly abusive terms and conditions'.[1993]

7.28 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

7.29 We have also had submissions from two academics, Dr George R Barker and R. Parisi, who were generally positive about CSDs.[1994]

[1988] 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.
[1989] ACT/The App Association, accessed 19 April 2024.
[1990] CCIA Members, accessed 19 April 2024.
[1991] Startup Coalition, accessed 19 April 2024.
[1992] [✂] submission to the CMA [✂].
[1993] Banking Provider 1's response to working papers and updated issues statement, pages 1-2.
[1994] 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 74. We note that Dr George Barker is a member of the Oxford Cross Disciplinary Machine Learning

Our assessment

7.30 We have 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

7.31 We have 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.

7.32 In relation to the prevalence of CSDs, we have set out below:

(a) cloud providers' views on the prevalence of CSDs; and

(b) evidence from data on AWS' and Microsoft's CSDs.

Cloud providers' views

7.33 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 [✂]% of total revenues from cloud infrastructure services in the UK and that this is so low that it is implausible that any rivals could be foreclosed.[1995]

7.34 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.[1996] 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.[1997]

Research Cluster (OXML). The OXML is supported by Microsoft (see page 1). R. Parisi, The Cloud Services Markets' Competitive Landscape: A contribution to the Competition and Markets Authority, page 14.
[1995] AWS' submission to the CMA [✂].
[1996] AWS' submission to the CMA [✂].
[1997] AWS' submission to CMA dated 31 July 2024, 'AWS Response to the CMA's Committed Spend Agreements Working Paper', paragraph 21.

Our assessment

7.35 We have considered that AWS' estimates of market coverage for its CSDs ([✂]% of revenues from IaaS and PaaS in the UK) are in line with our own estimates and that this is a material level of coverage.

7.36 We have also considered that 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) We have considered our assessment of market power within our framework for assessing CSDs to be 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').[1998] In determining whether there is an AEC, the CMA will consider a range of factors, which may include the levels of firms' market power.[1999]

(b) Additionally, market power may also be exerted by a single firm, or by several firms acting independently (ie unilateral market power).[2000] 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 higher market share of that provider) and as a share of total demand for that provider's services. AWS' and Microsoft's CSDs can have characteristics that lead to harm to competition in the market irrespective of whether AWS and Microsoft engage in 'tacit collusion'.

(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.

7.37 In addition, AWS submitted that its estimate does not include UK customers whose non-UK parent entity holds a CSD with AWS and do not include [✂].[2001] Therefore, we have considered that some relevant CSD agreements may be

[1998] See Our task chapter, section "our statutory duty".
[1999] 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.
[2000] See CC3 (Revised), paragraph 178.
[2001] AWS' submission to the CMA [✂].

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 revenues from these customers would be counted in AWS' total UK revenues (the denominator) but would be excluded from the UK revenues that are subject to CSDs (the numerator).

Data analysis

7.38 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).

7.39 We have calculated these metrics using data provided by AWS and Microsoft.[2002] 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.

7.40 We have found that the proportion of total annual cloud spend accounted for by customers with a CSD in 2022 was [✂]% for AWS and [✂]% for Microsoft.

7.41 Both of those figures have grown substantially over the years – from [✂]% in 2018 to [✂]% in 2022 for AWS, and from [✂]% in 2019 to [✂]% in 2022 for Microsoft.

7.42 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.

7.43 These figures represented a very small proportion of the providers' total number of UK customers in 2022 ([✂]% for AWS and [✂]% for Microsoft). This is mainly driven by the large number of smaller customers (ie those spending less than $500,000 yearly),[2003] that accounted for around [✂]% of total customer base for AWS and for around [✂]% for Microsoft.

7.44 When looking at customers in higher spend bands, the proportion of customers with CSDs is much higher. For example, in 2022, [✂]% of customers spending

[2002] AWS' and Microsoft's responses to the CMA's information requests [✂].
[2003] As set out in 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.

more than $500,000 per year with AWS had a CSD in place. The equivalent figure was [✂]% for Microsoft.

Impact of CSDs on customers' choices

7.45 We have 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.

7.46 In order to examine the impact of CSDs on customers' choices, we have considered:

(a) cloud providers' views on the impact of CSDs on customers' choices; and

(b) customers' views on the impact of CSDs on their allocation of workloads.

Cloud providers' views

7.47 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.[2004]

7.48 AWS submitted that customers responding to price differentials is evidence of effective price competition.[2005]

7.49 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.[2006]

Our assessment

7.50 With respect to these submissions, we have considered the following:

(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 to assess this is to ask customers directly the extent to which they think CSDs impact their allocation choices, which we did in our

[2004] Microsoft's response to the Competitive Landscape, Committed Spend Agreements and Egress Fees Working Papers, paragraph 69.
[2005] AWS' submission to the CMA [✂].
[2006] Google's response to the Committed Spend Agreement Working Paper, paragraph 6.

customer questionnaire, 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 have considered the key question to be whether CSDs reduce the ability and incentive of rivals to compete.

Customer evidence

7.51 We have 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.

7.52 All but one customer that answered these questions had either AWS or Microsoft as their main provider.

New workloads

7.53 We have 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’.

7.54 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.[2007]

7.55 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.[2008]

Existing workloads

7.56 We have also asked large customers about the effects of CSDs on the allocation of existing workloads. We have 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

[2007] Responses to the CMA’s information requests [✂].
[2008] Responses to the CMA’s information requests [✂].

switch any such workloads. The customers had to choose one option on a scale from 'very unlikely' to 'very likely'.

7.57 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.[2009]

7.58 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.[2010]

Evidence from the Jigsaw report

7.59 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;[2011] 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 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.[2012]

Impact of CSDs on competition

7.60 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

[2009] Responses to the CMA's information requests [✂].
[2010] Responses to the CMA's information requests [✂].
[2011] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 6.2.
[2012] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraphs 6.2.8-6.2.9.

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 providers' views

7.61 Microsoft said that over time, the value of discounts in percentage terms that Azure customers received has [✂]. Microsoft also said that its regression analysis of UK CSD contracts shows that [✂].[2013]

7.62 AWS submitted the following points, which it said are inconsistent with the claim of customer lock-in or rivals' foreclosure from AWS' CSDs:

(a) [✂].[2014] [2015]

(b) The structure and level of AWS' discounts is such that efficient rivals can compete on incremental demand.[2016]

(c) There is a large body of evidence of entry and growth by many rivals of all sizes.[2017]

(d) Market outcomes (repeated innovation, higher quality of services, decreasing prices, decreasing margins, and efficiencies passed on to customers) are consistent with effective competition.[2018]

7.63 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.[2019]

7.64 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

[2013] Microsoft's submissions to the CMA [✂].
[2014] AWS' submission to the CMA [✂].
[2015] AWS' submission to the CMA [✂].
[2016] AWS' submission to the CMA [✂].See also AWS' submission to the CMA [✂].
[2017] AWS' submission to the CMA [✂].
[2018] AWS' submission to the CMA [✂].
[2019] AWS' submission to the CMA [✂].

products, including its 'must-have' Office software productivity suite and client PC and work group server operating systems.[2020]

7.65 In response to that point, Microsoft said that it does not offer bundled discounts across Azure and its non-Azure services.[2021] 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.[2022]

7.66 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.[2023]

7.67 We also asked 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 will not generally impact the scope of the workloads for which cloud providers are trying to compete.[2024]

Our response to these providers' views

7.68 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 have considered 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.

7.69 Taking their analyses at face value, we have considered that there is another possible interpretation of their results, that is – although no weakening and/or marginalisation of rivals (and by extension, any resulting impact in terms of worse

[2020] Google's response to the Committed Spend Discounts Working Paper, paragraph 17.
[2021] Microsoft's summary of hearing on Tuesday 16 July 2024, paragraph 24.
[2022] Transcript of hearing with Microsoft [✂].
[2023] [✂] response to the CMA's information request [✂].
[2024] Responses to the CMA's information requests [✂].

price/discount deals for customers) has happened yet, this might be the result of a market that is still growing significantly. That would translate to a large injection of contestable demand in the market each year, which means a lower share of demand for AWS' and Microsoft's services that is sticky so long as this growth trend continues. However, as the market matures, the increase in contestable demand might slow significantly, and the associated share of demand for AWS' and Microsoft's services that is sticky might increase. At that point, and all other things being equal, the impact of AWS' and Microsoft's CSDs on smaller competitors might increase, and these competitors might be weakened and marginalised to a point that AWS and Microsoft would then be able to increase prices and/or reduce discounts.[2025]

7.70 We note that this could be the case even if AWS and Microsoft were still strongly competing against each other. This is because the reduction of a competitive constraint from some cloud providers in an oligopolistic market, eg due to these other cloud providers being weakened or marginalised, may affect competition overall, and in turn, market outcomes (eg prices, quality) even if competition between the AWS and Microsoft still exists.

7.71 Relatedly, 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.

7.72 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.

7.73 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.

7.74 Irrespective of whether there has been some entry and growth by rivals, our analysis has indicated that there is concentration in the supply of public cloud infrastructure services in the UK and that concentration has increased between 2019 and 2022.

[2025] A customer identified a concern along these lines in its response to the Committed spend agreement working paper ([✂] submission to the CMA [✂].).

7.75 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 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.

7.76 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

7.77 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.

7.78 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.

7.79 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. [✂].

7.80 Below, we have first explained our approach to the analysis. We have then discussed one of the key assumptions behind the analysis – that customers’ existing demand is sticky, and finally presented the results of our analysis.

Our approach

7.81 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.[2026]

7.82 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 have 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. If AWS' margin is larger than the discount that the rival would have to apply, then we would have inferred that a rival with costs similar to AWS would be able to compete for that incremental demand.

7.83 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'.[2027] In the following paragraphs, we have adopted AWS' terminology on margins.

7.84 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] [✂] AWS' average of [✂]% – this translates to a net margin of around [✂]%.[2028] [✂].[2029]

(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

[2026] AWS' submissions to the CMA [✂]; AWS' response to the CMA's information request [✂].
[2027] AWS' submissions to the CMA [✂]; AWS' response to the CMA's information request [✂].
[2028] [✂]. AWS' responses to the CMA's information requests [✂].
[2029] See Chapter 3, section on 'Our analysis of profitability', and Appendix K, sections on 'Cost of capital' and 'Our profitability analysis'. See also CC3 (Revised), paragraph 116.

at whether rivals have the incentives to compete for the whole incremental demand or just for the uncommitted incremental demand. [✂]. 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.

7.85 In practice, our analysis works in the following way. Starting from a data set with [✂] 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%, [✂]%, [✂]% and [✂]% of AWS' average net margin of [✂]%;

(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%.

7.86 This meant that each scenario looked like a version of the chart presented in Figure 7.2 above.

7.87 This resulted in [✂] scenarios.[2030] 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.

[2030] [✂].*[✂] = [✂]; where [✂]. is the number of AWS' CSAs in the starting data set, [✂] is the number of 1 percentage points increments in the range 10% - 150%, [✂] is the number of values we consider for the proportion of the incremental demand left uncommitted, and [✂] is the number of values we consider for the rival's net margin relative to AWS' net margin.

Share of demand for the incumbent's services that is sticky

7.88 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 the incumbent's services that is sticky, the higher the ability and incentive or rivals to compete for the incremental demand.

7.89 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.[2031]

7.90 In this section, we assessed our default assumption more in detail by looking at evidence on the share of customers' existing demand that is sticky.

Cloud providers' views

7.91 AWS submitted that its services are not a 'must-have' (ie unique to AWS), and therefore cannot be leveraged to foreclose competitors.[2032]

7.92 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.[2033]

7.93 Microsoft submitted that there will be constant and significant growth in contestable demand. According to Microsoft, this means that, all things being

[2031] 80/(80+20+50) ~ 100/(100+100) = 50%.
[2032] AWS' submission to the CMA [✂].
[2033] AWS' submission to the CMA [✂].

equal, the alleged 'contestable' demand should be significantly larger than any alleged 'sticky' demand at any given time.[2034]

7.94 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'.[2035]

Our assessment

7.95 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

7.96 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.[2036]

(b) We also found evidence that the additional costs for customers brought about by the presence of egress fees act as barrier to switching.[2037]

(c) We found evidence on barriers to switching from Microsoft's software licensing practices in Chapter 6.[2038]

7.97 The Jigsaw report noted that a few of the customers interviewed chose to switch between cloud providers.[2039] 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.[2040] The report also noted that switching cloud providers is seen as the equivalent of moving other kinds of infrastructure, such as 'moving house' or

[2034] Microsoft, Response to the competitive landscape, committed spend agreements and egress fees working papers, paragraph 62a.
[2035] Microsoft, Response to the competitive landscape, committed spend agreements and egress fees working papers, paragraph 70.
[2036] See Chapter 5, Technical barriers.
[2037] See Chapter 5, Egress fees.
[2038] See Chapter 6.
[2039] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 1.3.14.
[2040] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 1.3.16.

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 provider(s), or a sense of 'lock in', is a factor across all cloud providers as change brings cost and risk'.[2041]

7.98 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

7.99 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.

7.100 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 have considered this evidence in the round and alongside other evidence received and discussed above.

7.101 Based on our analysis, there was significant variation in customer responses to this question.[2042] 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.[2043] A few customers said that they would not be willing to switch a proportion of their AWS demand, ranging from 21-66%.[2044] Finally, a few customers said that they were willing to switch all of their AWS demand.[2045]

7.102 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.[2046] 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%.[2047] Finally, some customers said they were willing to switch all of their Microsoft demand.[2048]

[2041] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraph 1.3.13.
[2042] Responses to the CMA's information requests [✄].
[2043] Responses to the CMA's information requests [✄].
[2044] Responses to the CMA's information requests [✄].
[2045] Responses to the CMA's information requests [✄].
[2046] Responses to the CMA's information requests [✄].
[2047] Responses to the CMA's information requests [✄].
[2048] Responses to the CMA's information requests [✄].

7.103 In relation to the cloud providers' views set out above, we have considered that:

(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, we consider the 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 do 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

7.104 Generating a chart like the one in Figure 7.2 for a specific AWS' CSD yielded below. Here, we took a CSD with commitment value equal to a [representative] commitment value for all first-time AWS' CSDs in our data set.[2049] 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.[2050]

[2049] We identified [✂] first-time CSAs out of total [✂] in AWS' data set used for our analysis.
[2050] More details on our assumptions and calculations can be found in Appendix U: CSA – quantitative analysis.

Figure 7.3: Total profit to rival for different units of contestable demand it wins for a representative ([✂]) AWS' CSD

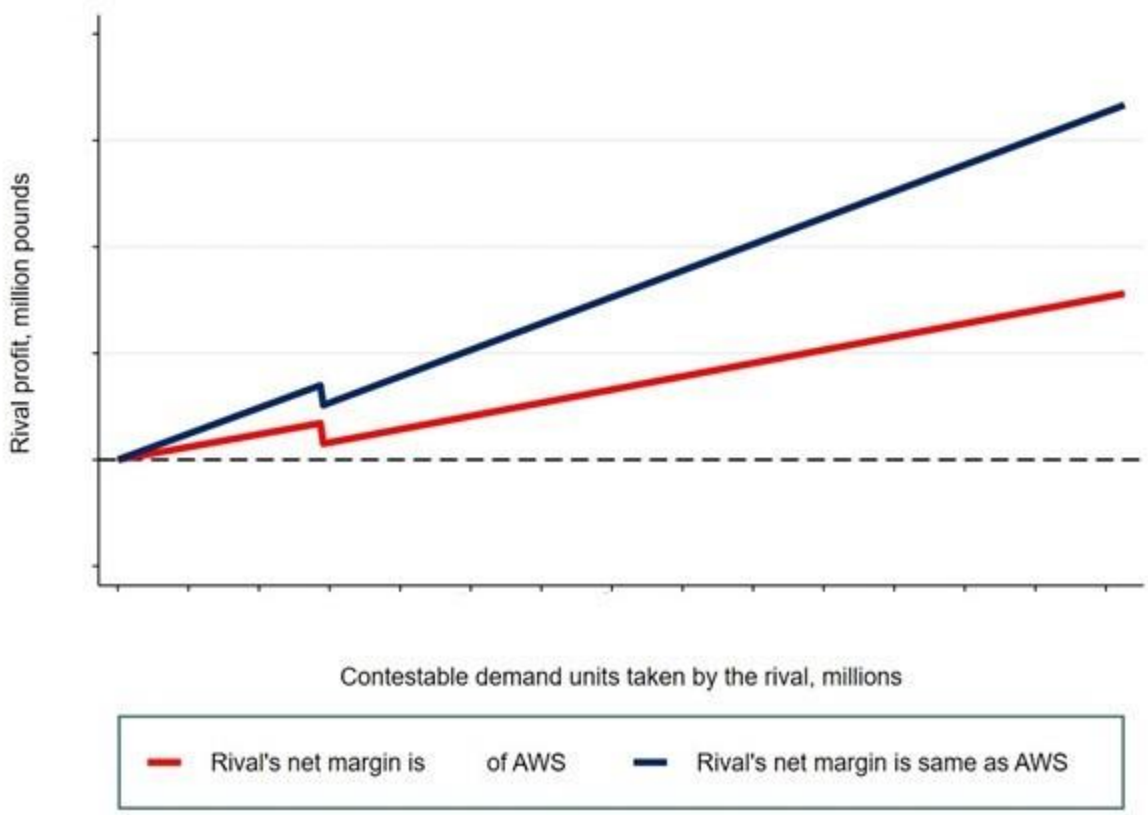

Source: CMA analysis.

7.105 The figure 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.

7.106 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's 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 in a mature market).

7.107 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).

7.108 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.[2051] 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.

7.109 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. This could be the case if we expected the cloud infrastructure services markets to 'mature' in the foreseeable future. As occurrences of assessment failures tend to be more likely when looking at low demand increments, all other things being equal, CSDs would

[2051] 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.

be more likely to harm competition if we expected the market to mature in the foreseeable future.

7.110 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, [✂].[2052]

7.111 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]

7.112 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 when the cloud infrastructure services markets mature.

Qualitative assessment

7.113 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.

7.114 We have complemented that assessment by looking at the proportion of customer demand covered by the commitment and the length of CSDs. We have considered that, despite not featuring directly in our quantitative analysis, these were also relevant factors that could lead CSDs to harm competition.

7.115 Overall, we have not found that 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 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.

[2052] Responses to the CMA's information requests [✂].

Potential benefits of CSDs

7.116 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.

7.117 Therefore, in this section, we have considered whether there are any potential benefits arising from efficiencies on competition associated with CSDs (otherwise known as 'rivalry enhancing efficiencies'). This includes 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.

7.118 In this section we have set out submissions from cloud providers on potential benefits and our assessment of these submissions.

Stakeholders' views

Benefits accruing to customers

7.119 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.[2053]

7.120 Microsoft and Google also stated that removing or restricting CSDs would reduce benefits accruing to customers.[2054]

Benefits accruing to cloud providers and incentives to invest

7.121 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.

[2053] 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.
[2054] Microsoft's submission to the CMA [✂]. Microsoft's response to the Competitive Landscape, Committed Spend Agreements and Egress Fees Working Papers, paragraph 59. Microsoft's submission to the CMA [✂]. Google, response to the Committed Spend Agreements Working Paper, paragraph 11.

7.122 AWS and Microsoft said that CSDs are one of several factors they consider when deciding the timing of investment in infrastructure.[2055]

7.123 A cloud provider said that its decisions to invest in infrastructure are not directly linked to committed volumes.[2056] 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.[2057]

Our assessment

7.124 Based on the evidence above, we have found that:

(a) in general, the provision of discounts can be beneficial to those customers who qualify for it. 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 decision short term, such benefits would have to be balanced against any potential longer term harm to competition; and

(d) there are other alternative structures which cloud providers could use to offer discounts to customers that would continue to allow customers to benefit from discounts but would not risk give rise to the same concerns about harm.

Provisional conclusions

7.125 We have provisionally found that AWS' and Microsoft's CSAs cover a significant portion of the cloud infrastructure services markets and influence customers' choices in relation to workload allocation.

7.126 However, in relation to the impact of CSAs on competition, our quantitative analysis has indicated that the existence of AWS' and Microsoft's CSAs, as currently structured, do not adversely affect rivals' ability or incentive to compete against AWS and Microsoft.

2055 AWS's and Microsoft's responses to the CMA's information requests [✂].
2056 [✂] response to the CMA's information request [✂].
2057 [✂] response to the CMA's information request [✂].

7.127 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 leading to harm to competition and/or customers.

7.128 In view of the above, we have provisionally found that 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.

7.129 Although we have not identified competition concerns, this does not mean that harm to competition might not materialise in the future. In particular, we have considered that the link between sticky and contestable demand in AWS' and Microsoft's CSAs may likely to lead to harm to competition if:

(a) the market matures such that 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.

8. Provisional decision on competition

8.1 This chapter sets out our provisional decision on 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 United Kingdom or a part of the United Kingdom.[2058]

Introduction

8.2 Cloud services are key inputs to many businesses and organisations across the UK economy with £9 billion spent on them in the UK in 2023. It is therefore vital for economic growth in the UK that this sector works well for its customers.

8.3 We have sought the opinions of a wide range of UK cloud customers and they have told us that cloud services bring them many benefits compared to on-premises IT: they are reliable, scalable, easy to use and maintain, providers support them, and they are innovative. Cloud services sometimes cost less for customers than their alternatives.[2059] However, our research found that customers are also aware of the potential high costs and risks of switching or using multiple clouds and so have some sense of dependence or lock in with their providers.[2060]

Our view of well-functioning cloud services markets

8.4 In identifying features or combination of features of the market that prevent, restrict or distort competition thereby giving rise to an AEC(s), we have to find a benchmark against which to determine how the market may be judged to be performing.[2061]

8.5 In the absence of a statutory benchmark, we use the benchmark of 'a well-functioning market' [2062] 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.

8.6 But 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).[2063]

[2058] Section 134(1) EA02.
[2059] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), page 25.
[2060] Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024), paragraphs 1.3.12-1.3.16.
[2061] CC3 (Revised), Guidelines for market investigations: Their role, procedures, assessment and remedies, paragraph 320.
[2062] 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.
[2063] CC3 (Revised), paragraphs 154 and 320.

8.7 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).

8.8 In our view, well-functioning cloud services markets are those in which customers do not become locked into their providers but rather are able to switch cloud providers 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 providers being more able to compete effectively for all workloads and across the range of cloud services, including those services that use Microsoft software.

8.9 If customers were able to switch and multi-cloud more easily, and had more choice of provider, they could choose the best products and services for their needs, whether by switching provider, by combining services best suited for their needs across providers through multi-cloud architectures, or both. Greater entry and expansion by cloud providers[2064] 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.[2065]

8.10 Therefore, in a well-functioning market we would expect there to be more switching and multi-cloud, and choice of provider, than we observe at present which may ultimately be expected to impact the price and quality of the cloud services and choice of cloud provider available to customers.

Market definition

8.11 We have provisionally found that there are separate markets for IaaS[2066] and PaaS: these are the relevant cloud services markets. Neither software as a Service (SaaS) nor traditional 'on-premise' IT or private cloud form part of the IaaS and PaaS markets, because most customers do not see them as close substitutes.

8.12 We have provisionally found that the geographic scope of the IaaS and PaaS markets is Europe-wide (that is, UK and EEA).

[2064] 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.
[2065] CC3 (Revised), paragraphs 205 and 207.
[2066] We have defined the market for IaaS services in aggregate with the exception of cloud infrastructure services based on accelerated compute infrastructure. For the purposes of this investigation, we refer to this market as IaaS.

Market concentration

8.13 We have provisionally found that the cloud services markets are highly concentrated and that this is likely to endure. All metrics on shares of supply in IaaS, PaaS and combined, show that AWS and Microsoft are the largest providers with Google, the third largest provider, much smaller according to our measures.

8.14 The IaaS market is particularly highly concentrated. While Google's market share has increased in recent years, its share remains significantly smaller than that of Microsoft or AWS whose combined market share remains high and stable. The lower and/or declining shares of other smaller UK providers indicates that they are likely to remain considerably smaller than AWS or Microsoft.

8.15 The PaaS market is also concentrated but less so than IaaS and with a longer tail of small providers. However, we note that although AWS, Microsoft and Google offer both IaaS and PaaS, most PaaS providers do not also offer IaaS, so any sale they make is also likely to be a sale for an IaaS provider, as IaaS is an input to ISVs' services. ISVs as PaaS-only providers therefore place limited constraint on firms supplying both IaaS and PaaS.

8.16 Across both markets, there are also a few smaller 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 AWS, Microsoft and Google.

Impact of AI on cloud services

8.17 We have provisionally found that partnerships between larger cloud providers and FM developers are widespread and that they may play an important role in shaping the competitive conditions in the supply of accelerated compute to FM developers and in the supply of access to FMs to other customers.

8.18 Although we note that AWS, Microsoft and Google each has a strong position in the supply of IaaS based on accelerated compute, based on the evidence we have considered, we have provisionally found that there is currently no significant direct impact from IaaS based on accelerated compute on competition in IaaS based on standard compute.

8.19 Providing access to FMs has emerged as a potential future driver of customers' choice of cloud service provider and cloud providers have differentiated strengths in this regard, but customers do not think that the current competitive landscape in respect of the supply of access of FMs is fixed. Based on the evidence we have considered, we have provisionally found that access to FMs by customers of cloud services is not currently a strong driver of customer choice and that the extent to which it will become a driver of choice in the future is uncertain.

Market outcomes

8.20 We have provisionally found that AWS and Microsoft have been generating sustained returns from their cloud services substantially above their cost of capital in cloud services for a number of years.

8.21 Prices paid by cloud customers for different cloud services have moved in different directions for different services, with some services and products increasing in price over time, while others are falling. Customers say that cloud services offer both quality and innovation to them. However we consider that a more competitive market would have sustained better market outcomes, including more consistently competitive prices, as well as further improvements in quality and innovation.

Barriers to entry and expansion

8.22 We have provisionally found substantial barriers to entry and expansion in the provision of cloud services, in particular for IaaS.

8.23 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, whereby the larger providers have 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. This disincentivises IaaS market entry and expansion.

8.24 Furthermore, given the scale of investment and expansion that large cloud providers have made to date in IaaS, any new entrant (or company seeking to expand) would need to invest substantially more than the large existing suppliers in order to close the gap in a timely way. The levels of investment that AWS and Microsoft are expecting to make in the coming years may raise these barriers even higher.

8.25 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 and expansion.

8.26 The wide product portfolios of the larger cloud providers 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. Customers also place importance on their provider’s reputation and this contributes to overall barriers to entry and expansion in cloud services.

Barriers to switching and multi-cloud

8.27 We have provisionally found that there are substantial technical and commercial barriers facing customers who wish to switch or use multiple clouds.

8.28 Some customers can and do successfully multi-cloud but we have found that technical barriers to multi-cloud negatively affect many customers' ability to use and integrate multiple clouds. This limits customers' ability to exercise choice of cloud provider.

8.29 Whilst customers are able to switch if they are willing to expend all the required resources to do so, we have found that the perceived costs currently outweigh the benefits for many.

8.30 Technical differentiation of features and interfaces in core and ancillary cloud services means that customers cannot easily compare or substitute products without technical effort. This harms customers' ability to multi-cloud and/or switch clouds. Further disincentives to multi-cloud include latency (the time it takes to transfer data between public clouds), a lack of transferable skills across clouds and insufficient transparency of cloud service features.

8.31 There are some mitigations to these technical barriers from customers and cloud providers, but these only mitigate the technical barriers to multi-cloud and switching experienced by customers or their effect on competition to a limited extent.

8.32 Egress fees that customers pay for transferring data from one cloud to another are another barrier faced by customers who may wish to switch and/or multi-cloud. These costs can restrict customers from exercising effective choice and from responding to attractive offers, leading to weaker competition between providers.

8.33 We have not seen sufficient evidence that these fees fund investment and innovation, nor is it clear that they result in passed-on cost savings to customers or deter inefficient egress usage by them.

8.34 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).

8.35 Technical and commercial barriers enable cloud providers to exercise some degree of power over their existing customer base, which reduces their incentives to compete for rivals' customers on price, service quality, and/or innovation.

8.36 We also consider that technical and commercial barriers, by making it harder for cloud providers to win customers from their rivals, deter entry and expansion and, as such, exacerbate the barriers to entry and expansion arising in these markets.

8.37 Finally, these barriers contribute to AWS and Microsoft being able to maintain high and stable market shares.

Microsoft licensing practices

8.38 We have considered whether Microsoft's software licensing practices partially foreclose its rivals in the supply of cloud services, particularly in competing for customers who purchase cloud services that use Microsoft software as an input.[2067]

8.39 We have provisionally found that Microsoft has the ability to harm its rivals: it 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 unable or unwilling to switch away from these products, there are limited alternatives and/or Microsoft has a moderate to high market share. In addition, Windows Server and SQL Server are a significant input to customers' cloud services for a range of customers. As such Microsoft has the ability to partially foreclose rivals who provide services to customers using the relevant Microsoft software products.

8.40 We have provisionally found that customers are not aware of quality differences between using the Microsoft software on Azure, compared to rival clouds, but they are aware of product differences in terms of access to security updates and other features which are important to them. Microsoft's productivity suites and Windows 10/11 are also a particularly important input to virtual desktop infrastructure which, while a small part of cloud services, is important to some customers.

8.41 In relation to Microsoft's incentive to harm its rivals in cloud services, we have considered Microsoft's conduct as this provides an indication of the incentives it faces because its actions are likely aligned with those incentives and pricing of its software inputs, as well as how its position in relevant markets may affect its incentives. We have provisionally found that Microsoft currently sets a high input price for AWS and Google to pay in order to host Windows Server and SQL Server in the UK. This cost has increased substantially since 2018 and is higher than the retail price that Microsoft charges its own cloud customers.

8.42 Microsoft, through Azure, has a significant market share in both IaaS and PaaS, and is a strong constraint on AWS and Google. This means that customers that

[2067] 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.

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.

8.43 In addition, customers are unlikely to switch away from the relevant Microsoft software products to other software on AWS and Google, particularly if they are coming from using those software products in a traditional IT environment and would incur switching costs. This would mean that the cost to Microsoft of raising rivals' costs is lower, meaning there is a greater incentive to partially foreclose.

8.44 We have considered whether Microsoft's licensing practices have the effect of reducing competition in cloud services markets.

8.45 For Windows Server and SQL Server, we have found evidence consistent with AWS' and Google's competitive offerings being worsened by Microsoft's licensing practices. Evidence from both AWS and Google is consistent with high input costs resulting in them charging a higher price and having a less competitive offering. Customers also perceive AWS and Google to be more expensive. 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 qualify for AHB, AWS has a mostly similar list price for certain instances of Windows Server VMs as Azure. Google's list prices on these instances are mostly at least 5% higher than Azure's. The evidence also shows Windows Server and SQL Server are used disproportionally more on Azure compared to on AWS and GCP.

8.46 In relation to Windows 10/11, Visual Studio, and the productivity suites, AWS' and Google's competitive offerings are directly affected by Microsoft's practices in relation to these software products. Microsoft does not make these available to AWS and Google through their respective SPLAs, and customers cannot make use of the BYOL route to deploy existing licences for this software on AWS' or Google's public cloud (with the exception of certain Microsoft 365 licences on Amazon Workspaces). These products are used via VDI in the cloud and, while VDI represents a relatively small proportion of usage of all cloud services, we consider that it is sufficiently material for Microsoft's conduct relating to these software products to contribute to an adverse effect on rivals' competitive offerings.

8.47 We have provisionally found that Microsoft's licensing practices have sufficient effect to reduce competition in cloud services.

8.48 While smaller cloud providers are not subject to the same restrictions that AWS and Google face, we have seen limited evidence that they benefit in terms of growth or scale from Microsoft's software licensing practices.

8.49 We have provisionally found that Microsoft's licensing practices are partially foreclosing AWS and Google.

Our provisional decision

8.50 We have provisionally found that high levels of market concentration and barriers to entry and expansion have enabled each of the two largest providers, AWS and Microsoft, to hold significant unilateral market power in these markets and to earn returns across the cost of their capital. This gives 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. This harm is exacerbated by the features we have found arising from technical and commercial barriers.

8.51 We have also provisionally found that there are technical barriers and commercial barriers in the form of egress fees to switching and multi-cloud that give rise to an AEC in cloud services in the UK by locking customers into their initial choice of provider which may not reflect their evolving needs and limiting 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.

8.52 We have provisionally found that Microsoft's licensing practices are partially foreclosing AWS and Google which is having an impact on their competitive positions and that this gives rise to an AEC in cloud services in the UK. It also exacerbates the AEC that we have provisionally found arising from high market concentration and barriers to entry and expansion in relation to Microsoft's significant unilateral market power.

Customer detriment

8.53 We consider that the AECs we have provisionally found may be expected to result in substantial customer detriment in cloud services in the UK.

8.54 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 quality and innovation are also important to customers.[2068]

8.55 We consider that the AECs we have provisionally found have had a material impact on customers' ability to switch, multi-cloud and exercise choice over their

[2068] CC3 (Revised), paragraph 104.

provider, which may ultimately be expected to impact the price and quality (including access to innovative new services) of cloud services.

8.56 We consider that detriment may manifest itself in terms of UK customers paying higher prices or benefiting from lower levels of quality or innovation for these services than they would if the markets were more competitive. By way of illustration, if prices are on average 5% above those in a well-functioning market, this would in aggregate lead to UK customers paying around £430 million more per year for these services than they would in a more competitive market.[2069] If quality or innovation were lower by the same degree, this would also have a material impact on customers.

[2069] Calculations are as follows: 9bn - (9bn/1.05) = 9bn - 8.57bn = 429m. £9bn was the value of IaaS and PaaS UK revenue in 2023. See revenue figures in chapter 2

9. Our proposed remedies

- This chapter covers the remedies that we propose to introduce to remedy, mitigate or prevent the AECs that we have provisionally found, and the customer detriment that may be expected to result from the AECs.
- We have assessed potential remedies that we could implement through the remedy-making powers provided by the Act, consisting of order-making powers or accepting undertakings from parties. We considered potential remedies in relation to technical and commercial barriers and Microsoft's licensing practices, and provisionally found that 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 provisions of the Act. Additional details on these considerations are set out in Appendix W.
- We consider that the recently introduced DMCC Act powers are better suited to addressing the concerns we have identified than the powers directly available to us in this market investigation. These 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, particularly as a result of their greater flexibility, including new powers designed to enhance the effectiveness of remedies, and better provisions for ongoing monitoring and oversight.
- We consider that focusing any interventions on the two largest providers of cloud services (ie AWS and Microsoft) would directly benefit the majority of UK customers as well as altering the wider competitive conditions and commercial pressures on other CSPs which, in combination, would effectively address the provisional AECs we have identified.
- We have provisionally concluded that the following potential remedies would represent as comprehensive a solution as is reasonable and practicable to the AECs and resulting consumer detriment that we have provisionally found:

    Remedy 1: a recommendation to the CMA Board to prioritise commencing an SMS investigation of AWS' digital activities in respect of cloud services, and if an SMS designation is made to consider imposing appropriate interventions such as those identified in this report; 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, and if an SMS designation is made to consider imposing appropriate interventions such as those identified in this report.

Introduction

9.1 We have identified competition concerns in the cloud services markets arising from high concentration, barriers to entry and expansion and the conduct of cloud service providers. Comprehensively addressing such concerns in digital markets has historically been difficult using the CMA's existing competition tools (eg through market investigations).[2070]

9.2 Subsequent to the reference from Ofcom, the UK Parliament passed legislation (the Digital Markets, Competition and Consumers Act, or DMCC Act) to create a new digital markets competition regime. The associated new powers commenced on 1 January 2025.[2071]

9.3 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.'[2072] 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.[2073]

9.4 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 this market investigation,[2074] as well as forward-looking conduct requirements which are not available to us using the CMA's remedy-making powers under the Act.[2075] Furthermore, the DMCC Act allows for greater flexibility than is currently available to the CMA via existing competition tools (eg it will generally be easier to iterate remedies), as well as introducing new 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 ex ante functions to investigate and monitor digital markets, we consider it better suited to provide an ongoing monitoring and oversight function particularly if this is required for an extended time period.

9.5 We consider that the introduction of the CMA's digital markets functions is relevant to our assessment as it provides a potential remedy to the concerns we have identified, in the form of a recommendation to the CMA Board to consider using its

---

[2070] Sarah Cardell: Ensuring digital market outcomes that benefit people, businesses and the wider UK economy.
[2071] 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).
[2072] A new pro-competition regime for digital markets, Consultation document.
[2073] Sarah Cardell: Ensuring digital market outcomes that benefit people, businesses and the wider UK economy.
[2074] See DMCC Act, section 51(1), which references Schedule 8 to the Act. (Schedule 8 sets out the permitted remedial measures that can be included in an order remedying any AECs found in a market investigation).
[2075] For example, conduct requirements can be imposed ex ante to prevent harm from arising.

DMCC Act powers. In our remedies assessment, therefore, we have considered the extent to which the nature and extent of the DMCC Act powers are better suited to addressing the concerns we have identified than the powers directly available to us in this market investigation.

9.6 We have provisionally found that our proposed recommendations to the CMA Board would effectively address the competition concerns we have identified, and that remedial action aimed at influencing the behaviour of AWS and Microsoft would be effective in addressing concerns observed on a market-wide basis. We consider that measures addressed to AWS and Microsoft would directly benefit the majority of UK customers. In addition, 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.

9.7 In this chapter, and the associated appendix, we have set out our assessment of potential remedies, consistent with our legal duties and the current guidance for market investigations.

9.8 We welcome views on our remedies assessment and proposals through this consultation process and will consider submissions ahead of taking our final decision.

Structure and context

9.9 This chapter covers the package of remedies that we propose to introduce to remedy, mitigate or prevent the AECs that we have provisionally found. The chapter is structured as follows:

(a) First, we set out the relevant framework, guidance and process we have followed. We then discuss how the CMA's new powers under the DMCC Act,[2076] factor into our remedy assessment.

(b) Second, we set out our views on the proposed recommendations. This includes explaining the relevant background, summarising stakeholders' views, describing the remedy and its intended effects on the AEC(s), explaining our design considerations, and setting out our assessment on effectiveness and proportionality.

(c) Third, we include an explanation of other potential remedies we considered but that we are not proposing to progress.

(d) Last, we describe our provisional decision on remedies.

[2076] Digital Markets, Competition and Consumers Act 2024.

9.10 Additional details on our remedies assessment and considerations are included in Appendix W.

9.11 We welcome views from parties on the design, implementation, effectiveness and proportionality of our proposed remedies. We will consider evidence and submissions on our proposed remedies, and if we find one or more AECs in our final report we will take all relevant evidence and submissions into account in developing our final remedies.

Framework for our assessment of proposed remedies

9.12 If we identify any AECs, 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 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.[2077]

9.13 In coming to a view on remedies, the Act 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'.[2078] 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.[2079]

1.2 A detrimental effect on customers is one that results, or may be expected to result, from any AECs and takes the form of:[2080]

(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

---

[2077] The Act, section 134(4); CC3 (revised), paragraphs 325 to 328.
[2078] The Act, section 134(6).
[2079] CC3 (revised), paragraph 329.
[2080] The Act, section 134(5).

(b) less innovation in relation to such goods and services.

9.14 Where more than one measure is introduced, the CMA will consider the way in which the measures are expected to interact with each other,[2081] which may be complementary in their effectiveness and costs, or they may be in tension in some areas. We would consider both the effectiveness of individual measures in the context of an overall package, and the potential package of remedies as a whole.

9.15 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 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 measures are not available, or as an interim solution while other measures take effect.[2082]

9.16 In assessing remedies, the CMA will consider the extent to which different remedy options are likely to be effective in achieving their aims, including their practicability.[2083] The effect of any remedy is always uncertain to some degree and in assessing effectiveness:

(a) we consider the risks associated with different remedies and will tend to favour remedies that have a higher likelihood of achieving their intended effect;[2084]

(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;[2085]

(c) we will generally look for remedies that prevent an AEC by extinguishing its causes, or that can otherwise be sustained for as long as the AEC is

[2081] CC3 (revised), paragraph 393.
[2082] CC3 (revised), paragraphs 330 to 333.
[2083] CMA3, Market Studies and Market Investigations: Supplemental guidance on the CMA's approach, paragraph 4.15.
[2084] CC3 (revised), paragraph 335 and CMA3, paragraph 4.16.
[2085] CC3 (revised), paragraph 336 and CMA3, paragraph 4.17.

expected to endure.[2086] We also tend to favour remedies that are expected to show results within a relatively short time;[2087] and

(d) where more than one measure is being introduced as part of a package of remedies, we will consider the way in which the measures are expected to interact with each other.[2088]

9.17 Where we have identified an effective remedy, we will consider its proportionality. In assessing proportionality, we are guided by the following principles set out in our guidance. 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;

(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.[2089]

9.18 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.[2090]

9.19 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

[2086] CC3 (revised), paragraph 337 and CMA3, paragraph 4.18.
[2087] 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.
[2088] CC3 (revised), paragraph 341 and CMA3, paragraph 4.24.
[2089] CC3 (revised), paragraph 344.
[2090] CC3 (revised), paragraph 348.

(c) if remedies extinguish Relevant Consumer Benefits (RCBs), the amount of RCBs foregone may be considered to be a relevant cost of the remedy.[2091]

9.20 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.[2092]

Proposed remedies: recommendations to the CMA Board

Background

9.21 In our potential remedies working paper, we stated that we would consider the CMA's new powers under the DMCC Act,[2093] as part of our consideration of any potential remedies.[2094]

Description of the new digital powers

9.22 The DMCC Act gives the CMA new powers to intervene in digital markets by establishing a new, targeted regime. It has been introduced in recognition of the fact that some digital markets share a combination of characteristics that can cause the market to 'tip' in favour of one, or a few firms.[2095] The new regime will strengthen the existing competition rules and allow the CMA to take a targeted and iterative approach to address concerns in certain digital markets.[2096]

9.23 The digital markets competition regime will apply 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.[2097]

[2091] The CMA may have regard to the effect of any remedial action on any RCBs of the feature(s) of the market(s) concerned (the Act, 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 (the Act, section 134(8)).
[2092] 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.
[2093] Digital Markets, Competition and Consumers Act 2024.
[2094] Potential remedies, paragraph 3.24.
[2095] A new pro-competition regime for digital markets, Consultation document.
[2096] Sarah Cardell: Ensuring digital market outcomes that benefit people, businesses and the wider UK economy.
[2097] Digital Markets, Competition and Consumers Act 2024, section 3(1).

9.24 Under the DMCC Act, for a firm to have SMS in respect of a digital activity it must have:

(a) substantial and entrenched market power in a digital activity which is linked to the United Kingdom;[2098]

(b) a position of strategic significance in respect of that activity;[2099] and

(c) global turnover of more than £25 billion or UK turnover of more than £1 billion.[2100]

9.25 Decisions in respect of the new digital markets competition regime are the responsibility of the CMA Board. The DMCC Act provides that certain decisions must be made by the CMA Board,[2101] including whether to begin an initial SMS investigation, whereas other decisions may be delegated. The CMA Board, or an appropriately authorised Board committee, will decide whether to make an SMS designation.[2102] An SMS investigation has a nine-month statutory deadline.[2103]

9.26 The DMCC Act sets out that once the CMA designates a firm with SMS in respect of a digital activity, it may impose conduct requirements (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.[2104] 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.[2105] PCI investigations will be subject to a nine-month statutory deadline.[2106]

9.27 The new regulatory powers provided by the DMCC Act commenced on 1 January 2025.[2107]

[2098] Digital Markets, Competition and Consumers Act 2024, sections 2(1) and (2).
[2099] Digital Markets, Competition and Consumers Act 2024, section 2(b).
[2100] Digital Markets, Competition and Consumers Act 2024, section 7.
[2101] Digital Markets, Competition and Consumers Act 2024, section 106.
[2102] Digital Markets, Competition and Consumers Act 2024, section 106. See further CMA194: Digital markets competition regime guidance, paragraphs 9.39-9.40.
[2103] Digital Markets, Competition and Consumers Act 2024, section 14(2). This is subject to the CMA's power to extend the deadline in the circumstances set out in section 104.
[2104] Digital Markets, Competition and Consumers Act 2024, section 19.
[2105] Digital Markets, Competition and Consumers Act 2024, section 46.
[2106] Digital Markets, Competition and Consumers Act 2024, section 50(1). This is subject to the CMA's power to extend the deadline in the circumstances set out in section 104.
[2107] The DMCCA Commencement Regulations, sections 1(2), 2(1) and 2(2).

Relevance to this market investigation

9.28 Under s134(4)(b) of the Act, we can recommend the taking of action by others for the purposes of remedying, mitigating or preventing (either on their own or in combination with other measures) an AEC. Such a recommendation can form part of a remedies package, or all of it.

9.29 Our potential remedies working paper outlined that we could make a recommendation to the CMA Board to consider whether to make an SMS designation and whether to impose CRs or PCIs where we considered that use of the CMA's pending new functions would be an effective remedy to any AEC(s) identified.[2108]

Stakeholder views

9.30 There were mixed views from stakeholders on the role of the CMA's DMCC Act powers in this market investigation:

(a) Microsoft submitted that the powers were 'ideally suited to intervene at some point in the future if and when outcomes merit intervention.'[2109]

(b) AWS submitted that the evidence gathered 'render[s] unwarranted an intervention under the DMCC Act.'[2110]

(c) Google considered that it would be 'more appropriate in these circumstances for the CMA to use its existing market investigation remedy powers.'[2111]

(d) One customer believed the current competitive landscape warranted intervention through DMCC Act powers and that doing so should be a 'high priority' for the CMA.[2112] We discuss this customer's more detailed suggestions in the section 'Other remedy options that we are minded not to ' below.

Description of the remedies

9.31 We are proposing:

(a) Remedy 1: a recommendation to the CMA Board to prioritise commencing an SMS investigation[2113] of AWS' digital activities in respect of cloud services,

[2108] Potential remedies, paragraph 3.23.
[2109] Microsoft Response to the Competitive Landscape, Committed Spend Agreements and Egress Fees Working Papers, paragraph 11.
[2110] AWS' response to the CMA's updated issues statement and working papers, dated 25 June 2024, paragraph 74.
[2111] Google's response to the CMA's Potential remedies working paper, dated 27 June 2024, paragraph 11.
[2112] [✂] submission to the CMA [✂].
[2113] This refers to an initial SMS investigation under section 9 of the DMCC Act to consider whether to designate an undertaking as having SMS in respect of a digital activity in accordance with DMCC Act, section 2.

and if an SMS designation is made, to consider imposing appropriate interventions such as those identified in this report; 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, and if an SMS designation is made, to consider imposing appropriate interventions such as those identified in this report.

9.32 In arriving at our provisional decision, we have taken into account the considerations relevant to making a recommendation and to determining its scope set out in the Act[2114] and in relevant guidance applicable to market investigations.[2115]

1.3 Key considerations are:

(a) Effectiveness: how recommendations to the CMA Board would be an effective and comprehensive remedy; and the likelihood that the recommendations will be acted upon by the CMA Board and, if so, over what time period; and

(b) Proportionality: why recommendations to the CMA Board are proportionate.

Effectiveness

9.33 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. [2116]

1.4 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.[2117] We have reflected these risks in our effectiveness assessment below, in particular by focusing on:

(a) how the recommendations will address the identified harms; and

(b) the likelihood of the recommendations being acted upon and, if so, the time period over which they are likely to be implemented.

[2114] The Act, sections 134(6) and (7).
[2115] These are the considerations set out in CC3 (revised) at paragraphs 379, 380, 390, 391 and Annex B, paragraphs 94 to 102.
[2116] CC3 (revised), paragraph 335 and CMA3, paragraph 4.16.
[2117] CC3 (revised), Annex B, paragraph 97.

How the recommendations will be effective in addressing the harm

9.34 As set out in Chapter 8, we have provisionally identified features that in combination are giving rise to AECs in cloud services in the UK. These features are high levels of market concentration, barriers to entry and expansion, technical and commercial barriers to switching and multi-cloud and Microsoft's licensing practices. 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.[2118]

9.35 As a result, the new digital markets competition regime and the powers provided by the DMCC Act are well suited to address the types of issues that we have identified in connection with the provision of cloud services, supporting the potential for recommendations to the CMA Board to represent an effective and comprehensive remedy. The most relevant aspects of these powers for our consideration of recommendations to the CMA Board include:

(c) a wide range of intervention powers;

(d) flexibility to allow for future variation and iteration;

(e) powers to test and trial potential interventions; and

(f) ongoing regulatory oversight including a wide range of investigatory powers.

9.36 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.[2119] 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.[2120]

9.37 The CMA will be able to impose CRs to guide the practices of SMS firms, and it has the ability to iterate CRs over time to ensure they are effective. The CMA will also have a duty to consider on an ongoing basis whether to impose, vary or

[2118] 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.

[2119] Whilst Conduct Requirements are largely limited to the digital activity itself, Section 20(3)(c) of the DMCC Act will allow 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.

[2120] The DMCC Act expressly references Schedule 8 of the Act 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.

revoke CRs,[2121] as well as a duty to review any PCO at a pre-identified date to determine whether to revoke or replace them thereby allowing the CMA to proactively iterate.[2122] This flexibility to respond to changing circumstances, and iterate to respond to firm behaviour, is well suited to cloud services.

9.38 CRs are also well suited to principles-based approaches. 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.[2123]

9.39 In relation to the potential remedies we have considered, we note that the ability for an intervention to iterate over time will support an effective and comprehensive solution:

(a) We considered that any remedies seeking to address technical barriers, including those that could be introduced under DMCC Act powers,[2124] need to be sufficiently well-specified but still flexible enough to remain appropriate for a wide range of uncertain future changes in cloud services, which is technically complex. These remedies 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 any such remedy design is the ability to iterate the remedy so that it is, and will continue to be, effective, adapting to changing market conditions and preventing any harmful unintended consequences. Further details on technical barriers potential remedies are included in Appendix W.

(b) We considered that any remedies seeking to address Microsoft's licensing practices, including those that could be introduced under DMCC Act powers,[2125] should be implemented as a combination of principles-based and rules-based interventions to comprehensively address the specific licensing practices engaged in by Microsoft. Given the complexity of the potential remedies package, a key element in the remedy design is the ability to refine, recalibrate or amend the remedy over time. Further details on our assessment of potential remedies to Microsoft's licensing practices are included in Appendix W.

(c) We considered that any remedies in these markets, including those seeking to address egress fees, would benefit from being considered as part of a

[2121] Digital Markets, Competition and Consumers Act 2024, section 25.
[2122] Digital Markets, Competition and Consumers Act 2024, section 55(1).
[2123] CMA194: Digital markets competition regime guidance, paragraph 3.29.
[2124] We note that a number of the permitted types of Conduct Requirements 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.
[2125] We note that a number of the permitted types of Conduct Requirements 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. Digital Markets, Competition and Consumers Act 2024, section 20.

single coherent package. As the harm to competition that we have provisionally identified arises from a combination of features, the ability to iterate remedies over time would provide greater flexibility in the design, implementation and monitoring of a complete and holistic set of effective interventions.

9.40 Regarding powers to test and trial potential interventions in the digital markets competition regime:

(a) The CMA will also be able to require SMS firms to perform a specified demonstration or test.[2126] This can include requiring a firm to demonstrate a technical process, such as how an algorithm operates, and can require a firm to vary its usual conduct, for example to assess the effect of different choice architecture and assess compliance with particular requirements. 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.[2127]

(b) These new testing and trialling powers would be relevant if the CMA were to consider the potential remedies we assessed for certain technical changes to reduce frictions. 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 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. Some of the technical barriers potential remedies discussed in Appendix W, in particular, would likely benefit from testing and trialling provisions.

9.41 The DMCC Act provides the CMA with a range of investigatory powers which it may use when administering and enforcing the digital markets competition regime. Monitoring will be a key part of the CMA's role in overseeing the digital markets competition regime. It will allow the CMA to respond quickly where firms fail to comply with competition requirements, in particular to consider whether enforcement action is warranted for non-compliance and to inform whether new or varied competition requirements may be necessary, as well as if it is appropriate to remove or narrow particular interventions. This flexibility would enable the CMA to take account of and react to any relevant changes in relation to cloud services.

9.42 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.

[2126] Digital Markets, Competition and Consumers Act 2024, section 69(5).
[2127] Digital Markets, Competition and Consumers Act 2024, see sections 51(3) - (4).

9.43 We have also considered whether our proposed 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 provisional 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 60-70% across IaaS and PaaS (and 80-90% of IaaS market).[2128] 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 markets in question.

9.44 We provisionally consider that the recommendations we propose above would effectively and comprehensively address the AECs we have provisionally identified.

Likelihood of the recommendations being acted upon and likely timing

9.45 Our guidance states that when evaluating the effectiveness of a recommendation as a proposed remedy, we 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.[2129]

9.46 In reaching this view, we must have regard to:

(a) the stated policy of the body to which the recommendation is to be directed; and

(b) the possibility that that stated policy may change, either in light of our recommendation or subsequent events.[2130]

9.47 To assess the likelihood that our recommendations will be acted upon, we set out below:

(a) the decision-making processes under the DMCC Act;

(b) a consideration of the criteria that the CMA will use when deciding whether to designate a firm as having SMS; and

---

[2128] Shares of supply are by reference to revenue, see Chapter 3 (market shares).
[2129] CC3 (revised), paragraph 391.
[2130] CC3 (revised), paragraph 391.

(c) timing expectations for our recommendations to be acted on.

9.48 The DMCC Act provides that certain decisions are the responsibility of the CMA Board including whether to begin an initial SMS investigation, and the CMA Board or relevant Board committee will decide whether to make an SMS designation and whether to impose CRs and/or PCIs.[2131]

9.49 We have not carried out an assessment of whether or not AWS and Microsoft meet the specific SMS designation criteria as the CMA Board is ultimately responsible for deciding whether to prioritise a firm's digital activity for SMS investigation, whether the tests for designation are met and whether CRs and/or PCIs should be imposed. However, we have considered the regulatory framework created by the DMCC Act and the CMA Board's stated policy in assessing the likelihood that any recommendation we make will be acted upon. We note that the CMA Board's stated policy is that it will have regard to its Prioritisation Principles[2132] when considering which firms and digital activities to prioritise for SMS investigations, including taking account of evidence gathered through its markets or enforcement functions.[2133]

9.50 We consider that 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.[2134] We have provisionally found that each of AWS and Microsoft has significant market power in the supply of cloud services in the UK and that is likely to endure.[2135] We consider that AWS' and Microsoft's respective positions in the provision of cloud services appear to be strategically significant. We note that the provision of cloud services itself is a strategically important digital activity[2136] 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.[2137] We also note that each of Amazon and Microsoft have a global turnover in excess of £25 billion.[2138]

2131 DMCC Act, section 106. See further CMA194: Digital markets competition regime guidance, paragraphs 9.39-9.40.
2132 CMA Prioritisation Principles.
2133 CMA194: Digital markets competition regime guidance, paragraph 2.83c.
2134 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.
2135 See chapter 8 (AEC decision chapter).
2136 The importance of cloud services has been recognised by a number of sector regulators (see Regulatory Barriers section of the Barriers to Entry and Expansion chapter), and the UK government has recently classed data centres as Critical National Infrastructure.
2137 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.
2138 Amazon and Microsoft's public accounts. We note that section 7(2) of the DMCC Act 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.'

9.51 On prioritisation, we anticipate that the work we have undertaken in this market investigation would be factored into the CMA Board's considerations. In reaching a final decision on remedies and prior to making any recommendations we will consult with the CMA Board. As noted in our guidance, such consultation would enable us to further understand the benefits and risks of implementing the recommendation, to inform decisions about the specification of any recommendation, as well as informing our judgement about the likelihood of the recommendation being accepted.[2139]

9.52 On timing, the new regulatory powers provided by the DMCC Act commenced on 1 January 2025.[2140] The subsequent timelines for any action under the DMCC Act are that the CMA would have to complete any SMS investigations within a statutory timeframe of nine months,[2141] and the CMA would be able to impose an initial set of CRs on designation or as soon as practicable afterwards,[2142] while any PCIs would follow a further nine-month investigation.[2143]

9.53 In view of the above, we provisionally consider it likely that our recommendation to the CMA Board would be acted upon in a timely manner.

Provisional view

9.54 We consider that the range of powers that will be available to the CMA under the DMCC Act are particularly well suited to address the concerns we have provisionally identified in this report in a timely manner, enabling a targeted and iterative approach to competition interventions that would avoid the specific risks we have identified with using our market investigation order making powers.

9.55 Our provisional view is that the ongoing regulatory oversight provided by designations of each of AWS and Microsoft under the DMCC Act would effectively and comprehensively address the features and AECs we have provisionally identified, including the unilateral market power of AWS and Microsoft, technical and commercial barriers and Microsoft's licensing practices.

Proportionality

9.56 In this section we set out our assessment of whether the proposed remedies would be proportionate to achieve their aim. We do so by considering whether each remedy would:[2144]

[2139] CC3 (revised), Annex B, paragraph 99.
[2140] DMCCA Commencement Regulations, sections 1(2), 2(1) and 2(2).
[2141] Digital Markets, Competition and Consumers Act 2024, section 14(2). This is subject to the CMA's power to extend the deadline in the circumstances set out in section 104.
[2142] CMA194: Digital markets competition regime guidance, paragraph 3.40.
[2143] Digital Markets, Competition and Consumers Act 2024, section 50(1). This is subject to the CMA's power to extend the deadline in the circumstances set out in section 104 of the Act.
[2144] CC3 (Revised), paragraph 344.

(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

9.57 We discussed above the effectiveness of the recommendations, and provisionally conclude that they would be effective in achieving their aim.

No more onerous than needed

9.58 Our provisional AECs have identified concerns arising with AWS and Microsoft, and so it is necessary to include both in our proposed recommendations. We do not consider that the design parameters could be changed to make the approach less onerous while retaining the effectiveness of the approach.

Least onerous if there is a choice between several effective measures

9.59 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 currently identified an alternative remedy that would be less onerous whilst still being a comprehensive solution to the AECs that we have provisionally found.

Does not produce disadvantages which are disproportionate to the aim

9.60 We consider that the recommendations would not produce disadvantages which are disproportionate to their aims. 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. For example, the CMA can only impose CRs or PCIs if it would be proportionate to do so.[2145,2146] It will be for

[2145] 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; Digital Markets, Competition and Consumers Act 2024, section 19(5)-(8).

[2146] 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; Digital Markets, Competition and Consumers Act 2024, section 46.

the CMA Board (or its delegates) to assess the costs associated with any potential interventions it considers imposing following any SMS designation.

Other remedy options that we are minded not to progress

9.61 During the course of this investigation, we have considered a number of additional remedy options which we do not currently propose to include in any remedies package. These are:

(a) measures we are not minded to progress under the remedy-making powers provided by the Act for:

(i) technical barriers;

(ii) egress fees; and

(iii) Microsoft licensing practices.

(b) structural remedies; and

(c) utility-style regulation.

9.62 We consider each of these to explain our reasoning.

Measures we are not minded to progress under the remedy-making powers provided by the Act but recommend the CMA consider if SMS designations are made

9.63 We consulted on other potential remedies in our working papers which we have provisionally decided not to pursue through remedy-making powers conferred under the Act. These were specifically aimed at:

(a) Technical barriers;[2147]

(b) Egress fees;[2148] and

(c) Microsoft licensing practices.[2149]

9.64 In our assessment in this chapter and in Appendix W, we have considered how to achieve as comprehensive a solution as is reasonable and practicable to the AECs arising from these features. Although we consider that implementing potential remedies using our market investigation powers could, in principle, be capable of addressing the features we have provisionally found are giving rise to

[2147] Technical barriers working paper, chapter 9.
[2148] Egress fees working paper, chapter 4.
[2149] Licensing working paper, chapter 7.

AECs, these would not result in as comprehensive a solution as is reasonable and practicable.

9.65 As noted above, it is recognised that our remedy-making powers in market investigations are more suited to one-off interventions and not as well suited to addressing certain issues arising in evolving digital markets. We find that this is the case in relation to the features we have provisionally found. In particular, we consider that implementing the potential 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,[2150] the inability to test and trial remedies of this nature,[2151] and/or 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 market investigation powers in 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 are provided in Appendix W.

9.66 However, we consider that many of the proposed remedies we have considered could form part of an effective and comprehensive solution if implemented through the digital markets competition regime under the DMCC Act which was specifically designed for digital markets such as cloud services. Therefore, we recommend that if the CMA designates AWS and/or Microsoft with SMS status in respect of cloud services that it considers imposing appropriate interventions such as those identified in this report.

[2150] We note that section 162A of the Act, introduced through s139 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 the Act. 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.

[2151] Sections 161B to 161E and 162B of the Act, introduced through 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 the Act. 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.

Considerations in the context of making recommendations

9.67 In the section 'How the recommendations will be effective in addressing the harm' above, we have discussed the DMCC Act powers, and the implications this would have for addressing the types of issues that we have identified.

9.68 We consider that if the CMA were to designate AWS and/or Microsoft with SMS and consider the possible remedies we have identified in Appendix W, it would have a wide range of intervention powers, flexibility to allow for future variation and iteration, powers to test and trial potential interventions, and ongoing monitoring via a wide range of investigatory powers, which we consider would likely address many (if not all) of the major risks we have identified in these potential remedies assessments. However, it will be for the CMA Board (or its delegates) to decide whether to commence an SMS investigation and, were it to designate AWS and Microsoft with SMS, whether to implement any of the remedies we have discussed in Appendix W.

9.69 We have also considered whether our proposed 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. As discussed in the 'How the recommendations will be effective in addressing the harm' section above, our provisional view is that interventions aimed at influencing the behaviour of AWS and Microsoft would be an effective remedy because they would directly benefit the majority of UK customers. In addition, 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.

Technical barriers

9.70 In Appendix W, we have set out our assessment and provisional views on each of the potential remedies to 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 the 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 trailing of standards so that they could be refined or recalibrated efficiently.

(c) Accordingly, a key element in the design of a standardisation remedy would be the ability to iterate the remedy in response to technical changes, new innovations or the re-introduction of technical frictions in an attempt to circumvent the aim of the remedy. 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.

(d) We consider that a transparency remedy could be implemented through a market investigation order but we are of the provisional view that this would not be an effective remedy on its own, as information transparency remedies would not reduce the prevalence or impact of technical barriers that are currently present. Furthermore, there is 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.

(e) We are of the view that requiring the largest cloud providers to offer abstraction layers, potential remedies to reduce latency and a remedy to require more cloud agnostic training (potential remedies 5, 6, and 8 in Appendix W) are currently likely to face greater challenges to their effectiveness, including risks that are unlikely to be addressed by expanded remedial powers such as the ability to iterate or to test and trial remedies in advance of implementation.

9.71 In light of these assessments, our provisional view is that, although in principle a package of remedies implemented through a market investigation order could address these technical barriers, 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 the Act.

9.72 We consider that if the CMA were to designate AWS and Microsoft with SMS in respect of cloud services and consider the imposition of appropriate interventions such as those considered in this report, it would have the ability to test and trial remedies, 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

9.73 We consider that a ban on egress fees for switching and multi-cloud, applied to at least AWS and Microsoft, could, in principle, represent an effective standalone remedy to the egress fees feature we have provisionally identified. However, any egress fees remedy needs to be considered in the context of the wider remedial action we are proposing, specifically the recommendations to the CMA Board and our expectation that the CMA Board would act upon these in a timely manner.

9.74 We are concerned about the effectiveness of a remedies package that included our proposed recommendations to the CMA alongside an egress fees remedy using our remedy-making powers under the Act. These concerns would arise from a process of implementing, monitoring and maintaining any remedies implemented under the DMCC Act in parallel with an egress fees remedy under the Act, 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) During any overlapping period between a remedy implementation phase using the remedy-making powers under the Act and any SMS investigations under the DMCC Act, AWS and Microsoft would be engaging with two separate parts of the CMA operating under different legal frameworks in relation to the same markets/activities. We consider that this would add considerable complexity associated with implementing remedies in these markets through these two regimes.

(b) Although we consider the risks associated with monitoring and enforcement of an egress fee remedy (in particular a ban) 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.

(c) Any increased regulatory complexity could undermine the effectiveness of interventions considered appropriate under the DMCC Act. 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.

9.75 We also consider that implementing and maintaining a remedy using our remedy-making powers under the Act in parallel with the CMA exercising its powers under the DMCC Act in respect of AWS and Microsoft in the same markets/activities could be an inefficient use of CMA resources, particularly given there is an overlap between the remedial powers in the two regimes.

9.76 In view of the above, we consider that implementing an egress fees remedy using our remedy-making powers under the Act would introduce risks to an overall remedies package that included our proposed recommendations to the CMA

Board, and this would risk undermining the effectiveness of the remedies package as a whole.

9.77 These risks would not arise were the CMA to consider the imposition of appropriate interventions to address egress fees following any SMS designation of AWS and/or Microsoft in respect of cloud services. 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 impose appropriate interventions to address egress fees such as those identified in this report.

Microsoft licensing practices

9.78 In Appendix W, we considered potential remedies aimed at addressing the AEC arising from Microsoft's licensing practices through our remedy-making powers under the Act. We identified three measures as being necessary to any potential remedial action in this regard, namely:

(a) Remedy A - Fair, reasonable and non-discriminatory pricing: This potential remedy would require Microsoft to apply a 'FRAND' approach in relation to pricing its software products regardless of which cloud they are hosted on.[2152] It would also include information transparency obligations.

(b) Remedy B - Product functionality and technical performance: This potential remedy would impose 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.

(c) Remedy C - Licence transfer: This potential remedy would focus specifically on contractual licensing practices relating to the transfer and/or deployment by end customers of previously purchased software products on the cloud of their choice.

9.79 We have identified risks with implementing this package of remedies using our remedy-making powers under the Act:

(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

[2152] 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.

address the potential for circumvention without the ability to iteratively develop these measures.

(b) Effects of complexity: due to the complexity of the remedies, any initial design would likely require refinement, recalibration, or correction, particularly over time.

(c) Monitoring and enforcement: we would expect that this package of remedies would involve an elaborate monitoring and compliance programme, which would involve establishing and maintaining a framework akin to the role envisaged for the CMA under the digital markets competition regime and introduce potentially overlapping regulatory regimes in relation to cloud 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 remedies package over time increases this concern.

9.80 In principle, it might be possible to address some of these risks in the design of a market investigation order. For example, in Appendix W we discuss (and dismiss) the potential for establishing a new implementation and oversight body. Alternatively, a market investigation order could seek to introduce greater flexibility through explicit review points for the specific remedy design (eg taking place on a regular time period or being triggered by particular events) to try and ensure that the remedies were effective on an ongoing basis. However, it is not clear that this would allow for sufficient flexibility to comprehensively address the AEC we have provisionally found.

Structural remedies

9.81 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.

9.82 In our working paper on potential remedies, we included the possibility of structural remedies, in particular requiring certain cloud providers to divest:

(a) their entire cloud operations;

(b) part of their cloud operations; or

(c) their IaaS or PaaS operations.[2153]

[2153] Potential remedies, chapter 5.

9.83 We also considered the potential for operational separation. Operational separation involves requiring companies to operate their cloud businesses independent of their other business units. Operational separation is not a structural remedy, as it does not involve a transfer of ownership, but it could potentially achieve similar aims to structural remedies.[2154]

9.84 In our working paper, we explained that structural remedies (and operational separation) raised particular risks including that they may:

(a) introduce inefficiencies by creating more, smaller market participants;

(b) disrupt legitimate, pro-competitive combinations between cloud and non-cloud services; and

(c) create distortions in the market by embedding artificial distinctions on how products are categorised. This is particularly relevant to the remedy that would require cloud providers to separate and divest one of their IaaS or PaaS operations.

9.85 There are also practical considerations, as it would be very difficult to separate cloud-specific assets from a company's wider business and there is a risk that without ongoing regulation over time the companies directly affected by any structural remedies either regain their previous scale or new providers gain their market power.

9.86 For these reasons, we stated that we were minded not to prioritise further consideration of structural or operational separation remedies.

9.87 We invited views on this in our potential remedies working paper, and no parties told us that structural remedies would be an effective and proportionate remedy, or provided evidence to support such a view. Hence, it is our provisional view that at this point in time structural remedies and operational separation would not represent an effective and proportionate remedy to address the harm arising from the effects of barriers to entry and expansion in the cloud services market.

Utility-style regulation

9.88 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 direct price regulation of some IaaS services, guaranteed minimum service standards and FRAND conditions on access to services.[2155]

[2154] Potential remedies, paragraphs 5.11 to 5.13.
[2155] [✂] submission to the CMA [✂].

9.89 We are concerned that 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.

9.90 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. We consider that the recommendations described above would represent a comprehensive remedy to the AECs. We note that, if the CMA designated AWS and/or 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.

Provisional decision on remedies

9.91 In response to the features and AECs that we have provisionally found, we have provisionally concluded that the following remedies would represent as comprehensive a solution as is reasonable and practicable to the AECs and resulting consumer detriment:

(a) Remedy 1: a recommendation to the CMA Board to prioritise commencing an SMS investigation of AWS' digital activities in respect of cloud services, and if an SMS designation is made to consider imposing appropriate interventions such as those identified in this report; 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, and if an SMS designation is made to consider imposing appropriate interventions such as those identified in this report.