Public cloud infrastructure services market investigation

Updated issues statement, 6 June 2024

Introduction

1. On 18 October 2023, we published an issues statement for the market investigation into public cloud infrastructure services (cloud services).[1] We are now providing an update on our emerging thinking and on the theories of harm that we have been investigating, based on the evidence we have received to date and our initial analysis. We have not yet reached any provisional conclusions and our views as set out in this statement and our working papers may change in light of our ongoing analysis as well as comments and further evidence we receive in response to these publications.

2. We have published six working papers containing a snapshot of our work to date on the following issues: the competitive landscape and market outcomes, egress fees, committed spend discounts, technical barriers, software licensing and potential remedies. We have also published primary research undertaken amongst customers of cloud services (the Jigsaw report).[2]

3. This document and our working papers are published to invite comments from interested parties on those issues that we consider are likely to represent the focus of our investigation as we progress to provisional findings.

4. If any interested parties wish to provide further relevant evidence and submissions to our inquiry or if they consider that there are any inaccuracies or gaps in our understanding or evidence base, we invite them to respond by email to CloudMI@cma.gov.uk by 27 June 2024.

5. We aim to publish submissions we receive in response to this update. We therefore request that, when submitting any response, parties provide a non-confidential version at the same time to avoid unnecessary delay in

[1] See Annex 1 to this paper for further information on the background to this inquiry
[2] See Cloud services market investigation

publication. If any party has concerns about publication, it should raise these at the time of submitting its response.

Our emerging views on the competitive landscape and market outcomes in cloud services

Nature of competition

6. Cloud services are increasingly important inputs to many businesses and organisations across the UK economy and across a range of different industries.

7. Evidence we have seen to date shows that a relatively small number of high-spend customers account for a large proportion of cloud providers' UK revenues and a relatively large number of low-spend customers are responsible for a small proportion of their revenue. In particular, the top 10% of customers account for a very large majority of revenues and the top 1% account for over half of revenues.

8. Customers buy the large majority of cloud services directly from cloud providers. Most customers have standard contracts that have been agreed without negotiation, but larger customers either engage in bilateral negotiations and occasional tenders and are able to negotiate terms that depart from standard contracts.

9. There are different models of multi-cloud use, and we cannot accurately measure the full extent of switching by customers, or the extent to which customers use multiple clouds.

10. However, the evidence we have seen to date suggests that, while there is some degree of multi-cloud use, it may be quite limited in scope and mostly found amongst larger customers. This evidence also suggests that switching between cloud providers is uncommon. We are continuing to consider the evidence of the prevalence of switching and use of multiple public clouds by customers.

11. Cloud providers compete on a range of factors and the factors that seem to be the most important to customers when choosing their main public cloud provider are service quality, price (including discounts and/or cloud credits), data sovereignty requirements, range of services and the number and location of data centres.

12. The supply of cloud services has evolved over time: AWS launched the first public cloud service in 2006; Microsoft made its cloud service Azure

generally available in 2010 and Google made Google Cloud generally available in 2011. Cloud services grew as providers competed to attract customers to the public cloud for the first time.

13. In recent years cloud services have grown significantly both in terms of revenues (UK cloud revenues more than doubled during the period from 2019 to 2022) and data centre capacity in the UK and Europe (more than doubled between 2020 and 2023).

14. The evidence we have seen to date suggests that the number of customers moving to public cloud for the first time has decreased and in 2021 and 2022 new customers made up a small proportion of year-on-year growth. Demand for cloud services is growing because existing customers are expanding their existing workloads and/or adding new workloads.

15. If existing customers face barriers to switching and multi-cloud, then the strength of competition is likely to be weaker than when competition was focused on customers moving to public cloud for the first time. We consider the extent to which there may be particular barriers to switching and multi-cloud below.

Market definition

16. Cloud services include both infrastructure as a service (IaaS) and platform as a service (PaaS).[3] Our emerging view is that both IaaS and PaaS are relevant markets. We set out below our considerations in coming to this view.

17. We considered IaaS and the extent to which PaaS is substitutable for IaaS and thus in the same relevant market. Our emerging view is that there is a relevant product market for the supply of IaaS and that, where relevant, PaaS should be considered as an out-of-market constraint.

18. In particular, we recognise that for some customers and for some workloads, IaaS and PaaS are substitutes. However, evidence we have seen to date from customers indicates that PaaS is not a good substitute for IaaS for most customers and workloads and most customers expressed that they are unwilling to substitute between the two, even if it may be technically possible to do so.

[3] IaaS covers services that provide access to processing, storage, networking, and other raw computing resources; PaaS covers services that can be used to develop, test, run and manage applications in the cloud

19. We considered whether software as a service (SaaS)[4] is substitutable for PaaS. The evidence we have seen to date on the extent of substitutability between PaaS and SaaS is mixed and limited.

20. We also considered the extent to which either traditional IT (that is, dedicated computing resources on-premises) or private cloud services (that is, a cloud which is exclusive to one customer) are substitutes for any of IaaS, PaaS or SaaS.

21. While some large customers of public cloud service may be able to react to a price increase by switching to private cloud or traditional IT, the evidence to date indicates that, even for large customers, any such switches would be unlikely due to the specific reasons they place workloads on public cloud and the costs and time associated with doing so. Therefore, our emerging view is that traditional IT and private cloud should be considered as out-of-market constraints where applicable.

22. Based on the evidence we have seen to date, our emerging view is that the geographic scope of the markets is more likely to be Europe-wide (that is, UK and EEA) than UK or global. To the extent that it is relevant, we will take into account non-European providers as out-of-market constraints. We will also take into account the way in which providers' strength in UK datacentres may be important in competing for some customers.

Shares of supply and market concentration

23. The evidence we have seen to date on shares by revenue, capacity and flows of new business suggests that AWS and Microsoft are the largest two cloud providers and significantly larger than Google, the next largest provider.

24. IaaS and PaaS are both concentrated and both have become more concentrated from 2019 to 2022 as the collective shares of AWS and Microsoft have increased in each: in IaaS, their collective shares were [70-80]%, in PaaS [50-60]% and across both IaaS and PaaS, they were [70-80]% in 2022. No other provider has a share over [5-10]%.

25. Many large customers told us that they do not see any suitable alternatives to AWS and Microsoft as their main cloud providers. They do not perceive smaller providers to have comparable offerings to AWS and

[4] SaaS are 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.

Microsoft, albeit Google is perceived as being closer than Oracle, IBM or others. However, smaller providers may still be seen as suitable alternatives for certain workloads.

Market outcomes

26. Our assessment of the profitability of cloud providers, based on evidence received to date, indicates that AWS and Microsoft have been generating returns from their cloud services above their cost of capital, and this could be expected to continue in the future.

27. Our assessment of prices, quality and innovation is less developed and requires careful interpretation, for the reasons given in our competitive landscape working paper.[5] Although we have seen evidence that in some areas cloud services are higher quality and more innovative than traditional IT services (which we consider to be outside of the market), it is unclear what the counterfactual should be in relation to quality and innovation: that is, what the level of quality or innovation would be in a well-functioning market and how that compares to the current level of quality and innovation in cloud services.

Barriers to entry and expansion

28. The evidence we have seen to date suggests that economies of scale which include the need for high levels of capital investment, the importance of the range of services, economies of scope and to a lesser extent network effects represent a significant barrier to entry and expansion in cloud services.

29. In particular:

(a) The supply of IaaS requires significant capital investment in fixed assets. These investments are mainly sunk costs that would not be recovered in full upon exit. While the upfront investments can be reduced by co-locating or leasing data centres, a new entrant or a competitor looking to expand would still need to invest in the servers, components and network equipment for a co-located or leased data centre.

(b) The largest cloud providers are planning significant further investment in their cloud infrastructure, in particular in accelerator chips (eg GPUs), to

[5] See Cloud services market investigation

meet growing demand for AI services.[6] This increases the capital investment required by a new entrant, should they choose to offer customers accelerated compute capacity.

(c) Large cloud providers benefit from economies of scale such as benefiting from bulk purchasing discounts on necessary equipment and from more efficient use of energy associated with large data centres and more efficient use of assets by targeting higher rates of utilisation.

(d) There is evidence that having a large portfolio of cloud services may give cloud providers advantages over their competitors. In particular, the range of first party products is an important factor for customers when choosing which cloud providers to use, there may be some economies of scope in supplying a range of services and there may be some indirect network effects between independent software vendors (ISVs)[7] and customers.

Impact of AI on competition in cloud services

30. We have seen that cloud providers are an important source of accelerated compute capacity for organisations developing and/or deploying Foundation Models (FMs), who in turn are becoming an increasingly important source of revenue for cloud providers and this growth is expected to continue. Partnerships between the large cloud providers and FM model developers are extensive and are likely to play an important role going forwards.

31. Accelerator chips are vital for cloud providers seeking to supply accelerated compute for FM development and/or deployment. A shortage of these chips has meant that cloud providers have been competing to secure them. The three largest cloud providers have also invested in self-supply of accelerator chips.

32. We are still gathering and assessing evidence on this which will enable us to judge whether the growth in AI workloads will allow smaller cloud providers to expand or whether it will result in a further barrier to entry and expansion.

[6] 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.
[7] ISVs build products on services from the cloud providers and compete directly with some of their services

Our emerging view on market power

33. The evidence and analysis set out in the competitive landscape working paper, alongside our other working papers, will inform our assessment of whether one or more cloud providers hold significant market power.[8]

34. Based on the evidence we have seen to date, our emerging view is that there are indicators of significant market power being held by the largest two providers, AWS and Microsoft. This is because:

(a) They both have high shares of supply in IaaS and PaaS and the collective share of other providers is falling;

(b) Potential rivals face significant barriers to entry and expansion, including high levels of capital investment and economies of scale and scope; and

(c) Whilst assessing current market outcomes is complex given the current stage of market development, our profitability assessment indicates that AWS and Microsoft have both been generating returns above their cost of capital.

35. We will continue to investigate the extent to which any such market power is likely to endure into the future.

Our emerging views on the hypotheses we have been investigating

36. Our issues statement set out that we would focus our investigation on four groups of high-level hypotheses, also known as theories of harm, based on both the structure of the market(s) that we will investigate and the conduct of relevant firms within these or other related markets.[9] [10] These are:

(a) Theory of harm 1: technical barriers make switching and multi-cloud harder and limit competition between cloud service providers.

(b) Theory of harm 2: egress fees harm competition by creating barriers to switching and multi-cloud leading to cloud service providers entrenching their position.

[8] The term 'market power' is used to denote the ability of a firm to influence aspects of competition; there are gradations of market power with many firms having limited or transitory market power and others having 'significant market power' which endures over time. CC3 (Revised), Guidelines for market investigations: Their role, procedures, assessment and remedies (publishing.service.gov.uk), paragraph 9.
[9] See paragraph 155 and following CC3 (Revised), Guidelines for market investigations: Their role, procedures, assessment and remedies (publishing.service.gov.uk)
[10] See Issues statement (publishing.service.gov.uk).

(c) Theory of harm 3: committed spend discounts raise barriers to entry and expansion for smaller cloud service providers by incentivising customers to concentrate their business with one provider.

(d) Theory of harm 4: software licensing practices by cloud service providers restrict customer choice and prevent effective competition.

Theory of harm 1: Technical barriers make switching and multi-cloud harder and limit competition between cloud service providers

37. We are investigating whether, and the extent to which, technical barriers prevent or restrict the ability of customers to multi-cloud and switch providers, and which may hinder competition among cloud providers. We are also investigating the extent to which these barriers can be overcome or are inherent in the technology.

38. The evidence we have seen to date relating to technical barriers when using multiple public clouds presents a varied picture. Some customers mentioned general concerns about the difficulty of integrating and operating multiple clouds, including issues such as additional complexity, operational overheads and the differences between clouds. However, some customers said that they face minimal challenges due to the availability of open APIs, cloud-agnostic ISV services and other workarounds.

39. Customers also told us about some specific factors that contribute to technical barriers they face when integrating and operationalising multiple public clouds. These are:

(a) The differences in interfaces of core services;

(b) The differences between how certain cloud infrastructure services integrate with other services from the same cloud provider and how they integrate with ISVs' services (ie asymmetry of integrations);

(c) The differences in interfaces of ancillary services and tools, and in particular IAM services and tools;

(d) The differences in skills required to operate and engineer within different public clouds; and

(e) The latency of connections between different public clouds.

40. The evidence we have seen to date relating to technical barriers when switching presents a clearer picture: it shows that customers experience

technical challenges relating to switching between public clouds, and this can affect their willingness to consider switching and the extent to which switching takes place.

41. The evidence also shows that customers must invest extra effort and resources to mitigate lock-in. Customers consider this as a trade-off when deciding their cloud strategy: some view lock-in as being necessary to gain the benefits of the cloud (eg by procuring highly abstracted proprietary, managed PaaS services), whereas other customers prioritise reducing lock-in as much as possible (eg by deploying and managing their own cloud-agnostic software). Overall, customer or provider-led mitigations may not effectively support efficient multi-cloud and switching for customers.

42. Whilst some specific software, such as Kubernetes and Terraform, as well as proprietary services such as Azure Arc, may help customers use more than one cloud, such software does not fully overcome the challenges of using multiple public clouds and switching between them.

43. The large cloud providers, AWS and Microsoft, told us that they are incentivised to make their public clouds interoperable in order to retain customers that value avoiding a lock-in and/or to win such customers from competing providers. However, it does not necessarily follow that, just because customers would value interoperability between public clouds, cloud providers are incentivised to lower technical barriers for multi-cloud and switching. While lowering such barriers may allow cloud providers to win more incremental customers and/or workloads from their rivals, it may also increase the risk that these cloud providers lose customers and/or workloads to their rivals, or that these cloud providers would have to offer more competitive prices, quality or levels of innovation in order to retain them.

44. Our emerging view is that cloud providers have a complex mix of incentives when deciding whether or not to support multi-cloud and switching. Therefore, in assessing cloud providers' incentive to lower technical barriers to multi-cloud and switching, we will consider the extent to which efforts to facilitate interoperability have already eliminated technical barriers or, conversely, to what extent technical barriers remain.

Theory of harm 2: Egress fees harm competition by creating barriers to switching and multi-cloud leading to cloud service providers entrenching their position.

45. Egress fees are charges to customers when they transfer data out of their provider's cloud, either to an end user or application, when moving data

between the cloud and its on-premises data centres, when moving data between different cloud providers, either as part of a multi-cloud architecture, or as part of switching between two cloud providers. Our focus is on egress fees incurred when transferring data to other cloud providers as customers are likely to do this when switching cloud provider or using multiple clouds.

46. We are investigating whether, and the extent to which, egress fees may make it more costly and difficult for customers to choose the best value offers for them and may deter customers from switching or using multiple clouds. In particular, egress fees are a cost to both switching and using multiple clouds and customers may be harmed when they are deterred from availing themselves of attractive competitive offerings due to egress fees and other costs of switching and multi-cloud.

47. Such costs may further harm competition in the long run where they make it more difficult for smaller rivals to expand, benefit from economies of scale, and compete with larger rivals on a stronger footing. These costs may therefore reinforce or increase the level of concentration in a market, and we are more likely to be concerned where such costs are present in markets that are more concentrated, especially when they are imposed by the larger providers in those markets.

48. In our assessment of this theory of harm, we considered first the prevalence of egress fees, ie the extent to which egress fees are payable as well as the extent to which they have been paid. Where egress fees are applied to only a small subset of customers – by number and by value – this may be indicative, but not probative, of how significant they are or whether they are acting as a barrier to switching and multi-cloud.

49. Our emerging view on the prevalence of egress fees is that, to date, they are common and widespread across all customer use cases: all UK public cloud customers are subject to egress fees if they transfer data out of their cloud provider infrastructure unless the amount of data transferred falls within a cloud provider's free tier.

50. Second, we considered a range of evidence on the role of egress fees in customers' choice to switch or multi-cloud including evidence on the likely magnitude of egress fees that customers would experience when switching or using multiple clouds. Switching costs are more likely to reduce customers' propensity to switch when they are large in magnitude.

51. Based on the evidence we have seen to date, our emerging view on the magnitude and role of egress fees on customers' choice is that egress

fees can be relevant to customers' decision making when switching or using multiple clouds. In particular:

(a) Our estimate of hypothetical 'one-off' switching costs indicates the total financial costs incurred as a result of the egress fees that customers would incur when switching all the data stored with their current cloud providers. These costs can also be materially higher for some customers. This additional cost could make it more difficult for customers to switch cloud provider. In addition, to the extent customers would need to run an integrated multi-cloud architecture for the duration of their switch, egress fees may represent an even higher proportion of their annual cost of cloud.

(b) Some customers identified egress fees as a challenge when considering switching. In addition, some customers said that egress fees had been a challenge to multi-cloud architectures and/or took them into consideration when taking their decision to set up a multi-cloud architecture; and

(c) One cloud provider's internal documents show how data transfer fees could be used, alongside other cloud services, to move customer demand away from capacity constrained regions. Also, another cloud provider's internal document suggests making egress fees free was a way to reduce customer lock-in.

(d) Our data analysis shows that a majority of the cloud providers' UK customers paid egress fees of less than 1% of their annual spend on cloud. However, a substantial minority paid higher proportions. We note that our analysis on current spend on egress fees must be interpreted cautiously. For example, an analysis of actual egress fee spend is likely to understate the importance of egress fees as where egress fees are high, customers may seek to avoid them by deciding not to switch or multi-cloud (or doing so to a lesser extent). In addition, issues with the data, including the inability to distinguish egress to end users means the estimates include costs that are unrelated to switching or the use of multiple clouds.

52. Finally, large cloud providers have told us that egress fees are reflective of costs, although they did not explain how the extent to which egress fees are cost-reflective relates to the question of whether egress fees are a switching cost that could distort competition.

53. Our initial analysis indicates that egress fees may not be reflective of the costs of providing egress for some cloud providers. We are continuing to gather evidence and refine our understanding of this issue.

54. We note that some providers have introduced free switching programmes in the UK in response to the European Union Data Act. We are continuing to gather further information to assess how these voluntary commitments should affect our consideration of the role of egress fees. We will also seek to gather further information on how cloud providers will comply with changes in European or other regulations that are yet to come into force.

Theory of harm 3: Committed spend discounts harm competition by reducing the ability and incentive of rival suppliers to compete for each other's existing customers and/or leading to the weakening or marginalisation of some suppliers, for example because they lose, or fail to achieve, economies of scale.

55. Committed spend discounts or agreements (CSDs or CSAs) are made 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. That is, they are a pricing structure that incorporates a conditional discount.

56. CSAs can be considered 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. One example of how such a conditional pricing structure may raise competition concerns 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

(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.

57. The concern under such circumstances is that the prospect of paying a higher price for the sticky demand 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.

58. Competition may be harmed to the extent that the conduct:

(a) reduces the ability and incentive of rival suppliers to compete for each other's existing customers; and/or

(b) leads to the weakening or marginalisation of some suppliers, for example because they lose, or fail to achieve, economies of scale.

59. Any harm to competition may eventually lead to higher prices or lower quality for customers overall. By virtue of the positions of AWS and Microsoft in the market(s), as outlined above, we consider that any impact on competition arising from their CSAs is likely to be greater than any impact from CSAs offered by smaller providers.

60. In our assessment of this theory of harm, we consider first the prevalence of CSAs. The evidence we have seen to date suggests that customers with a CSA represent a large share of each of AWS and Microsoft's total UK cloud revenues. It also suggests that, while CSAs are not common across all users of cloud services, they are much more common for customers with higher spend.

61. Second, we assess the extent to which CSAs affect customers' choices in relation to the allocation of workloads on public cloud. The evidence reviewed so far also suggests that CSAs have a material impact on the workload allocation decisions of many customers, being it new or existing workloads. It also indicates that CSAs influence, to some extent, the propensity of customers to stay with their current provider and lead customers to spend more than what they would if they did not have to meet the spend commitment.

62. Third, we look at the extent to which the pricing structure of CSAs, that is, how the link between the sticky and the contestable demand is set up in practice, has characteristics that may cause them to harm competition. To do so, we consider the following factors: the share of demand that is 'sticky', the proportion of customers' total demand covered by their CSA commitments, the discount rate provided by the CSA, and some points of context which might aggravate any effects of CSAs on competition, such as the length of the commitment and the existence of any economies of scale.

63. With respect to the share of demand that is 'sticky', our emerging view is that if we were to find that there are high barriers to switching and multi-cloud arising from one or more of technical barriers, egress fees, and Microsoft's licensing practices, then this would strongly suggest that the share of sticky demand is high. Our primary research (the Jigsaw report)

notes that switching cloud providers is a complex and costly exercise. Moreover, several large customers view a material portion of their demand on AWS and Microsoft as sticky.

64. With regard to the proportion of customers' total demand covered by their spend commitment:

(a) customers highlighted increases in length of the CSA contract and increases in committed spend as terms that were proposed as part of CSA negotiations with AWS or Microsoft;

(b) the internal documents reviewed to date are consistent with AWS having the ability to cover a large portion of their customers' demand with their CSAs because they know or estimate individual information about the customers. These documents are also consistent with AWS and Microsoft seeking to increase the portion of their customers' demand covered by their CSAs during negotiations;

(c) our data analysis shows that, for both AWS and Microsoft customers, there is a material portion of customers for which the commitment exceeded customer demand on either AWS and Microsoft (ie where the customer has not met the commitment) and a material proportion of customers for which the commitment covered a large portion of the customer demand on either AWS or Microsoft, although we noted that those customers might have expenditure with other cloud providers as well; and

(d) our data analysis also shows that, for both AWS and, to a larger extent, Microsoft, there is a considerable degree of variation in the discount rates offered for very similar levels of commitment, suggesting a degree of negotiation or bespoke pricing for individual customers.

65. On the discount rate, our analysis shows that increments in discount rate at CSA renewal for AWS and Microsoft are material. However, these figures should be read in conjunction with other analysis, such as the share of demand that is sticky.

66. On the length of CSAs, our data analysis shows that AWS and Microsoft CSAs vary significantly in length but are typically several years long.

67. Finally, we also considered whether there are any potential benefits that may arise from CSAs even if they may give rise to harm to competition.

68. Some cloud providers have told us that CSAs help with investment decisions, although based on evidence to date it is not clear to us that CSAs are the only means of achieving this.

69. We note that the provision of discounts can 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. Therefore, even if CSAs lead to these potential benefits, it does not preclude that CSAs also may give rise to harm to competition.

Theory of harm 4: Software licensing practices by cloud service providers restrict customer choice and prevent effective competition

70. We are investigating the exact nature of cloud providers' software licensing practices and whether these practices disincentivise customers from using rival providers, raise barriers to entry and consequently give rise to concerns about harm to competition.

71. Although we received some submissions relating to Oracle's software licensing practices, we have decided not to prioritise the investigation of these because:

(a) Oracle's market share of cloud is relatively small, 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 of it was secondary;

(b) We 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 have prioritised the use of our resources on the areas where there is the potential for greater harm to arise.

72. In light of the above, we have focused our investigation on Microsoft's software licensing practices. These practices may be more likely to harm competition in the markets for cloud infrastructure services if:

(a) the licensing practices relate to software products where Microsoft has market power, such that customers would find it difficult to switch away from them;

(b) Microsoft's rivals in providing cloud infrastructure services do not have an effective counter strategy; and,

(c) Microsoft's software products are provided at a higher price or lower quality to customers that choose one of Microsoft's rivals to be their cloud provider which may weaken competition between cloud providers.

73. Competition may be harmed such that it leads to foreclosure, whereby rivals are forced to exit from the market, are prevented from entering, or are materially disadvantaged, such that they consequently compete less effectively.[11]

74. Below we summarise Microsoft's licensing practices, our emerging views on the extent to which Microsoft has market power in software products and the extent to which the licensing practices may affect customer choice of cloud provider.

75. Based on the evidence we have seen so far, Microsoft appears likely to have a significant degree of market power in relation to its supply of the following products: Windows Server, Windows 10/11, SQL Server, Visual Studio and its productivity suites (these are the Microsoft 365 packages).

76. We received submissions setting out issues and concerns with regard to Microsoft's software licensing practices relating both to price and non-price factors.

77. Our analysis so far has focused on the potential impact of the price factors on consumers' choice of cloud provider. We are undertaking data analysis that seeks to estimate:

(a) the implied difference in the licensing costs for Windows Server and SQL Server on Azure compared with AWS or Google, if any;

(b) the proportion of cloud customers that use each of Windows Server, SQL Server, and MS365 or Office apps; and

(c) the relative usage of Windows Server and SQL Server on Azure compared with AWS and Google.

78. Evidence we have seen to date from customers on whether Microsoft's licensing practices had an impact on their choice of cloud provider shows that:

[11] CC3 (Revised), Guidelines for market investigations: Their role, procedures, assessment and remedies (publishing.service.gov.uk), paragraph 269.

(a) most customers we spoke to understand that using Microsoft software products is cheaper on Azure;

(b) the cost or ease and/or ability to use licences are selection factors for many customers we spoke to, and some particularly consider whether they can make use of their existing investment in licences in their choice of cloud provider;

(c) a few customers we spoke to do not consider licensing to be an important factor in their choice of public cloud. (For some of these customers, licensing was not relevant for their use cases.) and

(d) existing skills and familiarity with the Microsoft ecosystem were also very important selection factors for many Azure customers we spoke to.

79. The Jigsaw report found that participants considered that Microsoft's software licensing practices were not, on their own, the influencing factor in their choice of public cloud provider. The report particularly highlights that the original take up of Azure was often closely related to customers' pre-existing use of Microsoft products. However, some participants considered Azure as the natural choice for both technical and financial reasons and licensing terms appear to contribute to this.

80. We will consider our data analysis, once complete, in the round with the other relevant evidence.

The AEC test

81. We have not yet reached any provisional conclusions on whether there is a feature or combination of features that prevents, restricts or distorts competition in connection with the supply or acquisition of public cloud infrastructure services in the UK or a part of the UK. Therefore we have not yet reached any provisional conclusion on whether or not there is an adverse effect on competition (AEC) in cloud services.[12]

82. In coming to a provisional view, we will seek to establish whether or not any feature, or any combination of features, can be expected to harm competition when measured against a theoretical benchmark.[13]

83. Our guidelines state that a market feature may be intrinsic to the structure of the market, including high levels of market concentration and high barriers to entry, or may arise from the conduct of any market participant

[12] CC3 (Revised), Guidelines for market investigations: Their role, procedures, assessment and remedies (publishing.service.gov.uk), paragraph 319.
[13] CC3 (Revised), Guidelines for market investigations: Their role, procedures, assessment and remedies (publishing.service.gov.uk), paragraph 154

(whether supplier, acquirer or customer and whether or not in the reference goods or services market).[14]

84. This means that, in reaching our provisional conclusions, we will consider both our assessment of the competitive landscape and of the four theories of harm we have been investigating, all of which relate to the conduct of market participants.

85. In the absence of a statutory benchmark, we use the benchmark of 'a well-functioning market' as set out in our guidelines. A well-functioning market is one that displays the beneficial aspects of competition, rather than an idealised, perfectly competitive market. The benchmark will generally be the market envisioned without the features that are identified as harming competition. 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) or if the nature of competition in the market is defined by arrangements put in place by government.[15]

86. In well-functioning cloud services market(s), we would not expect every customer to split its workloads across multiple providers in a highly integrated manner, or to switch provider every year. Rather, we would expect customers to be able to choose between a range of alternatives and to be able to multi-cloud and switch between products/providers. We note that customers may still face some sources of friction when exercising their choice of cloud provider, even in a well-functioning market, due to any intrinsic features.

87. Lower barriers to multi cloud and switching would be expected to enable:

(a) customers to make effective decisions, readily access the most suitable products for their needs (including via switching or using more integrated multi-cloud), boost their bargaining position in relation to providers and increase their access to a diverse range of innovative products.

(b) providers to effectively compete for their rivals' customers and workloads, rewarding those who are best able to win on the merits of their services. Innovation may increase as the benefits of innovation might be higher because potential demand for new innovative products could be greater.

[14] CC3 (Revised), Guidelines for market investigations: Their role, procedures, assessment and remedies (publishing.service.gov.uk), paragraph 155 and 157.
[15] CC3 (Revised), Guidelines for market investigations: Their role, procedures, assessment and remedies (publishing.service.gov.uk), paragraphs 154, 30 and 320.

88. We would also expect lower barriers to entry and expansion. We might expect these to enable providers to enter, including with innovative new products and services, and expand if they are meeting the demands of customers. We might also expect providers to be able to compete for sub-sections of a customer's total workloads if switching and multi-cloud were easier to achieve.

89. Given that ISVs are both customers and providers in cloud services, they might contribute to a well-functioning market by perceiving and exploiting new market opportunities. For example, in a well-functioning market where there are lower barriers to switching and/or more integrated forms of multi-cloud, ISVs may be in a position to generate greater rewards for bringing new and innovative products to market, as they could compete effectively with integrated cloud providers at a PaaS level. ISVs might also be incentivised to build services on a wider range of providers' infrastructure (or develop cloud-agnostic services) as the cost of integrating with smaller providers could more readily be recouped if those smaller providers were able to grow and compete more effectively with larger providers. As ISVs integrate with a wider range of cloud providers this could spur competition between those cloud providers further.

90. In summary, where the demand and supply side interact effectively in a well-functioning cloud services market, the benefits of competition would be unlocked. These benefits can include lower prices, better quality, a broader range of services offered, better service and higher levels of innovation.

Potential remedies

91. If we decide that that there is a prevention, restriction or distortion of competition in the cloud services market, we will have found an 'adverse effect on competition' (AEC).[16] This would trigger a duty to decide whether the CMA should take action and if so what action should be taken, and/or whether it should recommend that others take action, to remedy, mitigate or prevent 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.[17]

92. Working papers on each theory of harm set out some considerations on potential remedies relating to each area and we have also published a

[16] As defined in section 134(2) of the Act
[17] As defined in section 134(4) of the Act

working paper on remedies alongside this document which sets out some collective considerations surrounding a potential package of remedies.[18]

Responding to this statement

93. Any party wishing to respond to this updated issues statement should do so in writing by 27 June 2024 by emailing CloudMI@cma.gov.uk.

94. We will hold hearings with AWS, Microsoft and Google to discuss the issues and potential remedies set out in this statement and we will publish summaries of those hearings.

95. We will publish a provisional decision report containing our provisional decision on whether there is one or more AECs; if we do find AEC(s), then we will set out our provisional decision on remedies.

96. Our administrative timetable can be found on the inquiry case page Cloud services market investigation.

6 June 2024

[18] See Cloud services market investigation

Annex 1: Background to the investigation

1. On 5 October 2023 the Office of Communications (Ofcom), in exercise of its powers under sections 131 and 133 of the Enterprise Act 2002 (the Act), made a reference for a market investigation into the supply of public cloud infrastructure services in the UK.[19]

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

3. For the purposes of the reference:

   (a) 'Cloud infrastructure services' means services that provide access to processing, storage, networking, and other raw computing resources (often referred to as infrastructure as a service, IaaS) as well as services that can be used to develop, test, run and manage applications in the cloud (often referred to as platform as a service, PaaS).

   (b) 'Public cloud computing' means a cloud deployment model where cloud services are open to all customers willing to pay, and computing resources are shared between them.

   (c) 'Multi-cloud' means a cloud deployment model involving the use of more than one public cloud provider by a single customer, where multiple clouds may or may not be integrated with each other.[20]

4. The CMA is 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.[21] The decision is made by a group of independent members constituted from its panel, on behalf of the CMA.[22] If the group decides that there is such a prevention, restriction or distortion of competition, it will have found an 'adverse effect on competition' (AEC).[23]

5. If the CMA finds that there is an AEC, it has a duty to decide whether it should take action and if so what action should be taken, and/or whether it should recommend that others take action, to remedy, mitigate or prevent

[19] Terms of Reference (ofcom.org.uk)
[20] Terms of Reference (ofcom.org.uk)
[21] Section 134 Enterprise Act 2002 (legislation.gov.uk)
[22] Cloud services market investigation Kip Meek (Inquiry Chair), Robin Foster, Paul Hughes and Colleen Keck.
[23] As defined in section 134(2) of the Act

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.[24]

[24] As defined in section 134(4) of the Act