# CMA CLOUD SERVICES MARKET INVESTIGATION

# MICROSOFT RESPONSE TO THE PROVISIONAL DECISION REPORT


**24 February 2025**

Linklaters
LATHAM&WATKINS

## Table of Contents

**Contents** **Page**

# Cloud Services Market Investigation

# Microsoft Response to the Provisional Decision Report

## 1 Executive Summary

### 1.1 Introduction

(1) Following its release, we have closely studied the Provisional Decision Report (the "**PDR**")[1] on the state of competition in the market for cloud services. It should come as no surprise that we disagree with substantial portions of the PDR.

(2) It does not reflect how the cloud computing market operates in practice. Instead of following the evidence, it grounds much of its analysis in hypothetical scenarios and focuses on largely theoretical issues. We are particularly concerned about its singling out of Microsoft and its accusation that Microsoft is unfairly using some of its software products to prevent Amazon Web Services ("**AWS**") and Google Cloud from competing effectively for UK customers.

(3) At the start it is notable that UK customers – businesses, enterprises, public sector bodies, non-profit organisations – do not seem to share the PDR's desire for intervention. As a group, UK customers, the real (ultimate) concern of the Competition and Markets Authority ("**CMA**"), have raised limited complaints about the intensity of competition in the cloud computing market over the course of two extensive and long-running investigations.

(4) While AWS clearly believes it is entitled to license all of Microsoft's software for its own benefit and on favourable terms (even though AWS provides none of its own software to Microsoft or anyone else), it has not, to our knowledge, complained about an inability to compete effectively. Nor could it. After all, AWS has consistently claimed the number one spot in cloud computing globally and in the UK for more than two decades. And the CMA seems to agree, given its provisional conclusion that AWS is too profitable and should be designated as a provider with Strategic Market Status ("**SMS**") in cloud computing.

(5) Google, on the other hand, has complained, at least in its statements to the CMA and competition authorities elsewhere. To other audiences, Google celebrates how fast its cloud business is growing. Take for example Alphabet CEO Sundar Pichai's recent statement to Wall Street analysts: "*We've seen incredible momentum in the Cloud business… In Q4, Cloud was at a $36 billion annual revenue run rate — that's five times our run-rate five years ago*".[2]

(6) With results like that, it bears reflecting whether the CMA must intervene to enable Google's growth in the UK market by softening competition from its competitors, particularly when these practices raise issues only among Google's lawyers and the smaller cloud computing providers in the so-called "trade associations" Google has formed and funds.

(7) That Google is talking out of both sides of its corporate mouth is troubling, and we will speak to it below, when we address the issues of egress fees, interoperability, and legacy software pricing. But, first, it is important to talk about what the PDR fundamentally misjudges: the AI technology driving the competitive dynamic in cloud computing today.

---

[1] Available here: Cloud Infrastructure Services: Provisional Decision Report (28 January 2025)

[2] Sundar Pichai, "Cloud Next 2024: More momentum with generative AI," Google Cloud Blog (9 Apr 2024) (available at https://blog.google/products/google-cloud/google-cloud-next-2024-generative-ai-gemini/#models-expanded-access)

### 1.2 Focusing on the future

(8) Since the bounds of this inquiry were set out by Ofcom in July 2022, every technology market, including cloud computing, has been upended by the explosive emergence of generative AI, prompted by the launch of ChatGPT.

(9) While the PDR touches briefly and inconclusively on one small element of this, the provision of accelerated compute by cloud computing providers to foundation model developers,[3] it misses the big picture. On both the supply and demand sides of the cloud computing market, generative AI has rapidly become a primary driver of strategic planning and investment and competition.

(10) This year alone, cloud computing providers will invest more than $250 billion[4] globally to expand and upgrade their infrastructure and meet customer demand for capacity, not only to train and deploy AI models, but also for their use at scale. And they are rapidly innovating and building, both individually and with a host of partners, new AI platform services and development capabilities.

(11) On the demand side, customers of all shapes and sizes are exploring how they can tap into the potential of AI to amplify the capabilities of their existing cloud applications and solutions and to transform their business processes, making strategic choices about which cloud platforms best meet their needs. Even Prime Minister Keir Starmer is looking to "*mainline AI into the veins of*" the UK.[5]

(12) The wave is lifting providers old and new. At the end of 2024, Oracle CEO Safra Catz reported that "*Record level AI demand drove Oracle Cloud Infrastructure revenue up 52% in Q2, a much higher growth rate than any of our hyperscale cloud infrastructure competitors*".[6] This past year, cloud computing startups with infrastructure optimised for AI workloads, such as CoreWeave, Lambda Labs, and Crusoe, have closed multi-million dollar funding rounds at multi-billion dollar valuations on the back of exponential growth. And many companies – e.g. ServiceNow, CyrusOne, CloudHQ, CoreWeave, and Nscale – have announced plans and investment to expand their data centre footprint in the UK.

(13) Microsoft is investing and partnering, to be at the forefront of AI research and integration into its products and services and to help its customers accelerate their AI transformation.[7] At the same time, Google is a widely acknowledged leader in the early stages of this technology, which it employs widely in its Search and Advertising businesses, as well as its cloud computing business. Mr. Pichai has been vocal about his company's unique advantages: "*First, AI infrastructure. Our sophisticated global network of cloud regions and data centers provides a powerful foundation for us and our customers, directly driving revenue. We have a unique advantage because we develop every component of our technology stack, including hardware, compilers, models and products. This approach*

---

[3] PDR, paragraphs 3.506-3.510

[4] See paragraph (67)(i)(iv) of this response below

[5] See Prime Minister sets out blueprint to turbocharge AI - GOV.UK

[6] "Oracle Announces Fiscal 2025 Second Quarter Financial Results," Oracle Investor News (9 Dec 2024) (available at https://investor.oracle.com/investor-news/news-details/2024/Oracle-Announces-Fiscal-2025-Second-Quarter-Financial-Results/default.aspx)

[7] For example, Microsoft's agreement with LSEG is leading to the co-development of highly innovative financial services data and analytics solutions. The project will result in up to $2.3 billion of investment in cloud by LSEG's top 250 customers and overall cloud investments of $135 billion across other industries over the next five years.

*allows us to drive efficiencies at every level, from training and serving, to developer productivity*".[8]

(14) And Thomas Kurian, Google Cloud's CEO, has been clear that AI represents a key opportunity to grow even more market share: "*The biggest market opportunity we see is moving customers from AI pilots and proofs of concept toward larger implementations*". Mr. Kurian said recently about Google Cloud's growth: "*Google Cloud partners are now seeing AI infused into virtually every customer engagement*".[9]

(15) If the subject is competition, the current and future focus of the market could not be clearer.

### 1.3 The new technology is holistic

(16) Rather than acknowledge that these developments and eye-popping investments by multiple providers demonstrate that competition is functioning well, the PDR concludes that AI infrastructure has limited relevance, is distinct from cloud computing, and these developments present, at most, a distant and uncertain prospect.[10] This is a fundamental mistake affecting the diagnosis in the PDR of the (alleged) problems in the cloud computing market.

(17) There are two essential points of focus here. First, AI infrastructure (i.e. accelerated compute) plays an essential role in powering the training, deployment and use of foundation models, but those processes also require compute, storage and networking, along with many other cloud computing services. It cannot be separated out and the PDR is wrong to segment the market and suggest that these products / services are not inter-related. Customers make purchasing decisions based on all the functionality they will need for their cloud deployments, and AI capabilities are quickly becoming the key functional consideration.

(18) Second, from commonly used business analytics to the Machine Learning AI that already pervades business computing, to the most advanced generative AI, rich and diverse data is as critical a component as any software or hardware. Google, with one of the world's richest data repositories, obtained in all sorts of ways, is an exceptionally strong, fast-moving competitor.

(19) By essentially ignoring such competitive dynamics, the PDR ends up looking backwards, and focusing on issues largely peripheral to cloud computing competition today: (i) egress fees; (ii) interoperability standards; and (iii) Microsoft's pricing of legacy software to Amazon and Google.

(20) While it is hard to see how intervention in these three areas will meaningfully boost competition in cloud computing (or, given Amazon's and Google's impressive growth over this period, how those elements ever hindered competition), here we will turn to the lack of merit in the PDR's concerns. Not simply because we believe that the CMA's provisional understanding needs deep revision, but because there is a real danger that intervening in the market based on these misunderstandings will backfire, leaving the UK with the opposite of the CMA's goal of a healthy, well-functioning market, rich in growth and investment.

---

[8] Alphabet 2024 Q4 Earnings Call Transcript (4 Feb 2025) (comments by Sundar Pichai) (available at https://abc.xyz/2024-q4-earnings-call/)

[9] https://www.crn.com/news/cloud/2024/google-cloud-ceo-on-huge-investments-ai-and-challenges-in-2024

[10] PDR, paragraphs 3.507 and 8.19

### 1.4 Egress fees

(21) One year ago, AWS, Microsoft and Google Cloud removed egress fees, worldwide, for customers switching from one cloud computing provider to another, prompted by impending requirements of the European Union Data Act (the "**EU Data Act**").[11] That this change did not lead to a material uptick in switching rates should give the CMA serious pause about rushing to further intervention and extending its regulatory regime beyond that of the European Union.

(22) There is a good reason why egress fees do not drive customer decisions about switching or about the number of different clouds that they use. For almost all customers, egress fees, as the PDR implicitly acknowledges, are not "*the main concern*" or a "*relevant factor*".[12]

(23) When making decisions, customers are focused on the various services available and their inherent cost, plus the implicit costs of training, operations, time, and security. Against these decisions, the size of an egress fee, which is necessary to recoup capital outlays by many providers, plays at best a very limited, if any, role. The CMA's own customer research confirms this, observing that "*Egress fees were rarely top of mind when participants were asked about their openness to switching cloud provider or using multiple clouds, or even when specifically asked about potential challenges for a switch of multi-cloud strategy*".[13]

(24) While the fees themselves are low and largely inconsequential to customer cloud choices, real harm exists in the PDR's suggestion that the CMA intervene to eliminate egress fees, over both the public internet ("**standard egress**") and private premium networks ("**premium egress**").[14] The PDR does a disservice by summarily brushing aside the real concerns shared directly with the CMA by cloud computing providers[15] – and their business executives – in favour of the PDR's "*analysis of hypothetical scenarios*" where it speculates "*customers may incur substantial costs from switching and multi-cloud*".[16]

(25) Private premium networks provide superior network performance for customers whose specialised needs demand it and ease congestion on the public internet by offlaying traffic. They are the product of billions of dollars of investment globally – including subsea cables, terrestrial fibre optic cables, and dedicated edge nodes – on the part of many providers, including Microsoft.

(26) As is the case on the public internet, cloud computing providers do not charge ingress fees on private premium networks. Competitive practices have driven out that charge, with no government intervention. The market has instead accommodated the egress fees that pay for the networks. If the CMA artificially prohibits cloud providers from charging egress fees on private premium networks, cloud computing providers will be faced with difficult choices regarding the care and growth of this essential infrastructure for UK customers.

(27) If customers perceive that private premium networks are free, data transfer on these systems will increase, crowding those networks even as their owners struggle to pay for their upkeep. Cloud computing providers will need to raise prices elsewhere, most likely with indiscriminate charges on other services that extend to all cloud customers regardless of whether and how

[11] 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

[12] PDR, paragraph 5.421

[13] Cloud Services Market Investigation Qualitative Customer Research Final Report (May 2024), paragraph 1.4.12

[14] PDR, W.160-W.161

[15] PDR, paragraph 5.508

[16] PDR, paragraph 5.422

much they use the network. And the incentive for cloud computing providers to offer "premium" experiences will disappear, collapsing network choices around the lowest common denominator. In the ensuing confusion, both UK infrastructure and customers will be worse off.

### 1.5 Interoperability standards

(28) The PDR indicates that the CMA could designate and require at least AWS and Microsoft, if not more cloud computing providers, to implement support for certain industry standards in one or more aspects of their cloud services to improve interoperability and presumably ease switching or multi-cloud usage.[17]

(29) Although we believe that standards can play an important role in healthy and competitive technology markets, including cloud computing, we see real challenges and complexity in supplanting what is a market-driven process with a CMA-mandated one.

(30) Today, there are more than 1,200 active information technology standardisation organisations,[2] ranging from small industry special interest groups to formal international standards organisations comprised of representatives from 100+ countries. Many of the technologies that positively affect daily life are implementations of standards created by the tens of thousands of technical experts who participate in the technical committees of these organisations. Cloud computing also incorporates many standards, including for networking, data exchange between diverse systems, and interfaces that enable information security practices, and support of common programming languages. Microsoft has hundreds of employees engaged across a spectrum of standardisation activities and implements many thousands of standards across its product portfolio, including in its Azure cloud services.

(31) There is a second pervasive method of market-driven interoperability for cloud services, and that is through open-source software. Essential cloud service technologies common to all modern cloud providers such as containers, Kubernetes orchestration, and Prometheus real-time cloud metrics have no written standard, yet are norms that promote interoperability. Much like standards, participation in the open-source software projects and implementation of the final solution are voluntary. Yet marketplace demands by customers drive cloud provider adherence to the normalised technologies.

(32) To its credit, the PDR acknowledges "*that standardisation remedies in cloud would typically be complex, technical and likely to need continuous oversight and refinement both in terms of the scope of the remedy (eg which services are included) and the specific standards (eg maintaining the design of the technical standard itself)*".[18] Although it concedes that a market investigation order cannot overcome these challenges,[19] it suggests that the CMA's digital markets competition regime would have the capability to do so. But the problem is not the scope of powers that one regime or another can bring to bear or the duration of its mandate, it is the lack of domain-specific expertise and foresight in this technically complex area.

(33) The CMA, including the Digital Markets Unit ("**DMU**"), is fundamentally an organisation of competition lawyers, and economists. It is not well-equipped to dictate a critical innovation roadmap for a fast-moving and technically demanding global market. While the CMA counts technologists among its staff, those technologists are generalists. They do not have specialised experience building and operating a public cloud at scale or across the multitude

[17] PDR, W.18-W.74

[18] PDR, W.68

[19] PDR, W.69

of services offered (e.g. security, analytics, identity, management, databases). Nor do they have a broad base of customer knowledge or insight into breakthrough technology work inside the many companies in the marketplace. On what sort of timeline would the CMA be dictating technical specifications for cloud computing providers? What will such a constraint mean for the competitive dynamics if it is applied asymmetrically, leaving certain providers, like Google Cloud, free to operate unhindered?

(34) It is critical to get these choices right. Technology moves forward fast. Without consistent and sustained evaluation of the state of the art, the CMA will likely create a situation where they are mandating old technologies that are not globally relevant and/or will force cloud computing providers and their customers to maintain unnecessary parallel versions of technology to meet both statutory and market factors simultaneously. The long-term implications of the failure to keep pace will likely result in impaired technical compatibility, additional costs for all cloud providers, raising prices for consumers, and barriers to innovation for the UK marketplace. The most active model of compulsory standardisation is the "GB" standards system of China which has long presented industry (both domestic and foreign) with outdated requirements and/or additional costs that do not modernise at the pace of the marketplace.

(35) Standardisation is not the same for all critical technologies in that some, such as mobile communications, require standards to be completed ahead of any marketplace activity. Others, such as artificial intelligence, has been entirely the opposite as the innovations and early marketplace activities have no standardisation but the work on responsible behaviour assurance has been both essential and can only be completed after there is an understanding of what the technology is. Cloud computing is a mix of these two with no obvious decision framework that would suggest the CMA is promoting interoperability or mandating unnecessary resource allocation to outdated technologies.

### 1.6 Pricing legacy software for hyperscale customers

(36) The PDR takes aim at how Microsoft makes its legacy software available for use by AWS and Google Cloud, concluding that the prices charged "*partially foreclose*" AWS, the clear market leader, and Google Cloud from competing in the cloud computing market.[20] We will delve into the different aspects and details of the claims made by Amazon and Google in the sections below. But at the outset, it is important to take a step back.

(37) In short, Amazon and Google are asking the CMA to intervene and constrain the price Microsoft can charge them when they use Microsoft's software to build and sell cloud services to their customers and the price Microsoft charges its clients on Azure for the use of Microsoft software. This is an extraordinary and unprecedented intervention, riding roughshod over Microsoft's intellectual property ("**IP**") rights. No other software provider in the industry would be subject to similar limitations.

(38) To what end? The PDR acknowledges that the only beneficiaries of this remedy are Amazon and Google.[21] But they do not need the CMA's help to compete against Microsoft. A report by CNBC on the same day that the PDR was released makes this fact crystal clear: "*[Microsoft]'s year-over-year Azure cloud growth has been little changed for the past few quarters, while Alphabet and Amazon have shown acceleration. . . Since the last quarter of 2023, Azure growth has increased by 2 percentage points. Meanwhile, top rivals Amazon*

[20] PDR, paragraph 6.543

[21] PDR, paragraph 9.78 and W.246

*and Alphabet have seen cloud growth over that stretch accelerate by 7 points and 13 points, respectively*".[22]

(39) An extraordinary remedy may be appropriate in an extraordinary circumstance. But this is not that.

(40) Owing to our legacy software business, we have many customers who, over the years, have paid and are paying licence fees for Microsoft products (e.g. Windows Server and SQL Server) that support enterprise workloads on servers in their private data centres. When they are ready to move those enterprise workloads and modernise them for the cloud, we want them to choose Microsoft Azure. But they have a lot of choices, because our software is also available on other public clouds, including AWS and Google Cloud.

(41) We compete to win their business in a variety of ways, including competitive pricing. One way we let them know that we intend to offer attractive prices in our marketing efforts is to provide a standing offer to discount Azure and offset the portion of the price attributable to Windows Server and SQL Server software when customers deploy workloads on Azure that rely in part on that software (also known as "Azure Hybrid Benefit"). This is not foreclosure; this is the essence of competition, and it benefits UK customers.

(42) Now, compare that with the stance taken by AWS and Google Cloud on their popular proprietary products. While Microsoft has Windows Server and SQL Server, Amazon has S3 (Simple Storage Service), Aurora (Relational Database Management System), Dynamo DB (serverless document database), and more. Google has BigQuery (data warehouse), Google Analytics, Looker (business intelligence and data analytics platform) and more. Do AWS and Google license their proprietary software to their competitors, at any price at all? They do not.

(43) You can see how we might feel unfairly singled out by the PDR in this matter.

### 1.7 There they "Google" again

(44) Given what we wrote above about Google crowing about its growing market share even as it complains to regulators about constraints, it is no surprise that there is similar hypocrisy on Google's part about the source of its growth. To the CMA, it is about the migration of workloads that rely on Microsoft software. To the rest of the world, it is all about the new world of AI.

(45) In June 2024, Google told the Market Investigation Panel: "*[W]e remain of the view that market dynamics will continue to be driven primarily by customers migrating existing on-premises workloads to the cloud in the short and medium term*".[23] Yet two months before, Alphabet CEO Sundar Pichai spoke of "*incredible momentum in the cloud business*" thanks to Google being "*at the forefront of the AI platform shift*".[24]

(46) Speaking to industry analysts at the same time, Google Cloud CEO Thomas Kurian said customers were choosing based on business transformation, not legacy adaptation. He added, "*the basis of how they thought about it in 2007, '08, '09 was all about "How do they help me either go faster by building apps or reduce my cost of data centers by allowing me to lift and shift workloads?" Now, they're thinking about it in a different way. They're looking at it as, "Can I use AI to transform my business? Who's got the best platform and tools to*

[22] See Microsoft's underperformance has investors looking to cloud for growth

[23] Google Response to the CMA's Competitive Landscape Working Paper (14 June 2024)

[24] See Google Cloud Next 2024: Gemini and generative AI updates

*help me do that?" Once we get them to use the AI platform and tools, it does drag in many of our other services*".[25]

(47) This is not merely an investment of time and words. Google has sunk thousands, if not millions, of dollars to stand up and support "trade associations", who unsurprisingly parrot its claims to competition authorities and government officials.[26] This false echo chamber has distorted the analysis, implying that Microsoft's software licensing and pricing impacts cloud providers beyond Amazon and Google. It does not.

(48) Google has been particularly persistent in weaponising local cloud providers on its behalf. In June 2024, it offered members of one of those organisations – CISPE – a package valued at over $500 million to scuttle a proposed agreement with Microsoft and continue its advocacy against how we price our software.[27] Its overtures were unsuccessful. Notably, for a company that claims it cannot match the discounts provided by Microsoft to customers, the amount offered to CISPE members was *more than Google pays Microsoft annually* to license Microsoft software to build and provide services to its customers.

(49) When its efforts with CISPE failed, Google stood up an alternative trade association: the Open Cloud Coalition.[28] Following its launch in November 2024, the coalition wasted no time, publishing a paper directed at the CMA and arguing that Microsoft's "*unfair software licensing practices*" hurt the industry. It did so even though, as the PDR noted, those practices do not impact any of its members (many of whom do not license Microsoft's legacy software products) apart from Google.[29]

### 1.8 The problem of (no) evidence, and the future

(50) Setting aside these misleading behaviours, we struggle to see any evidence of foreclosure in the marketplace, either from AWS, Google or the PDR.

(51) AWS has been the market leader in public cloud since its launch in 2006. AWS serves millions of instances running Microsoft software to thousands of customers. Its website boasts: "*Customers have been running Microsoft Workloads on AWS for over 16 years, longer than any other cloud provider. Customers choose AWS because we have the most experience with Microsoft applications in the cloud and we offer the best platform for Windows Server and SQL Server for higher performance and reliability, greater security and identity services, more migration support, the broadest and deepest capabilities, lower total cost of ownership, and flexible licensing options*".[30]

(52) In a similar vein, six months ago, in the midst of this market investigation, AWS proclaimed that "*a clear majority of Microsoft SQL Server users chose AWS as their primary cloud*

---

[25] "An Interview with Google Cloud CEO Thomas Kurian About Google's Enterprise AI Strategy," Stratechery (11 Apr 2024) (available at https://stratechery.com/2024/an-interview-with-google-cloud-ceo-thomas-kurian-about-googles-enterprise-ai-strategy/)

[26] *See, e.g.*, The Computer & Communications Internet Association (CCIA): Google and Amazon are members of the organisation and representatives of those companies serve on its Board of Directors; Microsoft is not a member) (https://ccianet.org/about/board-of-directors/)

[27] See Kunert, Paul, "Google offered millions to ally itself with trade body fighting Microsoft," *The Register* (28 Nov 2024) (available at https://www.theregister.com/2024/11/28/google_offered_millions_to_cispe/); Stolton, Samuel, "Google Offered €470 Million to Derail Microsoft Antitrust Pact," *Bloomberg* (15 July 2024) (available at https://www.bloomberg.com/news/articles/2024-07-16/google-offered-470-million-to-derail-microsoft-antitrust-pact).

[28] See https://opencloudcoalition.com/membership/

[29] PDR, paragraphs 6.539 and 6.542

[30] Microsoft workloads on AWS (available at https://aws.amazon.com/windows/?blog-posts-content-windows.sort-by=item.additionalFields.createdDate&blog-posts-content-windows.sort-order=desc)

*provider when polled. In fact, data shows AWS is the top cloud provider of choice (52%) – more than the next six cloud providers combined*."[31]

(53) For more than five years, Google Cloud has consistently grown revenue and operating margin. In this way, it is following the curve of Microsoft and AWS, in which heavy investment in early years yields profitability later on. It is now the second largest business – behind Google Search – in Alphabet's portfolio. Thus, it is disingenuous of Google Cloud to claim it is somehow hampered in the UK market by having to pay for the current licensing charges for the use of Microsoft software, particularly when it won't offer BigQuery, a forward-looking AI product, to AWS or Microsoft.

**Figure 1: Google Cloud Platform: Revenue and Margin**

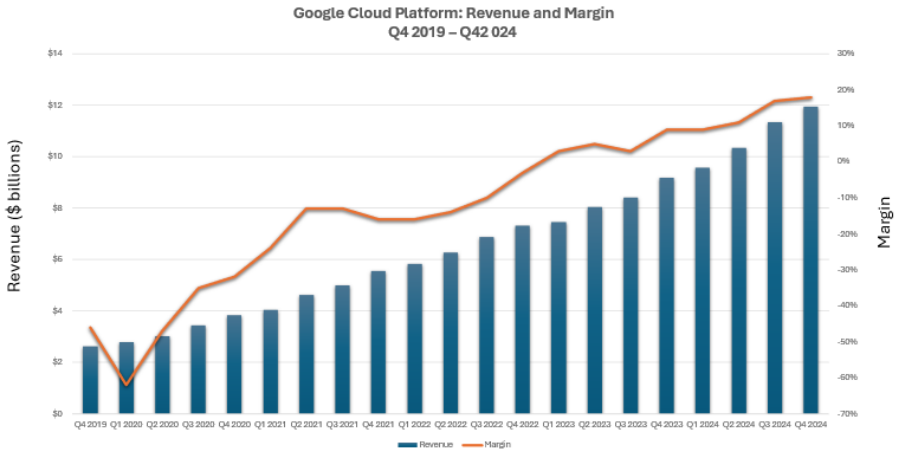


(54) Current and future competition in the cloud simply does not and will not turn on the price of Microsoft's legacy software. Popular as it may be on PCs and on back-office servers, open-source software dominates the public cloud. Most customer workloads in the cloud run on Linux, not Windows Server. And even those workloads that do incorporate Microsoft software also must include several other services (e.g. storage, networking, and many others), giving AWS and Google Cloud more than enough margin to profitably compete on price against Microsoft while absorbing the cost of Microsoft software.

(55) In addition, AWS and Google Cloud have responded – as one would expect in a healthy and competitive market – using different strategies and tactics to blunt the impact of Microsoft's licensing charges on their balance sheets. AWS has developed a wealth of training, content, and programs along with a network of partners to help customers assess, migrate, manage, and modernise their Microsoft workloads on AWS.[32] In 2023, AWS launched 108 product features for customers running workloads incorporating Microsoft software, including features aimed specifically at driving down licensing costs and, in the case of Amazon Q,

[31] Frank Wang, "New research shows AWS is the cloud provider of choice for SQL Server," Microsoft Workloads on AWS Blog (01 Aug 2024) (available at https://aws.amazon.com/blogs/modernizing-with-aws/new-research-shows-aws-is-the-cloud-provider-of-choice-for-sql-server/)

[32] See A Wat, C Dinjor, and K Lee, "Accelerate Microsoft Workload Modernization on AWS with Softchoice," AWS Partner Network Blog (14 Nov 2024) (available at https://aws.amazon.com/blogs/apn/accelerate-microsoft-workload-modernization-on-aws-with-softchoice/); and R Krishnan and C Brum, "Driving cost optimization and modernization of Microsoft workloads on AWS with SoftServe's application modernization platform," AWS Blog (4 Jan 2024) (available at https://aws.amazon.com/blogs/modernizing-with-aws/driving-cost-optimization-and-modernization-of-microsoft-workloads-with-softserve/)

leveraging generative AI to refactor. NET Framework-based applications to run on Linux.[33] In addition, AWS has invested continuously to improve the performance of its first- and third-party processors to deliver superior efficiency and consequently drive down customer licensing costs.[34] As a result, AWS claims that customers migrating to AWS can reduce the cost of Windows Server licences by an average of 77% and SQL Server licences by an average of 45%.[35]

(56) Google Cloud has taken similar steps.[36] It also benefits from its search and digital advertising services, which are essential to almost all businesses and enterprises, as well as other popular services. We know well from our own experience and fact finding by the U.S. House of Representatives and the press that Google Cloud often taps those businesses, offering customers ad credits and other incentives to win deals.[37] At the end of the day, customers want the best value, whatever the combination of gives and gets delivers it. AWS and Google Cloud have many different and unique tools in their arsenal to win UK customers.

### 1.9 Conclusion

(57) Microsoft has been a major investor in the UK economy for more than 40 years. Today, we are investing in digital infrastructure, working to enable the AI transformation of thousands of UK businesses, public sector organisations and non-profits, and helping to train the next generation of cloud and AI talent. We have consistently demonstrated our commitment to making the UK a more attractive destination for investment, a place where businesses can thrive, and where talent can flourish.

(58) Where there have been regulatory concerns in the past, we have worked constructively to find proportionate solutions that both address competitive dynamics and ensure that markets

---

[33] T Dull, "Year in review 2023: AWS for Microsoft workloads," AWS Blog (20 Dec 2023) (interview with Rajesh Rathod, Amazon EC2 product leader) (available at https://aws.amazon.com/blogs/modernizing-with-aws/year-in-review-2023-aws-for-microsoft-workloads/)

[34] See J Girven and C Lindeman, "Reduce your Microsoft licensing costs by upgrading to 4th generation AMD processors," AWS Compute Blog (4 Nov 2024) ("Amazon Web Services (AWS) and AMD have collaborated since 2018 to deliver cost effective performance for a broad variety of Microsoft workloads, such as Microsoft SQL Server, Microsoft Exchange Server, Microsoft SharePoint Server, Microsoft Systems Center suite, Active Directory, and many other Microsoft workload use cases. This post shows how the performance improvements of the latest generation AMD-powered Amazon Elastic Compute Cloud (Amazon EC2) instances can help you reduce licensing costs on Microsoft workloads running on AWS")

[35] See AWS Optimisation and Licensing Assessment (available at https://aws.amazon.com/optimization-and-licensing-assessment/)

[36] See Migrate and modernise Microsoft workloads on Google Cloud (available at https://cloud.google.com/windows?hl=en).

[37] See Kevin McLaughlin and Catherine Perloff, "How Google Cleaned Up Its Cloud Act," The Information (3 Feb 2025) ("Google has for some time been trying to use its power in advertising to find more cloud customers.") (available at https://www.theinformation.com/articles/how-google-cleaned-up-its-cloud-act?rc=5kmidk); Haranas, Mark, "Google Cloud Cuts VMware Cloud Prices," *CRN* (24 July 2024) ("Google Cloud is now offering customers bundles that package together Google Cloud Platform credits, YouTube ads, and the G Suite productivity software.") (available at https://www.crn.com/news/cloud/2024/google-cloud-cuts-vmware-cloud-prices-5-azure-aws-and-broadcom-things-to-know); Tilley, Aaron, "Google Bundles Products to Land Univision Cloud Deal," *Wall Street Journal* (26 April 2021) ("[Google Cloud], which announced the deal [with Univision] Monday, beat out rival services by packing its offering with benefits across its YouTube video platform, advertising and search services, according to company executives.") (https://www.wsj.com/articles/google-bundles-products-to-land-univision-cloud-deal-11619449740?msockid=2634a634d6416e992890b21bd7666fe9); and Subcommittee on Antitrust, Commercial and Administrative Law of the Committee on the Judiciary, "Investigation of Competition in Digital Markets: Majority Staff Report and Recommendations," p. 246 (2020): see also id. at pg. 242 ("*Google appears to leverage its dominant business lines, including popular APIs such as Google Search and Maps, along with machine learning services, to attract customers to its platform through discounts and free tier services. For example, according to internal strategy documents, in 2018 Google 'launched a program with the Play team to provide GCP credits to game developers based on their Play Store spend, to increase focus on Play and incentivize migration to GCP.' By harnessing Google's advantages in existing markets, GCP is undermining competition on the merits*")

remain open and competitive. That commitment has never wavered, and it remains our approach today.

(59) The stakes here are high – missteps in policy could weaken one of the UK's most dynamic, high-growth industries. Respectfully, much of the data underpinning this market investigation is already out-dated, outpaced by the speed of cloud and AI innovation, and largely out of sync with the real forces shaping this industry today. We find it difficult to reconcile the main concerns[38] in the PDR with the fast-moving realities of the cloud computing market – in particular the PDR's provisional recommendation to commence SMS investigations into Microsoft and AWS. The SMS regime is intended to target firms with substantial and entrenched market positions.[39] In contrast, after nearly two and a half years of investigation into the cloud market, the CMA has identified issues that are largely peripheral to competition today and are manufactured by Google and Amazon for their private advantage. The PDR does not do enough to show that these areas warrant further scrutiny under the SMS regime and the facts in front of the CMA tell a clear story: competition is intensifying, investment by multiple providers has never been higher, prices are falling, innovation is rapid, and providers across the UK are scaling their businesses quickly.

(60) For all these reasons, Microsoft believes that starting a **third** investigation – under the new Digital Markets, Competition and Consumers Act 2024 ("**DMCC**") – in relation to cloud computing after 30+ months of consideration by Ofcom, first, and the CMA, now, is a step in the wrong direction. Such an investigation would force the CMA down the path of intervening in a fast-evolving market on rapidly shifting sands, towards the inevitable outcome of the UK cloud computing market becoming the most regulated in the world.

(61) The Market Investigation Panel might consider, instead, to recommend that that CMA Board: (i) continue to actively monitor developments and outcomes for UK customers in the cloud computing market, particularly the ongoing growth of Google Cloud, the rapid evolution of AI technology and the impact, if any, of changes driven by the EU Data Act (supported by obligations on AWS, Microsoft Azure and Google Cloud to provide relevant information to the CMA); (ii) publish an annual review of these outcomes; and (iii) consult on any further action, to the extent required.

(62) An approach focused on monitoring of this "*vital input to businesses and organisations across the UK economy*"[40] is more proportionate and aligned with the CMA's strategy to deliver on its duty to promote UK competitiveness and economic growth.[41]

---

[38] It is notable that the closest the PDR gets to articulating a material harm to UK consumers is where it notes that "*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" (emphasis added)* – see PDR, paragraph 8.56. This is purely hypothetical and does not go into any detail of whether customers are currently paying more than they would do in a "well-functioning market"

[39] See New Bill to crack down on rip-offs, protect consumer cash online and boost competition in digital markets - GOV.UK which described the regime as being aimed at "*market **dominance [that] has stifled innovation and growth across the economy**, holding back start-ups and smaller firms from accessing markets and consumers*"

[40] PDR, paragraph 1

[41] See Mansion House 2024 speech - GOV.UK and Chancellor calls on watchdog bosses to tear down regulatory barriers that hold back growth - GOV.UK

## 2 Structure of this response

(63) The remainder of this response is divided into three parts:

(i) **Part A** responds to the PDR's emerging thinking regarding the competitive landscape in the cloud computing market;

(ii) **Part B** responds to the PDR's theories of harm regarding egress fees and technical barriers, and comments on the CMA's analysis regarding committed spend agreements or discounts ("**CSAs**" or "**CSDs**", respectively); and

(iii) **Part C** addresses the PDR's provisional conclusions relating to Microsoft's licensing practices.

## PART A: Response to Competitive Landscape

(64) This part of the response covers three topics regarding the PDR's conclusions on competitive landscape. First, it notes that instead of focusing on evidence which shows that prices are falling, and that innovation is prevalent, the PDR prioritises static measures of competition, which often provide only an incomplete picture; see **Section 3**. Additionally, **Section 4** notes that the PDR misses the bigger picture by disregarding the explosive emergence of generative AI and marginalising the impact of cloud providers' multi-billion-dollar investments in this area. Finally, **Section 5** covers how the PDR incorrectly accepts Google's arguments that it is a second-tier cloud provider when there is overwhelming evidence to the contrary.

(65) The UK Government has directed the CMA to prioritise UK economic growth[42] by driving innovation, increasing productivity and encouraging investment, particularly in key sectors on which the UK's industrial strategy is based. It has also emphasised the pivotal contribution of AI[43] in the context of growth. In response, the CMA has acknowledged[44] that, to deliver on this mission, it will act at pace and take regulatory actions in a predictable and proportionate way, with sound processes in place. The Government has also given a clear direction that the competitiveness of the UK on the international stage should be supported by a regulatory approach that is aligned to positions taken in other jurisdictions.[45]

(66) Against this background and bearing in mind the extremely dynamic nature of this market, described in detail below, it is Microsoft's proposal that the CMA focus on active monitoring ahead of intervention. This will provide the flexibility necessary for the cloud computing market to successfully transition through a pivotal moment of development, while also allowing the CMA to assess how to best to intervene, if, and when, it might be appropriate.

## 3 The PDR prioritises the wrong evidence and misses that competition is prevalent and delivering for customers

### 3.1 Evidence shows that competition in the cloud is delivering tangible and material benefits for consumers

(67) As explained above, there is a wealth of evidence that shows that the cloud market is working and delivering positive results for customers. While the PDR acknowledges many of these benefits it then incorrectly disregards them as part of its assessment.

(i) **Price declines.** Quality-adjusted real effective prices for the top 5 Azure products (VMs, storage, SQL database, Azure app service and backup), which constitute **[CONFIDENTIAL]** of Azure's total revenue, have fallen from 2018-22.[46] Real effective prices for the top 5 products fell **[CONFIDENTIAL]** as of 2022. These figures are consistent with a market with ongoing price competition and innovation. While the PDR acknowledges the decrease in Microsoft's prices,[47] it does not fully engage with Microsoft's pricing analysis. Instead, it only notes that some prices are

[42] Strategic steer to the Competition and Markets Authority - GOV.UK and Growth, opportunity, and prosperity for the UK at the heart of the CMA's next Annual Plan - GOV.UK

[43] AI Opportunities Action Plan - GOV.UK

[44] See New CMA proposals to drive growth, investment and business confidence – Competition and Markets Authority

[45] The FCA and PRA have recently become subject to secondary duty to align with relevant international standards

[46] Microsoft Response to the Competitive Landscape, Committed Spend Agreements and Egress Fees Working Papers Dated 1 July 2024 ("**WPs Response**"), paragraph 12(a)

[47] PDR, paragraph 3.286(c)

increasing while others are decreasing, without indicating the proportion of each. It also appears to treat empirical analysis submitted by Amazon in the same way.[48] This approach is incomplete, and it is critical that the CMA thoroughly considers and gives appropriate weight to evidence on pricing trends as part of its assessment.[49]

(ii) **No erosion in discounts.** The value of discounts that Azure customers received has remained **[CONFIDENTIAL]** (2022-23) while CSAs for UK customers **[CONFIDENTIAL]**.[50] This result is similar for Microsoft's competitors. Again, this shows that competition is fierce not only for new customers, but also existing customers. Cloud providers compete for every workload and are under pressure to retain existing customers through discounts. Similarly, the PDR notes that the discounts provided by large cloud providers do not weaken or marginalise smaller rivals, nor do they affect their ability to compete.[51] As with the pricing analysis, the PDR should give appropriate weight to this evidence in its assessment.

(iii) **Innovation via feature upgrades.** Microsoft's cumulative number of feature updates for the top 20 UK cloud services was 927 in 2023. Similarly, the number of **[CONFIDENTIAL]**.[52] Customer evidence supports this, and cloud providers are viewed as "*great innovators*". Additionally, the PDR notes that there is evidence of cloud providers innovating to win business from their competitors.[53] The PDR also acknowledges this innovation has benefits for customers, but otherwise appears to give it limited weight in its assessment and claims that this is not evidence of competition working fully effectively.[54]

(iv) **Capex that converts to customer benefits (scale, quality, innovation).** Since Azure's launch, the combined capex of Amazon, Microsoft and Google has been c. $700 billion. Microsoft's R&D and capex in particular has accelerated since Azure first launched in 2010, with a more than tripling in absolute R&D spend and a 12x rise in absolute capex.[55] Microsoft recently announced its plan to invest $80 billion in FY2025 "*to build out AI-enabled datacenters to train AI models and deploy AI and cloud-based applications around the world*";[56] Amazon's expected capital expenditure for 2025 is $100 billion, of which the "*vast majority"* will go towards making AI capabilities available to Amazon Web Services customers;[57] and Google has planned $75 billion of capital expenditure for 2025 to "*accelerate its AI ambitions*".[58] Similarly, following the UK's International Investment Summit in October 2024, a range of technology companies are investing £14 billion in AI infrastructure (data centres) in the UK.[59] The PDR recognises that spending on R&D

[48] PDR, paragraph 3.289

[49] PDR, paragraphs 3.277-3.278

[50] PDR, paragraph 7.61 and WPs Response, paragraph 12(b)

[51] PDR, paragraph 7.126

[52] PDR, paragraph 3.293(d)

[53] PDR, paragraph 3.300

[54] PDR, paragraph 3.300

[55] WPs Response, paragraph 12(d)

[56] See The Golden Opportunity for American AI - Microsoft On the Issues

[57] See Amazon doubles down on AI with a massive $100B spending plan for 2025 | TechCrunch

[58] See Alphabet brushes off DeepSeek fears with mammoth $75B capex spend | TechCrunch. Microsoft had previously committed to spending £2.5 billion on UK AI infrastructure between 2023-2026, while Amazon has pledged to invest £8 billion in UK digital and AI infrastructure between 2024-2028

[59] See https://pitchbook.com/news/articles/tech-firms-commit-14b-to-uk-ai-data-center-plan

can enhance quality and increase efficiency, allowing customers to benefit from higher-quality, more innovative cloud services and / or lower prices[60] and, as per the UK Government's draft strategic steer to the CMA, such business and international investment should be encouraged as it "*is a critical driver of growth*".[61] However, the PDR does not give any material weight to these developments, nor to the positive outcomes for customers, as part of its assessment.

**Figure 2: Capex by Cloud Providers:**

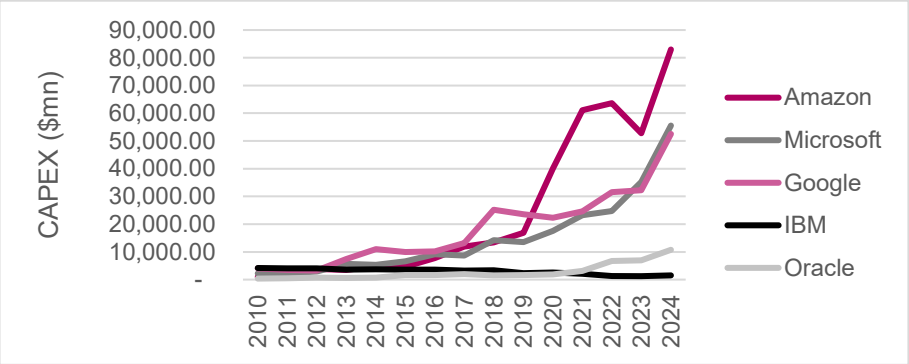


*Source: S&P Capital IQ*[62]

(v) **UK capacity expansion.** Microsoft has expanded the UK capacity of its dedicated Azure data centres by **[CONFIDENTIAL]**.

(68) But these are not just statistics – they feed through to *direct benefit for UK customers*. For example, last month Nationwide announced that its use of Azure OpenAI had truly transformed its back office operations with direct benefits for its customers in the form of a 66% reduction in the time it takes to deal with customer queries.[63] Other customers have been able to develop entirely new products using our technology, like Placing Platform Limited's use of a range of Microsoft tools to build a new trading hub enabling transition "*to a data-first digital trading environment [that] will help drive London Market growth and innovation*".[64]

(69) It is unclear what additional evidence the CMA would require to conclude that competition is working effectively, when prices are declining, quality is improving, and on all metrics, there is evidence of competition working in practice to deliver positive outcomes for customers. Notwithstanding this, the PDR primarily relies on evidence of relatively stable market shares and its ROCE analysis to reach its conclusions that "*a more competitive market would have sustained better market outcomes*".[65]

(70) We respectfully disagree - the evidence on market outcomes points in one clear direction: the cloud market is delivering positive results for customers, in the form of lower prices, consistent discounts, rapid and material innovation, and ongoing (very significant) capacity expansion. This comes at a time in which the explosive emergence of generative AI is

[60] PDR, paragraph 4.64

[61] See Strategic steer to the Competition and Markets Authority - GOV.UK

[62] Notes: The graph spans calendar years 2010 to 2024. Oracle data is reported at their fiscal quarter level, which does not align with the calendar quarters. Hence, Oracle's revenues have been computed by summing over nearest fiscal quarter data, which spans from November of the previous calendar year to November of the current calendar year

[63] See: 'Copilot, not autopilot': How Nationwide is boosting award-winning customer service with data and AI - Microsoft UK Stories and How Virgin Money is Humanising Digital Experiences with AI - IBM UK & Ireland - Blog

[64] See: PPL launches new trading hub built with Microsoft tech

[65] PDR, paragraph 3.505

upending the industry and lifting cloud providers old and new. In comparison to other market investigations where the CMA has elected to intervene, these are not market conditions that would typically justify intervention.

### 3.2 The evidence that the CMA relies on does not properly reflect the competitive conditions

(71) The CMA's profitability assessment assumes that cloud profitability over the past five years (2020-2024) represents a steady-state equilibrium.[66] However, this assumption does not capture the true competitive conditions for two reasons: first, Azure became profitable for the first time only in FY2020[67] and second, because of the rapid pace of technological change in cloud computing in recent quarters with the developments around AI.

(72) Forrester Research showed that the global cloud market grew by almost three times between 2015 and 2020 (from $87 billion to $236 billion), representing a growth rate of 22% year-on-year.[68] The industry is still evolving, particularly with the emergence of AI-driven cloud services, such as AI compute and accelerated workloads, which have recently become a major competitive battleground for attracting customers to the cloud.[69]

(73) By treating recent profitability trends as indicative of long-term stability, the CMA misses the companies' ongoing investment cycles. Microsoft's AI-driven capex on GPUs is expected to increase from $2.7 billion in FY2023 to **[CONFIDENTIAL]** in FY2027, a **[CONFIDENTIAL]** compound annual growth rate. This massive capital shift indicates that cloud services are entering a new investment phase rather than operating in a steady state.

(74) There are also a number of other flaws in the PDR's profitability assessment:

(i) **The PDR does not adjust for accounting-driven increases in EBIT and ROCE:** Changes in accounting estimates rather than underlying competitive conditions can partly explain the upward trend in EBIT and ROCE. Both Microsoft and AWS have extended the useful life of their cloud infrastructure assets, reducing annual depreciation expenses and increasing reported profitability. These adjustments do not inform market power or competition, yet the CMA does not normalise for these accounting changes when assessing profitability. By failing to adjust for these changes, the CMA's ROCE estimates overstate how sustainable Microsoft's profitability is and attribute accounting-driven gains to market power rather than financial reporting changes.

(ii) **The PDR underestimates WACC by using company-level regressions:** ROCE is calculated using operating profits from the narrowest reporting segments, specifically AWS for Amazon and Azure for Microsoft. However, WACC is computed using an equity beta derived from the entire Amazon and Microsoft businesses, regressed against the S&P 500. Both Amazon and Microsoft are major constituents of the S&P 500, collectively accounting for approximately 10% of the index.[70] Because beta is calculated based on the correlation of stock price movements with the broader market, their significant presence within the S&P 500 distorts their beta estimates. As a result, the beta estimates used in the WACC calculations may be

[66] PDR, E.19 and paragraph 3.242

[67] PDR, Appendix E, Table E.4

[68] See Total size of the public cloud computing market from 2008 to 2020, Forrester Research (Statista)

[69] See Cloud revenues poised to reach $2 trillion by 2030 amid AI rollout, Goldman Sachs

[70] See Top 25 companies by index weight, Investopedia

biased downward, underestimating the true systematic risk of AWS's and Microsoft's cloud businesses.[71]

(iii) **The CMA overstates Microsoft's ROCE by not taking into account the full historical cost of building Azure:** Microsoft Azure only became profitable in FY2020, but the investments that enabled its success, including data centres, networking equipment, and foundational cloud infrastructure, were made long before 2018 (the start of the period considered in the CMA's assessment). The CMA claims these investments were reflected in the capital employed in calculating ROCE.[72] However, because many of these assets reached the end of their useful life before 2018, they no longer appear on Microsoft's balance sheet and are therefore not included in the CMA's capital employed estimates. A more accurate assessment would require adjusting capital employed to reflect the full investment timeline of cloud infrastructure rather than relying on a snapshot of the balance sheet that excludes past capital expenditures.

## 4 The PDR incorrectly dismisses the impact of AI and accelerated compute, and misses that it is a key source of innovation and competition in the cloud market

(75) The distinction drawn in the PDR between IaaS based on standard compute and IaaS based on accelerated compute[73] means the CMA incorrectly disregards the dynamic nature of investment in cloud computing infrastructure.

(76) As noted above, while it is true that demand for AI workloads has increased the use of accelerated compute, this does not mean that it operates as a separate market from standard compute. On both the supply and demand sides of the cloud computing market, generative AI has rapidly become a primary driver of strategic planning and investment. AI is just one of many applications of cloud computing, and the same customers that require accelerated compute also require standard compute for various workloads.[74] Cloud providers also invest in common infrastructure that supports both types of compute, such as data centres, networking and storage.[75]

(77) Evidence shows that accelerated compute is not a separate market:

(i) **Customers generally do not buy accelerated compute in isolation:** Customers, including Foundation Model ("**FM**") developers, purchase compute and other cloud services based on workload requirements.[76] The key difference between the two types of compute is in the hardware configuration (e.g. CPU versus GPU),[77] but

[71] See Estimating risk parameters, Damodaran, "*The conventional estimate of this relative risk, measured by regressing stock returns against a market index, is flawed for three reasons - the market index can be dominated by a few stocks, the beta estimate can be noisy and the firm itself might have changed during the course of the regression*"

[72] PDR, paragraph 3.261

[73] PDR, paragraphs 3.447, 3.496 and 3.507

[74] "*In systems that use GPUs, the CPU still plays an important role in managing the GPUs' tasks and all other processing tasks that, while not as resource-intensive, are still integral to the computer's functions*" – see CPU vs. GPU for Machine Learning | IBM

[75] There has been a shortage of accelerator chips relative to demand which is attributable to long lead-times partly driven by third-party suppliers and the need for R&D to develop proprietary chips. It is unclear how substitutable the two types of compute are if chips were readily available

[76] "*Structurally, a GPU is extremely similar to a CPU. It comprises essentially the same components: DRAM, control unit, cache and an ALU*" – see CPUs and GPUs: What to use when for AI/ML workloads | Rafay

[77] See GPU vs CPU - Difference Between Processing Units - AWS

customers mix and match these compute resources to suit their workload. Similarly, customers do not procure AI compute in isolation – they integrate it with other PaaS services and IaaS services, including storage, networking, and security, which are shared across both standard and accelerated compute instances.[78] The mere fact that AI has increased demand for certain types of compute does not mean those compute resources constitute a separate market.

(ii) **Demand-side substitutability between standard and accelerated compute:** While accelerated compute is often used for AI workloads and high-performance computing, many tasks that initially rely on standard compute can migrate to accelerated compute as business needs evolve. Even if it is early days for AI, customers often reallocate workloads between the two, depending on cost-efficiency, availability, and performance considerations. Similarly, AI inference tasks, which begin on accelerated compute, may shift to standard compute as models become optimised for CPU-based execution.

(iii) **Supply-side substitutability:** Cloud providers do not build isolated infrastructures for standard and accelerated compute. Instead, they invest in scalable data centres that support both types of workloads.[79] There has been a shortage of accelerator chips relative to demand,[80] but this is attributable to long lead-times, partly driven by third-party suppliers and the need for R&D to develop proprietary chips.

(78) By placing accelerated compute in a separate market, the PDR misses a key source of innovation in the cloud market. The growth of AI and the demand for accelerated compute is driving competition between cloud providers across all areas of their businesses. Microsoft and their competitors cannot sit still, and this is why they are investing billions of dollars in growing their capacity and developing foundation models, all while innovating on new tooling, developer services, and applications.

(79) The growth of AI and accelerated compute cannot be put to one side in the context of this market investigation. Cloud providers are locked in a battle to roll out AI capabilities and solutions as quickly as possible and it is a key parameter of competition between them.

## 5 Placing GCP in a 'second tier' of cloud providers is out of date and inaccurate

### 5.1 The market shares that the PDR relies on to classify GCP as a second-tier cloud provider do not present as clear a picture as the CMA suggests

(80) The PDR's choice to categorise GCP as a second-tier cloud provider appears to primarily rely on GCP's market share, which the PDR says is "*significantly smaller*" than AWS and Azure. As we set out in the following sub-sections, we respectfully submit that this approach is out of date and ignores the wealth of other evidence which indicates that GCP is not only a strong competitor now but will become an even stronger competitor going forward.

(81) However, even the market shares present a more mixed picture than the PDR's conclusions indicate. The PDR's UK shares of IaaS revenue growth (Table 3.6) and IaaS / PaaS revenue

[78] See Infrastructure for AI: Why storage matters | IBM

[79] For example, Microsoft's CFO made the following statement in its Q4 FY2024 earnings call "*Cloud and AI related spend represents nearly all of total capital expenditures. Within that, roughly half is for infrastructure needs where we continue to build and lease datacenters that will support monetization over the next 15 years and beyond. The remaining cloud and AI related spend is primarily for servers, both CPUs and GPUs, to serve customers based on demand signals*" – see Microsoft Fiscal Year 2024 Fourth Quarter Earnings Conference Call

[80] PDR, paragraph 3.38

growth (Table 3.7) show that GCP is growing while Microsoft's share is falling, and AWS's shares remain stable. This indicates that, while still the smallest hyperscale cloud provider, GCP is growing faster than AWS and Azure and will be a larger competitor going forward.

(82) Additionally, these shares likely understate GCP's growth over this period.[81] As has been widely reported, GCP's lower profitability in the past is due in part to its own internal accounting issues.[82] The PDR does not appear to account for this in its methodology and we would encourage the CMA to address this issue before the final report.

(83) Similarly, the shares likely overstate Microsoft's recent growth, as its share of new customers in Table 3.8 and Table 3.9 are inconsistent with the fact that new customers who joined Azure between January and December 2023 only generated **[CONFIDENTIAL]** of 2023 Azure revenues in their first year of spend.[83]

### 5.2 Recent financial announcements from Google since the publication of the PDR show the inconsistency of placing GCP in a second tier

(84) As we note in our Executive Summary (see Section 1), the market shares cited by the PDR provide an incomplete picture of GCP's strong competitive position. Since the publication of the PDR, Google has released its latest financial results, in which the growth of Google Cloud was a key component:[84]

(i) Google Cloud grew by 30% in the fourth quarter of 2024 (which was higher than any of its quarterly growth rates in 2023).

(ii) Google Cloud's annual revenue was $43.2 billion in 2024, up 30.7% compared to 2023.

(iii) In 2024, the number of first-time commitments to GCP more than doubled compared to 2023.

(iv) GCP closed several strategic deals over $1 billion, and the number of deals over $250 million doubled from the prior year.[85]

(v) Growing even faster than its revenue, Google Cloud's profits were over $2 billion in Q4 2024 (from $864 million in Q4 2023), and its operating margin grew at a similar rate to 17.5% (from 9.4% in Q4 2023).[86]

(85) These results paint a clear picture of Google Cloud growing at pace and taking a growing market position. This was reinforced by comments from Google's leaders – such as Sundar Pichai, Google and Alphabet CEO who said, "*we delivered another strong quarter in Q4, driven by our leadership in AI and our unique full stack approach… we are set up well for continued growth*".[87] In short, Google sees itself as taking an ever-larger position fuelled by its privileged position in AI and becoming a more significant force going forward.

---

[81] Microsoft notes that although recent reported Google Cloud growth has been at similar levels to Azure (30% in Q4 2024 versus 31% for Azure), Google Cloud comprises of GCP and Google Workspace and GCP has been growing at a rate much higher than Google Cloud overall – see Alphabet's 2024 Q4 Earnings Call

[82] See How Google Cleaned Up Its Cloud Act

[83] Microsoft further notes that the share of new customers analysis conducted by the CMA is also susceptible to error due to inconsistencies in how cloud providers record customer data

[84] See Alphabet Announces Fourth Quarter and Fiscal Year 2024 Results

[85] See Alphabet Q4 earnings call: CEO Sundar Pichai's remarks

[86] See 2024 Q4 Earnings Call - Alphabet Investor Relations

[87] See Alphabet Q4 earnings call: CEO Sundar Pichai's remarks

(86) Google's external statements are consistent with its results in the influential Gartner Magic Quadrant for Strategic Cloud Platform Services. In 2024, GCP made a "*notable jump*", with Gartner placing GCP in the Leader quadrant effectively on par with AWS and Azure for its ability to execute and completeness of vision.[88]

(87) Google's result and public statements, and third-party reports are in stark contrast to the positioning of the PDR, which portrays GCP as a distant third party rather than the strong and fast-growing competitor that it really is.

### 5.3 The PDR fails to account for how Google's dominance in other markets gives it a clear advantage in cloud computing

(88) AI is also a major driver of Google's cloud business growth, but the PDR does not take account of Google's obvious advantages in AI and the future importance of those AI workloads. Google's CEO has emphasised Google's deep vertical integration in this regard: "*We have a **unique advantage** because we develop every component of our technology stack, including hardware, compilers, models and products*".[89] Google also enjoys significant market power as a result of its platform, which puts its AI service (Gemini) "*into the hands of billions of people around the world*".[90] Google's control of critical training data (such as YouTube) and consumer distribution channels (Android and the Google Play Store), as well as its complete vertical integration and platform power, affords it a competitive advantage that neither Amazon nor Microsoft enjoy.

(89) Similarly, the PDR assessment does not account for the unique benefits Google derives from its dominant position in digital advertising and other services. That, as the PDR notes, these are outside the cloud computing market is irrelevant: Google brings them to bear in that market. As noted previously by the CMA in its Online Advertising and Digital Platforms Market Study Final Report[91] and in our Working Paper responses, Google leverages its digital advertising and other businesses to win customers for GCP, not only for advertising services, but also for their other cloud-computing needs too. It is not surprising that GCP's largest industry segment is media companies who depend heavily on digital advertising for monetisation.[92]

(90) In summary, Google is already a strong competitor to AWS and Azure and has all the tools to further strengthen its position – significant financial resources, an already growing market share, the ability to leverage its dominance elsewhere to grow its cloud share and a leadership position in relation to AI. However, despite this overwhelming evidence, the PDR persists with the view that GCP is in a distant third place and should be treated differently from AWS and Azure.

---

[88] See Google is a Leader in Gartner Magic Quadrant for Strategic Cloud Platform Services | Google Cloud Blog

[89] (emphasis added) See 2024 Q4 Earnings Call - Alphabet Investor Relations

[90] See 2024 Q4 Earnings Call - Alphabet Investor Relations

[91] *See, e.g.,* CMA, Online Platforms and Digital Advertising Market Study, App O: Measurement Issues in Digital Advertising, paras 123-24 ("*Google Campaign Manager (its advertiser ad server) has no API to export exposure data and generally prohibits exposure data from leaving Google. Using Google's [Ads Data Hub] is the only way to access all of Google's post-campaign data and measure analytics services in a single environment, and ADH does not permit customers to export the information to any other measurement or ad tech partner they may prefer. These actions by Google have placed pressure on advertisers to choose between: (a) using Google for all of its ad tech and analytics services (Google ADH and Google Analytics) … or (b) using independent ad tech"*)

[92] See Google Cloud Market Share and Buyer Landscape in 2024, HG Insights

### 5.4 The exclusion of Google from any ultimate remedies would be unreasonable, highly distortive and amount to picking winners

(91) The PDR's provisional course of action (which would exclude Google from any intervention and oversight) would simply amount to the CMA picking winners in the market, providing Google with an unjustified leg-up, ultimately distorting competition in the UK:

(i) Google would have commercial freedom to invest in UK infrastructure without the threat of being obliged to provide unlimited data for egress services for free.

(ii) Google would also be able to depart from the rigid standards imposed by the CMA and would therefore be able to develop differentiated (including premium) products.

(92) This is particularly surprising given GCP (with its unique access to data) is best positioned of all cloud providers to benefit from the growth of AI and should exacerbate concerns about the distortive nature of the proposed remedies.

(93) Given – as described above – Google's leading and growing competitive position in the UK cloud market, as well as its self-proclaimed "*unique*" advantage in AI infrastructure, which is widely considered to be fundamental to current and future cloud workloads (as well as its vital offerings in search and digital advertising and beyond), there is no basis on which the CMA should discriminate between Microsoft, AWS and Google when scoping its intervention given that this would otherwise result in significant competitive distortion between the three comparable cloud providers in the UK.

## PART B: Response to the Egress Fees and Technical Barriers Theories of Harm and Comments on CSAs

## 6 Egress fees

### 6.1 The PDR's proposed remedy of banning egress fees is unprecedented and disproportionate

(94) Bandwidth services provide demonstrable value to cloud customers by allowing them to transfer significant volumes of data in and out of the cloud at desired levels of prioritisation, latency and security. Microsoft incurs significant fixed costs for the network infrastructure and ongoing data transfer costs (e.g. transit and peering costs) to supply bandwidth services.[93] The market has set the price of data ingress at zero but allowed for fees on data egress, i.e., egress fees.

(95) The PDR's proposal to prohibit the ability of Microsoft and AWS to charge fees[94] for the bandwidth services used when UK customers switch clouds or egress data in connection with their use of multiple clouds ("**multi-cloud**") is the most interventionist remedy of those considered on this topic. If adopted, it would curtail the commercial freedom of cloud providers to charge for all forms of standard and premium egress. It is hard to see how this remedy is proportionate given the consistent evidence that egress fees play a marginal, at best, role in customer decisions to switch or multi-cloud and its predictable negative unintended consequences.

(96) Before addressing both of these concerns, it is worth noting that the PDR's proposal goes further than the EU Data Act, which only bans egress fees in respect of: (i) standard / ISP-

[93] As evidenced in its internal documents – see WPs Response, paragraph 97

[94] Excluding direct connections and CDN transfers (PDR, W.173)

routed options (i.e. excluding premium egress); and (ii) switching, while permitting egress fees for multi-cloud provided they do not exceed the costs incurred for such services.[95] The PDR's proposed remedy would therefore mean an additional regulatory burden, representing a substantial limit on the ability of global cloud providers to recover the costs of operating and to investing in building out networking infrastructure in the UK. It would also place UK cloud providers at a competitive disadvantage compared to their global / European peers.[96]

### 6.2 The PDR's very interventionist remedy is not warranted based on the evidence, given the hypothetical analysis of a peripheral issue

(97) The PDR dismisses the use of real market data on the prevalence and actual spend by UK customers on egress usage, as it argues customers may be deterred from switching or multi-cloud due to the high level of egress fees.[97] However, the PDR provides no evidence of such customer "*avoidance*" behaviour and ignores the market realities that customers opt not to switch or multi-cloud for a number of more material reasons unrelated to egress fees[98] (and the fact that cloud providers often offer discounts for egress).

(98) In addition, the UK cloud market dynamics – which has led to free ingress due to intense competition for workloads and falling quality-adjusted prices for egress – demonstrate that no intervention is required. While the PDR places limited weight on the real-world evidence of falling egress prices,[99] the CMA's Market Investigation Guidelines are clear that (*inter alia*) price is an important and measurable market outcome for customers, which helps to determine whether there is an adverse effect on competition ("**AEC**") and in the consideration of remedies.[100]

(99) The PDR also disregards the reality that there has been a low uptake on the major cloud providers' elimination of egress fees for switching over the past year;[101] if egress fees were a meaningful barrier to switching, then a material use of such offer would be expected. The PDR's consideration of broader factors affecting customer switching decisions – including the fact that such programmes are (in the PDR's view) "*voluntary*", customers' awareness of the programmes and requirements to submit switching applications in advance –[102] is

[95] The PDR expresses concern as regards the regulatory risks of imposing parallel remedies under the market investigation and the DMCC – see PDR, paragraphs 9.77-9.79 and W.209, but it does not give such weight to the risks of regulatory fragmentation under the DMCC and EU Data Act. The fact that global cloud providers can charge different regional prices for their egress services is not particularly relevant, given that the PDR's proposed remedy does not merely affect the level of egress prices but the ability to charge for egress fees at all, thus fundamentally affecting the multi-product cloud pricing model

[96] If, for example, egress fees are driven to zero in the UK cloud market for all providers due to competition, while global competitors are still able to charge for such services ex-UK (and ex-EU as regards banned egress fees for switching)

[97] PDR, paragraphs 5.315, 5.333-5.334, 5.425 and 5.363

[98] As well as disregarding the fact that the PDR's benchmark of highly-integrated multi-cloud architecture may be misguided – see WPs Response, paragraphs 39, 43 and 89(e)

[99] See PDR, paragraphs 5.324, 5.424 and 5.468. As noted in Microsoft's March 2024 pricing submission, egress list prices have fallen since 2018 with the expansion of the egress free tier in Q2 2022, and the introduction of internet egress routed via public ISP network

[100] PDR, paragraph 103. The PDR does 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. However, falling real prices indicate competition. Further, as acknowledged by the PDR (paragraph 5.443), Microsoft's investment in its networking infrastructure affects multiple cloud services. As such, any cost savings from such investment is likely to be spread across all relevant services rather than specific ones

[101] PDR, paragraph 5.385

[102] PDR, paragraphs 5.386-5.387

inconsistent with the PDR's one-dimensional assessment of egress fees' influence on customer switching / multi-cloud decisions.[103]

(100) As a consequence of disregarding actual market data, the PDR relies on an analysis of one-off costs for (in its own words) "*hypothetical scenarios of switching and multi-cloud architectures*".[104] The PDR accepts that its hypothetical analysis is "*purely illustrative and therefore does not necessarily account for the boundaries of what might be practical*".[105] The realities of switching and multi-cloud are more complex than indicated by the PDR's simplistic analysis.[106] There are also a number of empirical errors with the PDR's quantitative analysis of one-off switching costs.[107]

(101) **The CMA cannot escape the reality that egress fees are a very small and negligible proportion of UK customers' cloud bills**. The PDR found that egress fees as part of a hypothetical "*one-off*" switch for an average customer based on a full transfer of their data may exceed only **[CONFIDENTIAL]** per year (if amortised over the average length of a contract as customers typically do) for less than only **[CONFIDENTIAL]** of customers.[108] The PDR's quantitative analysis does not evidence that it is a meaningful barrier to switching / to multi-cloud and, as such, provides no basis on which to base a robust AEC finding.

(102) The PDR's customer evidence (which is consistent with its Jigsaw Research)[109] indicates that egress fees are not typically the "*main*" nor a "*significant*" concern for customers nor a relevant factor for a "*majority*" of customers, but nevertheless erroneously considers it

[103] WPs Response, paragraph 82 and footnote 78

[104] PDR, paragraphs 5.334; 5.342, and 5.358

[105] (*emphasis added*) PDR, M.7-M.10

[106] WPs Response, paragraphs 8 and 82

[107] **The PDR used list prices** from the cloud providers websites for data transfers (see PDR, paragraphs 5.342(b) and 5.359(b)), **despite customers often negotiating discounts on egress fees** (WPs Response, paragraph 89(e) and PDR, paragraphs 5.337 and M.15). There is a strong incentive for Microsoft to negotiate such discounts on egress fees to maintain long-term / future customer relationships. Microsoft's Azure website advises customers should "*contact for prices over 500TB*"

The PDR's switching analysis **extrapolates the impact for customers wanting to lift and shift the entirety of their cloud spend**, despite the cloud providers' **free egress programmes eliminating all egress fees** where customers transfer all of their data from one provider to another

The PDR's switching analysis uses unweighted averages of the switching cost share of total spend (PDR Annex M). Based on our calculations, a weighted average (i.e. shares are weighted by overall revenues) using the PDR methodology and data **shows that one-off switching cost is less than [CONFIDENTIAL] across all years (and would be even lower where such cost is amortised across three years**)

The **PDR's hypothetical heatmap analysis is unrealistic and grossly overestimates the volume of monthly data egressed by UK Azure customers**. The average monthly data egressed by UK Azure customers was **[CONFIDENTIAL]** in 2022 (see Annex MSFT_CMA_Cloud_Data_0000046 submitted in response to Q12 of RFI s174 dated 2 November 2023; we estimated the monthly volume by dividing the annual egress volume by 12). This is just **[CONFIDENTIAL]** of the average monthly data stored summed across all UK customers in 2022

Using the PDR's heatmap data (PDR, Table M.8), the observed egress rate of **[CONFIDENTIAL]** of data stored by UK customers results in an **egress spend share of less than [CONFIDENTIAL] across all UK Azure customers in 2022**. If the PDR considered that the total volume of data egressed would double under its "*ideal*" multi-cloud (i.e. customers transfer **[CONFIDENTIAL]** of the monthly volume of data stored), using the PDR's heatmap **egress costs are still less than [CONFIDENTIAL] of total spend across all UK customers**

The **egress products considered by the PDR does not capture cheaper (per GB) alternatives that may be available to the customers** – e.g. customers transferring large volumes of data may potentially use Azure's Express Route unlimited data plans to reduce their egress costs (see WPs Response, paragraph 87)

[108] Such switching costs may exceed **[CONFIDENTIAL]** of annual spend if incurred in a single year (PDR, paragraphs 5.345-5.346). According to the PDR's analysis, the **majority of AWS, Microsoft and Google's UK customers would have to pay [CONFIDENTIAL] of their total annual spend in 2022 for a total switch of cloud provider** (PDR, paragraph 5.343(b))

[109] The PDR's Jigsaw Research, which found that almost no customers view egress fees as a significant barrier to switching or adopting multi-cloud (see the WPs Response, paragraph 89)

sufficient to find an AEC on the basis that they constitute a "*meaningful*" barrier to switching / multi-cloud for a "*substantial minority of customers*".[110]

(103) The PDR's dismissal of the real market data / dynamics and customer evidence in favour of a **hypothetical, unrealistic and erroneous analysis of egress fees' impact on switching and multi-cloud is not a robust basis of evidencing an AEC finding requiring the very interventionist remedy** proposed by the PDR. The PDR's evidence as it stands shows that egress fees are a hypothetical, or at most peripheral, concern for certain customers in the UK seeking to switch or multi-cloud.

### 6.3 The PDR's proposed intervention will reduce incentives to invest in premium network infrastructure in the UK and drive up prices on all customers.

(104) The CMA acknowledges in its Market Investigation Guidelines that a direct regulatory intervention on market outcomes such as the price "*may reduce economic efficiency (including dynamic incentives to invest and innovate) and adversely affect the economic interests of customers over the longer term*".[111] This is why remedies must be proportionate and any interventions that control market outcomes should be considered as a last resort.[112]

(105) The negative consequences of barring cloud providers from charging fees when bandwidth is used to egress data are easy to predict and flow from basic economic tenets and common business sense.

(106) First, barring egress fees will significantly diminish cloud providers' incentives to maintain and invest in the connection network assets which are the product of billions of dollars in investment by cloud providers. As recognised in the PDR, Microsoft had invested **[CONFIDENTIAL]** in cloud fixed assets as of 30 June 2023.[113]

(107) While cloud network quality is important to Microsoft's other cloud services,[114] what would be lost is the incentive to develop and offer premium bandwidth services that cater to the needs of fewer but more demanding customers. In fact, Microsoft may not be able to afford to offer any premium bandwidth services to UK customers at all. If such services must be offered at zero price, there is a real danger that all UK customers will choose the premium bandwidth service (why not? It's better and free!), driving up costs and overtaxing the network. In contrast, outside of the UK (including in the EU post-EU Data Act), Microsoft and AWS will remain free to invest in the high-performance infrastructure and services required to supply premium egress to customers. This proposed remedy would thus result in a **deterioration of the innovation, quality, variety, and availability of egress services and cloud infrastructure in the UK**.

(108) Second, to recoup their network costs, cloud providers will be forced to increase the prices of other cloud services in the UK. The real risk of such "*waterbed effects*" in multi-product firms is acknowledged by the CMA in its Market Investigation Guidelines,[115] as well as in the

[110] PDR, paragraphs 5.320, 5.399 and 5.420-5421

[111] (*emphasis added*) CMA's Market Investigation Guidelines, paragraph 352(a)

[112] CMA's Market Investigation Guidelines, paragraph 89

[113] The PDR also recognises that the larger cloud providers in the UK "*have invested significantly more in global infrastructure than smaller cloud providers*" – see PDR, paragraph 4.17(b); see also Microsoft's Response to the Issues Statement Dated 9 November 2023, paragraph 2.4(16)

[114] PDR, W.174

[115] CMA's Market Investigation Guidelines, paragraph 91(d)

academic literature on regulation of network industries.[116] This means that all customers will bear the network costs, whether or not they use the network and regardless of the extent to which they use the network.

(109) Finally, reducing egress fees to zero will inevitably lead to inefficient use of network infrastructure, driving up costs.[117] This is an incontrovertible consequence when the price for a service that provides intrinsic value to customers is reduced to zero.[118] Inefficient use is different from the switching or multi-cloud use that the CMA is seeking to promote. Cloud customers would have no incentive to be efficient in their data transfer practices and (as outlined above) would be incentivised to use the highest performing premium data transfer service that is available for free (thereby overusing that service in particular).[119]

(110) The CMA is under a statutory duty to consider the significant adverse effects on UK customers and the UK cloud market of the proposed ban on egress fees[120] to avoid imposing unnecessary burdens on businesses in the UK and to ensure that its remedies are proportionate to address the AEC identified.[121] Given that the PDR's proposed remedy is the most costly and the most harmful option on incentives to invest, it also **runs contrary to any policy objectives of encouraging business investment and growth of the UK tech economy** (which is one of the UK government's key sectors).[122]

### 6.4 Google must be subject to the same intervention (if the CMA insists one is required) to avoid distortion and detriment to the UK cloud market

(111) The PDR does not adequately explain why it is not necessary, proportionate or practicable[123] to include Google within the scope of the egress fees remedy, which is particularly concerning for UK customers given the PDR's calculations that **Google's customers face [CONFIDENTIAL] egress fees on average than AWS's and Microsoft's customers**.[124]

---

[116] See (i) *Testing the "Waterbed" Effect in Mobile Telephony by C Genakos and T Valletti*: "*Under quite general conditions of competition, theory suggests that lower termination charges will result in higher prices for mobile subscribers (the "waterbed effect")… although regulation reduced termination rates by about 10%, this also led to a 5% increase in mobile retail prices*"; and (ii) *Mobile Call Termination by M Armstrong and J Wright*: "*High FTM termination charges are a means of transferring surplus from fixed callers to mobile recipients (and in part to mobile networks if the waterbed effect is not complete)*"

[117] As evidenced in its ordinary course documents – see the WPs Response, paragraphs 95-96

[118] Academic studies have shown that when products are offered at zero prices, consumers may engage in wasteful or inefficient behaviour because they cannot internalise some or all of the associated costs. This behaviour is evident in contexts like traffic congestion and road overuse without tolls, and in airport scenarios where take-off and landing slot pricing does not reflect economic costs. See (i) *Antitrust in Zero-Price Markets: Foundations by J M Newman*, page 186; (ii) *The Economics of Welfare by A C Pigou*, which argues that roads are overused because drivers can externalise the congestion costs they create; and (iii) A Market in Airport Slots by K Boyfield et al

[119] The "*zero price effect*" is where people, when choosing between two products with one being free, tend to overvalue the free option. This suggests that a price of zero not only reduces the cost of purchase but also enhances the perceived value of the product. This effect is observed in various contexts, such as health insurance, where zero pricing encourages healthier children to use more services and increases the inappropriate use of antibiotics. See (i) *Zero as a Special Price: The True Value of Free Products by K Shampanier et al*; (ii) *Is Zero a Special Price? Evidence from Child Health Care by T Iizuka and H Shigeoka*; (iii) *Marketing Science 26.6 (2007), paragraphs 742-757*; and (iv) *American Economic Journal: Applied Economics (2002), paragraphs 381-410*

[120] As outlined above, the banning of egress fees will result in a range of adverse consequences broader than the reduction of Microsoft's profits – such that the PDR cannot disregard the adverse negative effects of its proposed ban on egress fees (CMA's Market Investigation Guidelines, paragraph 353)

[121] CMA's Market Investigation Guidelines, paragraphs 348, 353 and Annex B, paragraph 54

[122] See Invest 2035: the UK's modern industrial strategy - GOV.UK

[123] CMA's Market Investigation Guidelines, Annex B, paragraph 54

[124] PDR, M.23. In 2022, the average cost for a "*one-off*" switch as a percentage of total spend in 2022 was **[CONFIDENTIAL]** for Google's customers at **[CONFIDENTIAL]** compared to **[CONFIDENTIAL]** for Microsoft's customers and **[CONFIDENTIAL]** for AWS' customers (PDR, M.19 and Figures M.1-3)

(112) The PDR justifies its position on the basis that other cloud providers (e.g. Google) would follow any zero egress fees implemented by Microsoft and AWS.[125] This remains a speculative basis for designing a remedy and relies on future voluntary actions by cloud providers (precisely the type of voluntary actions that the PDR discounts when considering the cloud providers' free egress programmes).[126]

(113) If Google were outside the scope of the PDR's proposed remedy, then it would retain the commercial freedom to recoup its costs and ROI through egress fees, it would not be obligated to provide effectively unlimited egress services for free and, as a result, it would be the **only major player in the UK cloud market to be in a position to make significant investments in egress services** (in particular, premium egress) and the network infrastructure for UK customers.

(114) As noted above, Microsoft and AWS would have no incentive to invest in high-end egress services because they would be forced to offer them to all UK customers for free. More critically, they would not be able to operate effectively in practice because all customers would use whatever premium service is available: it would not be rational for a customer to choose a lower speed egress service if it is the same price (i.e. free) as a higher speed egress service. In practice, this means premium egress networks from Microsoft or AWS – even to the extent they exist – would not reliably operate as such.

(115) By contrast, Google – which will be able to recover its investment – will not only have the incentives to invest in, but be in a position to effectively operate, premium egress services. This would have the perverse impact of decreasing customer choice from the position today where there are multiple providers of innovative egress services. This is contrary to the aims of market investigation remedies which – per the CMA's Market Investigation Guidelines – should be to "*significantly increase competitive pressures in a market within a reasonable period of time*",[127] and to "*facilitate economic growth and increase the choice available to customers*".[128]

(116) Any proposed remedy on egress fees **must not discriminate between Microsoft, AWS and Google** as this would otherwise be **highly distortive to competition in the UK cloud market to the detriment of UK cloud customers** (as outlined further in Section 6.4 above). Given the Data Act applies to GCP, any intervention that captures only Microsoft and AWS would sit uncomfortably with recent CMA comments that it will seek to avoid divergence from regulatory outcomes in other jurisdictions, especially in relation to remedies.[129]

# 7 Technical barriers & interoperability standards

## 7.1 Introduction

(117) Microsoft supports improvements to technical standards / interoperability in cloud (and other technology markets) where they benefit customers. However, we firmly believe that technical standards in the cloud market need to be developed by market participants and industry bodies to ensure they make sense for customers, given they have the right incentives, resources and experience and avoid reducing incentives to innovate. Placing responsibility on the CMA (e.g. via the Digital Markets Unit) to determine which standards should be

[125] PDR, W.170

[126] PDR, paragraph 5.388

[127] CMA's Market Investigation Guidelines, paragraph 330

[128] CMA's Market Investigation Guidelines, paragraph 331

[129] New CMA proposals to drive growth, investment and business confidence – Competition and Markets Authority

implemented and supported risks depriving UK cloud customers of access to new features and innovation, and risks forcing cloud providers to divert resources away from developing standards that truly benefit customers.

(118) The PDR places undue weight on the commercial reality that technical differentiation exists as a result of competition for innovation between cloud providers. The PDR has failed to provide robust evidence justifying recommending the CMA takes forward a remedy on standardisation in cloud. Its case for imposing standards is based merely on anecdotal and qualitative evidence and ignores the fact that the market is already overcoming technical barriers where customers want it. The fact of product differentiation (and fierce competition based on such differentiation) will always mean there are some technical barriers to interoperability between cloud services, in the same way that there is always such a barrier in all areas of computing. No form of regulation can remove all such barriers without creating barriers to innovation. In the absence of infinite resources, the market is best placed to indicate where greater interoperability is most desired and how it is best achieved from a technical and customer perspective, and where suppliers ought to focus their efforts.

(119) Microsoft accepts that the DMCC has introduced additional flexibility to the CMA's remedies powers; however, this does not translate to the CMA being in a position of technical expertise as regards public cloud infrastructure to design or designate interoperability standards that would potentially hinder innovation when there is no cogent evidence of customer harm. The UK Government's draft strategic steer warns that the development of markets by such new and emerging technologies (such as cloud and its numerous fast-developing applications) is complex and not always predictable, thus the CMA "*should take particular care to ensure growth and innovation benefits are prioritised*".[130]

(120) This concern is particularly acute here. As noted above, while some technologies require standards to be completed ahead of any marketplace activity (e.g. mobile), the opposite is true for others (e.g. AI). Cloud computing is a mix of these two with no obvious decision framework that would suggest the CMA is promoting interoperability or mandating unnecessary resource allocation to outdated technologies.

### 7.2 There is no compelling reason to believe the CMA will make better technical judgements than the market – in fact, this risks harming the provision of cloud services in the UK

(121) The PDR does not provide any compelling reasons why the CMA is well placed to set standards for the global cloud market, let alone better placed than the 1,200 active information technology standardisation organisations and the tens of thousands of technical experts who participate in the technical committees of these organisations. It is difficult to imagine how any single regulatory body could have the relevant expertise to set standards in cloud is accepted in the PDR.[131] Additionally, the CMA does not have the ability to drive forward the development of open source software, which is the second pervasive method of market-driven interoperability for cloud services.

(122) Market players are best placed to make complex commercial and technical judgements to drive technical standards based on their extensive experience developed in their involvement in the cloud market since its origins[132] (and this is why Microsoft promotes the

[130] See Strategic steer to the Competition and Markets Authority - GOV.UK

[131] PDR, W.51-54

[132] Relative to e.g. the CMA which has been in shadow operation for only a few years and acquired its legal powers on 1 January 2025 with the entry into force of the DMCC

open source community and is active in various industry bodies to facilitate solutions that enable portability and interoperability in the cloud).[133] Standard-setting requires specific deep technical and engineering expertise of cloud infrastructure (rather than of data science) to set standards that address customer needs and are practical for stakeholders.

(123) Cloud providers as well as other bodies / stakeholders (e.g. in the open-source community) have already implemented a vast number of technical standards as well as developing numerous workarounds to technical barriers / differentiation in response to customer need, often through open source software.[134] The PDR highlighted many of these – such as Infrastructure as Code, self-managed open-source software (using IaaS), improved API integrations, containers and Kubernetes, cloud-agnostic services and tools and adaptors.[135] Many cloud providers (such as Microsoft) are members of the CNCF organisation,[136] which promotes the growth and development of cloud-native technology ecosystems and projects (e.g. container technology).[137] The PDR also recognised that the development of common IAM standards by cloud providers has helped facilitate customers' ability to integrate and switch between clouds.[138]

(124) There are intrinsic market incentives that drive devising through trial and error and investing resources in successful technical standards valued by customers. The fact that the CMA has the ability to test, trial and iterate remedies[139] does not mean that it has the technical capabilities to do so, or even the resources to do so in a timely fashion. It also gives rise to the risk of starting an "*open season*" among other regulators to set their own national standards, as well as allowing a UK regulator to set a global standard without local accountability or consultation that (by the simply global nature of the services involved) impacts the provision of cloud services in other territories. The likelihood of such standards diverging creates an untenable situation.

(125) It should be common ground that, once standards have been embedded into regulation, they may become difficult to remove and / or change if they become obsolete – regardless of the fact that the technology they relate to has moved on, or that customers are not adopting them. The PDR acknowledges that "*The specific set of cloud services where the likely benefits of requiring a common standard exceed the distortion risks are […] likely to change over time*".[140]

(126) **Placing responsibility on the CMA to set technical standards risks doing more harm than good for UK cloud customers and providers.** The PDR accepts that common standards could "*reduce the ability and/or incentive of cloud providers to innovate through improving and differentiating their services*".[141] Given that the market is already delivering solutions to technical frictions where there is customer demand to do so, requiring cloud providers to invest resources into adhering to regulator-imposed standards that may not be used by customers diverts resources away from investments which would improve outcomes

[133] See paragraph (30) above as well as Microsoft's submission on the CMA's Conceptual Remedies Framework Dated 23 August 2023, paragraph 7

[134] PDR, W.10

[135] PDR, paragraphs 5.281-5.282 and 5.266

[136] See Members | CNCF

[137] See Graduated and Incubating Projects | CNCF

[138] PDR, paragraph 5.193

[139] PDR, W.42

[140] PDR, W.47

[141] PDR, W4.45

for customers and competition in the UK. As recognised in the PDR, technical standards that are mis-designed or are too "*low*" a level may "*suppress natural developments in the market… **forcing suppliers to use a suboptimal standard**, and in turn **prevent a superior standard from being adopted***"[142] and "***reduce or limit the scope for differentiation on certain parameters**, which in turn could suppress **incentives to innovate***".[143]

(127) Standardisation would also impose additional costs on multinational cloud providers operating and investing in the UK (as per the CMA's experience in previous market investigation standardisation remedies).[144]

### 7.3 The PDR does not provide robust evidence of a material technical barrier to switching and multi-cloud beyond inherent technical differentiation

(128) Technical differentiation also helps players stand out in the competitive cloud market. The PDR accepts that cloud providers' independent R&D efforts and intense competition in terms of developing new and improved services have led to technical differentiation in cloud services in the UK market (in terms of, e.g., security features, developer tools, scalability and compliance certifications).[145] All of Microsoft, AWS and Google have highly diversified and sophisticated cloud offerings with numerous parameters of technical differentiation (e.g. as recognised by the PDR, Microsoft's products are highly secure).[146] Oracle submitted that "*it is differentiated from other cloud providers in part because it optimises for speed and performance*".[147] There are also, for example, certain smaller cloud providers that specialise in the provision of infrastructure optimised for AI workloads (e.g. CoreWeave and Lambda Labs).[148]

(129) From our experience, customers will weigh up using innovation / technical complexity against the ability to switch / multi-cloud with minimum friction. It is entirely the customers' choice whether to trade off complexity for switch-ability.[149] The PDR recognises the existence of this trade-off in respect of PaaS and other core services based on evidence from customers and stakeholders,[150] and also acknowledges the operational risks (e.g. to security) to customers of multi-cloud architecture in certain circumstances.[151] **The PDR's proposal to implement standards fails to adequately reflect the harm to customers that may be caused by such a regulatory standardisation**. The negative impact on innovation in such a nascent and evolving market such as cloud should not be overlooked.

(130) Moreover, the PDR's conclusion that there exist technical barriers causing customer harm is based on anecdotal evidence relating to such generic technical differentiation, rather than meaningful evidence of anticompetitive harm related to specific products / services. This reflects the fact that many of the "*technical barriers*" cited by the PDR are merely inherent

[142] PDR, W.31

[143] PDR, W.26

[144] For example, while up to 1.5 years was initially allocated from the final report in Retail Banking (2016) to establish the open banking standard, the implementation period lasted much longer and was ultimately finalised in September 2024. This incurred substantial monitoring and enforcement costs and posed complications for all industry participants

[145] PDR, paragraph 5.87

[146] Microsoft's industry-leading security framework boasts over 100 compliance offerings, which is the most comprehensive compliance coverage of any cloud service provider. See Trust your cloud | Microsoft Azure

[147] PDR, paragraph 5.82(e)

[148] PDR, paragraph 3.402

[149] PDR, paragraph 5.125

[150] PDR, paragraphs 5.111 and 5.125

[151] PDR, paragraph 5.52

features of technical differentiation in a high-tech and competitive market (as described above) rather than resulting from any anticompetitive features – for example, the most common "*technical barrier*" cited by customers in the Jigsaw Research was the need for software engineers to rewrite a large amount of code to recreate or integrate applications across different cloud environments.[152] Where the PDR actually assesses the competitive effects of particular cloud products raised by customers such as Microsoft's IAM services and tools (e.g. Entra ID and Active Directory), it finds provisionally that there is insufficient evidence of harm or that they do not create additional technical barriers to customer multi-cloud or switching.[153]

(131) The PDR acknowledges that its evidence is "*largely qualitative*",[154] and the evidence it describes relates almost entirely to anecdotes, primarily from customers.[155] These anecdotal examples do not represent convincing evidence that customers' incentives to switch / multi-cloud are harmed by any technical barriers or differentiation. For example, the PDR found that "*customers expressed differing views on the extent to which there were benefits to multi-cloud*",[156] and that at least some customers do **not** perceive there to be significant benefits to multi-cloud or integration.[157] The PDR acknowledges that customers undertake cost-benefit analysis when deciding whether to switch or multi-cloud, with some ultimately concluding that it is not worthwhile.[158] This **PDR's hypothetical scenario of frictionless switching and multi-cloud does not exist, and is in any event not what customers want.**

(132) The evidence adduced in the PDR does not constitute concrete evidence of an AEC in the cloud market as result of any technical barriers erected by cloud providers beyond inherent technical differentiation beneficial for customers in an innovative and competitive market – therefore, imposing an SMS designation and dictating the technical roadmap for public cloud services on this basis is not justified.[159]

## 8 Committed spend agreements or discounts

### 8.1 The PDR correctly finds that CSAs are not detrimental to competition in the cloud services market in the UK

(133) CSAs are a vital and beneficial component of the competitive process in the cloud market. They are part and parcel of fierce price competition between cloud providers, encouraging customer switching / multi-homing and, at a broader level, supporting investment in public cloud infrastructure.

(134) Having undertaken a detailed assessment, the PDR broadly accepts this view and finds that CSAs do not harm competition in the provision of public cloud infrastructure services. We welcome this finding and support the PDR's conclusion that no intervention is required with

[152] PDR, paragraph 5.50

[153] PDR, paragraph 5.199

[154] PDR, paragraph 5.89

[155] PDR, paragraph 5.47

[156] PDR, paragraph 5.93

[157] PDR, paragraphs 5.61 and 5.94

[158] PDR, paragraphs 5.93 and 5.96

[159] CMA's Market Investigation Guidelines, paragraphs 348 and 353; and CMA's Market Investigation Guidelines, Annex B, paragraph 54

respect to CSAs. However, there remain areas of the PDR's analysis that we disagree with, and which should be addressed in the final decision:

(i) The PDR's speculation that CSAs could potentially become detrimental to competition in future is purely theoretical and is in fact undermined by the CMA's own analysis.

(ii) The PDR gives insufficient weight to the pro-competitive aspects of CSAs – they support investment in UK cloud infrastructure and are a key driver of customer switching.

### 8.2 The PDR's view that CSAs could potentially become detrimental to competition in future is purely theoretical

(135) We agree with the PDR's conclusion that CSAs do not raise concerns today. Indeed, the PDR concludes that competitors who are as profitable as AWS would have the ability to profitably compete for incremental demand in 99.9% of instances and the incentive to do so in 99.2-99.6% of instances. Similarly, when the PDR repeats the assessment, but with a competitor that is significantly less profitable than AWS, the results are broadly comparable with a rival having the ability to profitably compete in 97.7-98.4% of instances and the incentive to do so in 94.7-97.7% of instances.[160] In summary, **having looked in detail at the evidence, in almost all scenarios, the CMA found that CSAs had no potential negative impact on competition**.

(136) Nevertheless, the PDR speculates that CSAs may harm competition "*in the future*" if: (i) the market matures such that "*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.[161] The PDR does not assess or calibrate how likely either of these scenarios are or what the impact on its analysis would be (e.g. how the ability / incentive percentage figures above would change). We disagree with this conclusion in two ways:

(i) First, for the reasons outlined in our Working Paper responses, we do not accept the premise of "*sticky*" demand in this market. The underlying evidence shows customer demand is contestable on a durable and long-term basis.

(ii) Second, the public cloud market is unlikely to mature or stabilise, in a conventional sense, in any realistic or foreseeable timeframe as shown by the very large (and increasing) investments by the cloud providers. As the PDR notes, the UK cloud services market has grown at an annual average rate of 33% in recent years and this is expected to continue.[162] More generally, only **[CONFIDENTIAL]** of all Microsoft's customers and **[CONFIDENTIAL]** of customers spending more than $500,000 have CSAs, meaning that even within providers' existing customer base, there remains a huge pool of customers who do not currently have CSAs.[163] These percentages **[CONFIDENTIAL]** for AWS.[164] On any metric, the cloud services market is expected to grow rather than mature. Furthermore, as noted above, cloud

[160] PDR, paragraph 7.107

[161] PDR, paragraph 7.129

[162] PDR, paragraph 2.86

[163] PDR, paragraphs 7.43-7.44

[164] PDR, paragraphs 7.43-7.44

providers continue to innovate and there is no reason to think that competition for the next workstream will weaken in the future.

(137) Each time the CMA has considered whether CSAs are detrimental to competition, it has looked at the underlying evidence and found no reason for concern. The same is true for its potential concerns about the impact on future competition. There is no evidence to suggest that current competitive conditions will evolve in a way which is detrimental to competition.

### 8.3 The PDR should give more weight to the pro-competitive nature of CSAs – they facilitate investment in UK cloud infrastructure and customer switching

(138) While the PDR recognises that CSAs can deliver some potential benefits for both customers and cloud providers, it suggests that these benefits need to be balanced against any longer-term harm to competition. As noted above, we consider the PDR's concerns regarding the potential for CSAs to be detrimental to future competition to be speculative and not supported by evidence.

(139) In contrast, the benefits that CSAs deliver to market participants today and going forward are tangible and real, which are dealt with only in passing.

(i) **CSAs enable cloud providers to invest in better infrastructure, support and services for their customers.** The commitments in CSAs enable cloud providers to make significant upfront investments in providing new and / or improved infrastructure, technical support and solution development. Cloud providers, such as Microsoft, have the confidence to make investments when they know they have guaranteed revenue from the CSA. Ultimately this investment has positive implications for other market participants (who often benefit indirectly from advances made with other customers) and also the wider UK economy in the form of improved infrastructure and greater skills / training.

(ii) **CSAs provide customers and cloud providers greater certainty of pricing and demand, which in turn supports customer-specific investment decisions.** Similar to cloud providers, the certainty created by CSAs enables customers to take material investment decisions that benefit not only themselves but the wider market. Fixed duration agreements, like CSAs, provide customers with assurance that the guaranteed terms will persist over the life of their cloud contract. This certainty allows customers to embark on more expansive or technologically challenging projects – e.g. not just switching from one cloud provider to another but also building new experiences and solutions on the cloud. For example, Microsoft's 10-year agreement with LSEG enabled it to build its next-generation data analytics and modelling solutions with Azure, and to migrate key technology infrastructure to the cloud.[165]

(iii) **CSAs do not hurt but instead enhance competition between cloud providers:** CSAs allow cloud providers to subsidise and help customers overcome any switching costs associated with switching from rivals. As with investments in infrastructure (described above), this is only possible with the commitment aspect of a CSA. These commitments help facilitate customers moving from one cloud provider to another and are generally supportive of switching. For example, **[CONFIDENTIAL]** and **[CONFIDENTIAL]** received **[CONFIDENTIAL]** in order to facilitate switching to Azure.

[165] See LSEG and Microsoft launch 10-year strategic partnership for next-generation data and analytics and cloud infrastructure solutions; Microsoft to make equity investment in LSEG through acquisition of shares - Stories

## PART C.I: Microsoft's Licensing

### 9 Microsoft's licensing of its IP to Amazon and Google does not foreclose them as effective competitors – key considerations

(140) Chapter 6 of the PDR provisionally finds that Microsoft's licensing of its IP to Amazon and Google[166] forecloses them as effective competitors for particular workloads alleged to have high Microsoft IP costs for Amazon and Google, relative to Azure. These alleged high-Microsoft-IP-cost-for-rivals workloads are sometimes referred to in shorthand as *Microsoft workloads* or similar.

(141) The key relevant Microsoft software subject to the IP licensing cost dispute are Windows Server, an operating system, and SQL Server, a database management system:

(i) Windows Server software in an input that, combined with CPU hardware, forms a Virtual Machine ("**VM**") a cloud infrastructure or IaaS component: a [WINDOWS SERVER VM].

(ii) SQL Server can also be an IaaS component: [SQL SERVER].

(142) Chapter 6 of the PDR bases its financial modelling on hypothetical purchasing scenarios of one or both of these IaaS components. These scenarios do not reflect how customers buy and use cloud services, as the PDR itself recognises in other chapters.

#### 9.1 Customers purchase cloud *workloads* of services that work together – they do not purchase one or two (Microsoft-IP based IaaS) *workload components* at a time.

(143) Cloud competition is for units of customer demand of *at least one individual workload*, though customers frequently *aggregate* workloads into "*deals*"[167] and further into CSA.[168] The larger the customer spend, the more likely this workload aggregation (and the use of CSAs) is. Azure had **[CONFIDENTIAL]** large UK customers that spent $20+ million each on Azure in 2024. Of these, **[CONFIDENTIAL]** had CSAs and the remaining one, **[CONFIDENTIAL]**. Each of the **[CONFIDENTIAL]** large UK customers consumed **[CONFIDENTIAL]** common services,[169] in addition to VMs, in both 2023 and 2024 as part of wider consumption of ~**[CONFIDENTIAL]** services on average.

(144) Conversely, there is no discrete demand at the level of infrastructure (IaaS) *components*. It is meaningless to make important downstream competition inferences at the component level, such as [WINDOWS SERVER VM] or [SQL SERVER].

(145) **The CMA's first major mistake is factual:** "*a Windows Server VM is the narrowest set of downstream products that Windows Server can serve as an input for".*[170] In fact, the narrowest set is a *workload.* Discrete IaaS components are not workloads, and a [WINDOWS SERVER VM] is not a workload. Nor does it become one merely by adding a [SQL SERVER] component.

[166] The CMA excludes concerns with respect to other cloud service providers, so the foreclosure focus is exclusively on the two rival hyperscalers

[167] PDR, paragraphs 6.448(a), 6.458(a) and 6.253

[168] See PDR Chapter 7

[169] The common services were: **[CONFIDENTIAL]**

[170] PDR, paragraph 6.267

**9.2 The CMA hypothesises that the largest most-sophisticated customers hive off workloads composed largely of one or two Microsoft IP-based IaaS components. This is unevidenced, unrecognisable to Azure, and economically puzzling.**

(146) Bypassing the reality that large customers purchase a large swath of different services, the CMA hypothesises, without even a footnote, that cloud competition works in a way that it does not.

(147) The CMA's conjecture is that the "*larger*" $20+ million spend multi-cloud customers[171] are "hiving off" some demand into "*single purchasing*" of "*Microsoft workloads*" to claim that Microsoft licensing costs reflect such a large portion of the cost and revenue opportunity for these hived off workloads that Amazon and Google cannot effectively compete.

(148) **The CMA's second major mistake is unsafe conjecture: "***Microsoft software workloads likely comprise a larger share of the total spend that large customers allocate to Azure within a single purchasing decision*".[172]

(149) It is *conjecture* because the CMA cites no evidence why this is "*likely*". It is *unsafe* and *unsupported* conjecture because it is:

(i) **Factually invalid.** As noted, and at a basic factual and conceptual level, [WINDOWS SERVER VM] and [SQL SERVER] are not themselves a workload, in principle or in practice, individually or together. They cannot comprise a "*single purchasing decision"*. Labelling them a "*Microsoft workload*" or similar, as one-dimensional shorthand, does not change the facts.

(ii) **Unrecognisable to Azure.** In the ordinary course of business Azure does not recognise the hypothesis that $20+ million spender large customers, as a class, behave in the way the CMA hypothesises.[173] Of course, Azure wants to win workloads, however they may be labelled: as explained in our first CMA meeting, all Azure consumed revenue ("**ACR**") is good revenue. But in order to compete, Azure does not market itself for so-called "*narrow purchasing decisions*"[174] as if its USP were "narrowly" as a destination for "*Microsoft workloads*" (which would mean less ACR) rather than a "*broad"* destination for "*broad purchasing decisions*" of the full suite of Azure services (which would mean more ACR) – such as the **[CONFIDENTIAL]** common services, or the **[CONFIDENTIAL]** total services that on average Azure's own $20+ million spender customers actually consume, almost all of which are *not* the subject of the SPLA licensing term dispute.

(iii) **Unexplained hypothesis of rational behaviour by large customers**. The CMA does not explain the economic puzzle of why, to any material degree: (i) the most sophisticated cloud customers who (ii) spend $20+ million overall on a large range of services overall, and who (iii) typically sign CSAs to obtain discounts over such large, aggregated levels of spend (comprising many workloads), would "hive off" a

[171] For "*large*" customers, the CMA has in mind *"customer groups with a higher propensity to multi-cloud, such as higher spend customers"* (PDR, paragraph 6.268) and by "*higher spend*" it is apparent that the focus, or high water mark of the AEC case, rests on the highest spend bracket, namely $20+ million: see spend brackets in PDR, Tables 6.3 to 6.6. See also PDR footnote 1746: "*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"*

[172] PDR, paragraph 6.298(d)

[173] Nor can Microsoft discern, from any available Microsoft or Google submissions, that either rival hyperscaler is the source of the large-customer conjecture

[174] PDR, paragraph 6.266, "*In each case,* [total customer spend on Azure services] *and* [customer spend on Windows Server VMs and SQL Server licensing IP on Azure] *represent narrow and broad purchasing decisions, respectively*"

"*Microsoft workload*". Doing so would: (1) incur transaction costs; and on the CMA's theory both (2) weaken the workload's margin attractiveness to Amazon and Google; and (3) mean paying Microsoft a "higher price" for the workload, stand-alone. The CMA must explain why its hypothetical rational customer would not source all the "*Microsoft workload*" components together or add the associated consumption to AWS or GCP committed spend to pocket the CSA-wide discount (and where, for Amazon and Google, it will be a small fraction of their CSA-level revenues, margins and costs).[175] It should also explain why rivals would not encourage this, in their own economic interest, as well as that of the customer.

(150) The combination of a factual mistake and an unsupported and unsafe conjecture infects the robustness and undermines the meaningfulness of the remainder of the CMA's "raising rivals' costs" ("**RRC**") analysis premised on both. The focus on [WINDOWS SERVER VM] and [SQL SERVER] does not tell the CMA whether Amazon and Google are truly "*unable to compete effectively*"[176] for cloud workloads. Much of the resulting evidence is therefore not probative of a properly-grounded AEC concern.

### 9.3 The CMA's margin analysis – specifically its embrace of Amazon's and Google's self-serving ask that their margins are "*too low*" to compete effectively – is demonstrably wrong in the light of the new Keystone analysis

(151) More decisive than the (flawed) RRC analysis, however, is the margin analysis of the CMA. This is because (healthy) gross margins drive incentives to compete (hard).

(152) The decisive point from the CMA is that: "*we consider that positive and low margins … 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*".[177]

(153) This claim makes little sense and does not even attempt to address why customers in the UK would be better off with AWS and GCP having higher IaaS component-level margins. Moreover, a proper analysis of actual costs and margins, shows no such challenge in competing for customers who incorporate Microsoft software into their workloads.

(154) While Google, Amazon, and in turn the PDR, may label certain workloads as being "Windows Server" or "SQL Server" or "Microsoft" workloads, this one-dimensional label establishes nothing more profound than that the workload includes one or two components for which Microsoft IP is an input. The labels do not speak to the substance: the share of total cost of that IP and, more importantly, the size of available gross margin across the workload as a whole.[178]

(155) The total workload margin (the "pot of total gross profit") is what incentivises competition. The CMA finds that *"[m]ost customers we spoke to said they use a mix of IaaS and PaaS for their cloud workloads*"[179] So: for Google, for "*most"* workloads, the pot of total gross

[175] If the CSA were only with Azure, and none with AWS or GCP, why "unbundle" and hive off that demand and pay separately, rather than keeping it within the flat CSA-wide discount? As to multi-cloud customers scenarios without any CSAs, the PDR says that an advantage of multi-cloud is: "*[c]ompetitive 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*" (PDR, paragraph 2.82(a)). The CMA must explain the puzzle why its hypothetical multi-cloud customers would apparently do the opposite

[176] CMA's Market Investigation Guidelines, paragraph 270

[177] (emphasis added) PDR, paragraph 6.462

[178] If revenue were used to compute the share of SPLA cost, it would incorporate information on the margin on the whole workload

[179] PDR paragraph 3.45

workload-level margin comprises (1) total IaaS services margins plus (2) total PaaS services gross margins (e.g., Google data, Google analytics, Google AI, etc.)

(156) The total subset for (1) of IaaS gross margin, for Google, where there are Microsoft IP inputs is (a) gross margin (mark-up) on an IaaS component, a [WINDOWS SERVER VM], with the associated IP license cost, and (b) the gross margin on Google's own IaaS components (such as Google storage and networking).

(157) A newly-extended Keystone analysis, as an extreme worst-case, shows ample total margin to incentivise Amazon and Google to compete hard.

(158) Keystone's standard "as-efficient-competitor test" results show that Google would make healthy margins even if the workload is limited to include just a WS VM and two technically inevitable complementary IaaS components – *Google* storage, and *Google* networking – any time that it pays the [WINDOWS SERVER VM] IP cost (SPLA) and matches the Azure Hybrid Benefit ("**AHB**") discount. The overall pot of gross margin on this subset of IaaS components is in the **[CONFIDENTIAL]%** range. If [SQL SERVER] is added for PAYG customers, the equivalent is still **[CONFIDENTIAL]%**.

(159) All these numbers are a worst-case or lower-bound extreme. They relate to minimum complementary IaaS components (storage, networking) but not *other* near-essential IaaS services such as management and security services. They also take no account, as a lower-bound, that most workloads will have "*a mix of IaaS and PaaS*" with additional gross margin for Google. All the same points above that apply to the #3, Google, obviously apply to the #1, Amazon.

(160) Given these facts, no exceptional UK government intervention in Microsoft's IP rights is warranted, let alone an intervention designed to increase Google's and Amazon's margins because of a concern that both are "*less likely to compete fiercely*" enough "*for Microsoft workloads*".[180]

## 10 Critical flaw in the CMA's AEC analysis

### 10.1 It is obviously wrong to use a [WINDOWS SERVER VM] component lens for analysing (alleged Amazon and Google foreclosure from) contestable purchasing decisions

(161) It is common ground that "*[c]loud providers do not compete to supply a single downstream product with a single unit cost; [r]ather, they compete to supply bundles of services that vary in composition and value*"[181] and that these so-called service bundles are known in the industry as **workloads** that solve a discrete set of IT needs.

(162) Workloads are the narrowest discrete unit of demand; as Amazon puts it "*customers assess their IT needs on a workload-by-workload basis*".[182] The individual workload is therefore also the narrowest contestable unit of competition for that demand, although most commonly customers make deployment decisions for multiple workloads at one time.[183]

[180] PDR, paragraph 6.462

[181] PDR, paragraph 6.243

[182] PDR, paragraph 3.77

[183] This finding is replete throughout the PDR. The PDR notes: *"[c]ompetitive 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*" (PDR, paragraph 2.82(a)); "*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*" (PDR, paragraph 2.62); and *"[e]ven customers that do not multi-cloud may choose to deploy some workloads on one cloud and later add others [i.e.*

#### 10.1.1 A Windows Server VM cannot be a workload or even a large part of one

(163) Yet in Chapter 6, the PDR focusses on a single IaaS component – [WINDOWS SERVER VM] – which it considers "*the narrowest set of downstream products that Windows Server can serve as an input for*".[184] It fails to consider that, even as a subset of a workload, the narrowest set of downstream products to make the VM even marginally functional is at minimum multiple essential IaaS components: a VM, storage and networking.

(164) As noted above, a workload with Windows Server VM must necessarily also include at a minimum its two essential complements: storage and networking. These are non-substitutes for each other (as the CMA finds)[185] but also true economic **complements** – in the sense that they must be used together.[186]

(165) When assessing the significance of Microsoft software inputs to Google and Amazon competing for cloud customers,[187] the CMA relies on the prism of input costs to consider *"narrow*" and "*broad"* purchasing decisions by customers (i.e. Microsoft software will be significant when competing for "narrow" demand for [WINDOWS SERVER VM]).[188] How customers define and structure their demand for a discrete workload is what dictates what suppliers' input cost proportions will be. The PDR artificially takes a workload's components [WINDOWS SERVER VM] and assesses the input cost in relation to that component instead of the workload comprising the input. It there, perhaps inadvertently, implies that high supplier input costs effectively govern how we should think about customers purchasing workloads.

(166) Customers (regardless of size, or whether they multi-cloud) do not make purchasing decisions with reference to – let alone because of – the input costs of their cloud provider. Rather, customers will *"assess their IT needs on a workload-by-workload basis*".[189] The relative supplier costs are irrelevant and unknown[190] to them.[191] They are focused on the overall value of their contracts for all the cloud services they are likely to use.

#### 10.1.2 The 'large customer behaviour' conjecture in the PDR is unevidenced, unrecognised by Azure, and economically puzzling

(167) The CMA hypothesises that multi-cloud customers often structure a subset of their workloads effectively around a high Microsoft IP component and, because of this demand pattern, Amazon and Google will lack the incentive to compete as hard for these workloads. The PDR says: the "*significance of the licensing input costs in competing for any single*

*at least one other] in an independent decision*" (PDR, paragraph 6.253). Microsoft is in accord: "*Microsoft has submitted that it competes for each workload"* (PDR, paragraph 6.253)

[184] PDR, paragraph 6.267

[185] PDR, paragraph 3.16

[186] If compute were not networked (purchased with networking), it would be an unconnected box in the cloud. That would make no sense: the cloud is about remote access to rental services. If storage were not purchased from the same vendor as compute, it would be in a different location and introduces latency issues. In any event, hypotheticals aside, the Keystone review of Azure data eliminate any scope for debate whether these services could for the sake of argument in fact be purchased separately for normal cloud use, at least when supplied by the hyperscale competitor set of Amazon, Google, Microsoft

[187] PDR, paragraph 6.267

[188] PDR, paragraph 6.266

[189] PDR, paragraph 3.77

[190] Cloud workload pricing is not done on some sort of negotiated open-book "cost-plus" basis as if it were the Ministry of Defence negotiating a nuclear submarine contract with a defence contractor

[191] They may, of course, have regard to pricing in determining how wide to make their workload (e.g. how many services to add to a core set), but the AEC analysis cannot wrongly assume that which it is seeking to test: whether IP input costs actually affect prices paid at the workload- or deal- or customer-spend level, after discounts are applied at any of these levels.

*customer … is likely to be closer to the upper bound*" when it is "*customer groups with a higher propensity to multi-cloud, such as higher spend customers*"[192] For these customers "*Microsoft software likely comprise a larger share of the total spend that large customers allocate to Azure within a single purchasing decision*".[193]

(168) The PDR has no footnotes for these hypotheticals and CMA cites no evidence of examples in support of these assumptions.

(169) This absence of evidence is consistent with Microsoft's own ordinary course business experience that $20+ million spender large customers, as a class, do not disaggregate their decision making in the way the CMA suggests. Nor, in order to compete, does Azure market itself for "*narrow purchasing decisions*" as "*narrowly*" having the USP as a destination for "*Microsoft workloads*" rather than a "*broad*" destination for "*broad purchasing decisions*" of the full suite of Azure services – such as the **[CONFIDENTIAL]** services that on average its own $20+ million spender customers actually consume, almost all of which are not the subject of the SPLA licensing term dispute.

(170) The CMA's logic is also puzzling as a general proposition of economically rational, large-spender, multi-cloud customer behaviour. It is incumbent on the CMA, not least given the central role of this alleged behaviour to its AEC theory, to explain why customers are not more likely to *aggregate* their demand to improve competition than "*likely*" to disaggregate to worsen competition for workload components or narrow (Microsoft-centric) workloads. The CMA acknowledges as common ground that customers aggregate demand further into CSAs to achieve both certainty and better prices, and that both Google and Amazon compete using CSAs.

(171) In addition, the CMA does not explain why Amazon and Google would not rationally have the incentive to encourage customers to structure the aggregation of their workloads in ways which dilute the significance of Microsoft IP cost in the overall opportunity to increase the competitive tension and their potential profits, not concentrate it. Indeed, Amazon and Google have already invested significantly in programmes to reduce the significance of Microsoft IP cost to their workloads, through optimisation and modernisation.[194]

(172) And the CMA has made no attempts to quantify such efforts from Amazon or Google, for example, any alternatives to Microsoft IP which Amazon or Google offer or any efforts to migrate a customer from products for which rivals pay for Microsoft IP to services where they do not (Linux VM-powered, and other substitutes).[195]

### 10.2 In relying on (1) a factual mistake coupled with (2) an unevidenced conjecture, the AEC analysis loses reliability as a likely predictor of foreclosure effects

(173) The first part of Chapter 6 primarily focuses on shares of input costs, for Amazon and Google, of Microsoft IP cost relative to other costs. This is the domain of the RRC test.

---

[192] PDR, paragraph 6.268. The PDR notes, at footnote 1746 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* (i.e. >$20+ million; see also PDR, Tables 6.3 to 6.6) *using more than one cloud*"

[193] PDR, paragraph 6.298(d)

[194] See Modernize Windows Workloads | Amazon Web Services; AWS Optimization and Licensing Assessment | Amazon Web Services

[195] Microsoft notes that such efforts by Amazon and Google would be "rivalry enhancing", which the PDR notes they should take into account (PDR, paragraph 6.45) but does not actually consider. In the confidentiality ring, Google provided customer examples (PDR, paragraph 6.501). While the CMA does not place much weight on these one-off examples, they do illustrate that rivals employ both technical and commercial strategies to compete for WS workloads. **[CONFIDENTIAL]** (See Annex 74.2 of the confidentiality ring)

However, for the reasons given above, the CMA has misidentified the significance of IP costs as a share of total costs by backing out narrow customer demand for workloads that contain a lot of IP cost.

(174) The analysis is, in effect, circular: if a workload only contained Microsoft IP cost it would not be attractive to Amazon and Google and the customer would have foreclosed supplier interest in its own demand. But, as explained, that is not how workload demand is structured, is unevidenced, and would be economically irrational of sophisticated larger-spending multi-cloud customers in particular.

(175) But the CMA correctly goes beyond simply **input cost proportions** in an RRC test to focus on **available margins** for incentives to compete. This is important because, while IP costs for Amazon and Google are variable, cloud infrastructure services are generally low-variable cost, high-margin services. This has been a key point in Microsoft's submissions and is consistent with the CMA's findings elsewhere on suppliers' EBIT and ROCE.

### 10.3 The missing 'foreclosure-free' counterfactual – Google's critical margin needed to compete

(176) On margins, the CMA ultimately concludes as follows: "*we consider that positive and low margins … 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*".[196] In other words, it is not that AWS or Google lose money (negative margins) on this component part of the sale. It is that their positive margins are too low, it is said, for Google and Amazon to compete "*fiercely*".[197]

(177) Both Google and Amazon have respective and bespoke Hyperscale SPLA licensing agreements with Microsoft that cover Windows Server IP. While Amazon is also a complainant, the CMA's Licensing AEC stems from a margin dispute and foreclosure argument put forth by Google. And indeed, if Google's foreclosure case as the #3 is demonstrably shown to be a hollow construct, it can hardly be said that Amazon, the #1, might still be a "foreclosure victim". (The obvious but head scratching implication of the PDR's finding is that in a "well-functioning" market, AWS would / should have yet more market share and more profitability?) For these reasons, Microsoft focuses for present purposes on Google although the same points apply, with even more force, to Amazon.

(178) Google claims that, once it has to pay Microsoft for its own commercial use of that IP, its remaining margins on that one component, while positive, are "*too low*". Google presents the CMA with one "*illustration*"[198] and "*five hypothetical customers*" that "*reflect real-world scenarios*"[199] of this "*too low*" margin phenomenon in action.[200] In accepting Google's argument that its positive margins (on a component, rather than workload) are "*too low*", the Licensing Chapter declines to take a position on at what point Google's margins would be satisfactorily "*high enough*".

---

[196] (emphasis added) PDR, paragraph 6.462

[197] It is not clear whether competing fiercely is the same or different term from the Guidelines' requirement that they be able to compete "*effectively*". CMA's Market Investigation Guidelines: "*By being subjected to higher input prices—of which an extreme form is a 'margin squeeze' —downstream competitors may be unable to compete effectively*" (paragraph 270); see also IP remedies discussion at CC3 paragraph 374

[198] PDR, paragraph 6.501(b)

[199] PDR, paragraph 6.502

[200] (*emphasis added*) The PDR places "very limited weight" on Google's examples because they are "anecdotal" (PDR, paragraph 6.503)

(179) This is not a rhetorical question: it is one of robustness and legality. The AEC analysis avoids any articulation of the alleged counterfactual absent the foreclosure effect, a counterfactual in which there is no foreclosure because Google's margins are "*high enough*" to compete fiercely.

(180) That is, the CMA has failed to articulate a "foreclosure-free" counterfactual. To show a foreclosure effect, the CMA is obliged under the Enterprise Act, as reflected in its own Market Investigation Guidelines, to compare the market outcome with the alleged foreclosure effect against the market outcome without it: *"In identifying some features or combination of features of the market that may give rise to an AEC, the CC has to find a benchmark against which to determine how the market may be judged to be performing. In the absence of a statutory benchmark, the CC defines such a benchmark as 'a well-functioning market'… The benchmark will generally be the market envisioned without the features*."[201]

(181) That is the legal test of causation – a market feature or practice that *causes* an AEC. In the context of licensing, where the AEC is premised on whether hyperscale rivals' margins drive their incentive to compete "fiercely", the foreclosure-free counterfactual (and any hypothetical remedy that generates it) "*has to find a benchmark*": the critical margin threshold at which Google will commence to compete effectively, or "*compete fiercely*", for so-called "*Microsoft workloads*".

(182) In the absence of any such margin benchmark, and considering Google continues to pour billions per quarter of investment into Google Cloud Platform, it is difficult to imagine how much more "*fiercely*" it might compete in its higher-margin counterfactual.

## 11 Strong evidence for a no-AEC finding

### 11.1 The Google Cloud value proposition to customers in IaaS / PaaS does not remotely turn on Microsoft-IP-based IaaS component pricing issues

(183) As noted above, the total workload margin (the "*pot of total gross profit*") is what incentivises competition. The CMA finds that "[*m]ost customers we spoke to said they use a mix of IaaS and PaaS for their cloud workloads*".[202]

(184) A mixture of IaaS and PaaS is also why customers would pick Google. Customers pick hyperscale clouds for their value-added differentiated services as a hyperscaler, not for their ability to offer a bare-bones IaaS workload stripped down mainly or only to Microsoft IP components.

(185) For example, the CMA says: *"Customers identified the following benefits …: '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 … two key challenges to using GCP: its complex billing and the lack of configurability relative to competitors like AWS".*[203]

[201] (emphasis added) CMA's Market Investigation Guidelines, paragraph 320

[202] PDR paragraph 3.45

[203] PDR, paragraph 2.30

(186) At the same time: *"Google's internal documents show that Google's strategic priorities for its cloud business are: Leading in the AI space [ ], winning startups, public sector, FTSE 100 & SI customers; and expanding Google's partners network*."[204]

(187) Neither customers nor internal documents suggest that reselling Microsoft IP via the incorporation of certain IaaS components is a key or untapped part of the Google Cloud value proposition.

### 11.2 Keystone shows that – even as an extreme lower bound or worst-case – an equally-efficient competitor has significant margin on workloads incorporating Microsoft IP

(188) Keystone, retained by Microsoft, provides new evidence annexed to this submission. It is already established that, when properly considered based on what customers actually consume (on average **[CONFIDENTIAL]** services in 2024 for Azure customers who spend $20+ million, and **[CONFIDENTIAL]** services common to each of these customers), Google has large margins with which to compete, and its cloud revenues and growth prove this to be true. This new economic analysis shows that, even assuming an artificially small set of "core" compute services are the ***only*** services used in a workload, as of today Google's total available gross margin on this bare minimum IaaS workload is not "*too low*" to incentivise Google to compete for these three components "*fiercely*" especially considering that the customers will always consume other services in addition to the core IaaS services.[205]

*Workload subset: (1) networking, (2) storage, (3) [WINDOWS SERVER VM]*

(189) Keystone has run a new as-efficient-competitor test analysis on the bare minimum IaaS subset of components needed just to run and use a Windows Server VM, Google storage and Google networking. It is therefore a lower-bound before adding a higher margin on other services (Google's Linux VMs, Google management, Google security, Google PaaS). It also accounts for (1) paying SPLA costs and (2) matching the AHB discount. On just this subset of a workload, the available gross margin for Google is **[CONFIDENTIAL]%+** (AHB customer usage only) and in the **[CONFIDENTIAL]%** range for all customers (all customer spend). The average gross margin on all other IaaS / PaaS component services with no Microsoft IP cost will likely be higher still, and drive up the overall workload- or deal-level gross margin.

*Workload subset: (1) networking, (2) storage, (3) [WINDOWS SERVER VM], (4) [SQL SERVER] PAYG*

(190) Keystone has also run the same analysis adding [SQL SERVER] PAYG.[206] The resulting available gross margin on this subset of workload is in the **[CONFIDENTIAL]%** range for all customers, and still **[CONFIDENTIAL]%** for WS AHB only. Any actual workload with additional IaaS / PaaS will drive up the total workload margin. For reference, Google Cloud's EBIT margins jumped from 3% in 2023 Q1 to 18% in Q4 2024, despite the alleged foreclosure from "cloud services".

---

[204] PDR, paragraph 2.31

[205] This Keystone analysis is part of a standard economic test for foreclosure, the "margin squeeze" or "as-efficient competitor" test. It must hold true based on inferences from Azure costs, even as the CMA awaits Google's own cost data, because Google is not an "inefficient" competitor compared to Azure (Amazon still less so). The evidentiary burden must shift to Google that it is allegedly "inefficient" because its costs are too high, and its gross margins too low, including on storage, and networking, and a variety of other services that go into workloads. It has not done so, Microsoft contends, because it cannot show these things. Otherwise, Google's case would be so much more straightforward and there would be no need for artificial constructs and lenses that appear calculated to produce "foreclosure" optics

[206] Only SQL PAYG is relevant because BYOL means that Google / Amazon do not pay incremental IP cost for other SQL IaaS customers

(191) If Google has the available margin to compete for the worst-case scenario, then it can certainly compete (even more fiercely) for any realistic or representative workload scenario. And if Google is "[able] *to compete effectively*" for workloads, then it is not foreclosed from cloud competition and there is no AEC. Google is, of course, the #3: the same, needless to say, applies to the #1, Amazon.

### 11.3 The economic evidence pointing to no-AEC is strongly corroborated by the lack of any concrete examples from the CMA that demonstrate the supposed harm – tangible foreclosure effects

(192) Inferences can be drawn not only from positive evidence but also the lack of evidence that ought to exist if the AEC hypothesis in question were in fact true. Such evidence can be summarised as follows:

#### 11.3.1 No examples that show – according to the CMA itself – Google or Amazon losing out.

(193) While for good reason placing "*very limited weight*" on anecdotal (or outright hypothetical) examples from Amazon and Google[207] the CMA cites generally its "*own*" customer evidence but then does not identify a single customer example of where the CMA believes that, in its view, that customer would have wanted to choose Google or Amazon but for the (perceived) licensing cost factor, such that this factor was the cause of Google's or Amazon's foregone revenue opportunity.

#### 11.3.2 No material expression of customer concern – according to UK enterprises themselves.

(194) Foreclosure that harms competition, by definition, harms customers. Despite finding that UK customers are harmed, and that UK customers are among its most sophisticated and largest companies, the CMA does not cite any examples of customers submitting that, in their view, licensing factors harmed their ability to exercise choice, or resulted in higher prices or worse terms.

(195) Out of 118 customer recipients of a CMA request for information relating to licensing, the CMA in the PDR cites no examples of customer "concern" in so many words. It cites only one example (at paragraph 6.481) of a customer view that, taken at its highest, suggests licensing made it difficult to choose other hyperscalers – where 'licensing' refers only to the discount that Microsoft offers certain Azure customers on Windows Server VMs, and not to the unavailability or restriction of Windows Server or SQL Server on rival clouds. (And, of course, by "licensing" the reference is to the price to be paid given that Windows Server is available on other cloud providers.) Other customers are cited as saying licensing is "*unimportant*", "very *unimportant*", "*irrelevant*".[208] Only a small number – and not "*many*" as the CMA implies – said it affected choice (without it being clear that they perceived choosing Azure, or choosing an AHB discount, as being a "bad thing").

(196) The CMA dismisses a handful of (predominantly hypothetical) Google and Amazon examples as "anecdotal", noting that: "*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*".[209]

[207] PDR, paragraph 6.503

[208] PDR, paragraph 6.486(b)

[209] PDR, paragraph 6.503

(197) However, this "*direct*" evidence, as summarised above, is no more supportive of a "*representative*" class of customer concern or adversely affected customers either.

(198) The striking paucity of concrete examples in either category above is, to Microsoft's knowledge, thoroughly unprecedented for a market investigation AEC or Phase 2 substantial lessening of competition ("**SLC**") merger case in a B2B market. Taken together, they strongly corroborate the economic evidence on hyperscaler margins, costs, and ability to compete that points to no AEC.

(199) Indeed, even simply on their own terms, they are warning signs to reconsider the robustness of the AEC.

## PART C.II: Supplementary Detail on Licensing Analysis and Evidence

### 12 How customers structure demand into workloads and aggregate workloads into deals and CSAs

(200) Competition in cloud services can be thought of as suppliers competing for "share of wallet" of enterprise customer infrastructure cloud spend, where customer demand is composed, in principle, of contestable workloads that are typically aggregated into contestable customer cloud "projects" for which suppliers offer "deals" and "deal discounts" to win those projects and which are further aggregated into CSAs covering multiple current and future workloads.

#### 12.1 Competition for a larger 'share of cloud wallet' of enterprise customers – which at the largest Azure customer spend bracket is almost exclusively through CSAs

(201) The widest unit of demand is at the total customer wallet level. There is strong competition to have the largest share of customer wallet and increase share including via CSAs that cover multi-year spending spanning a sizeable range of projects / workloads.

(202) The CMA recognises that the use of CSAs is pro-competitive giving both parties certainty and giving customers discounts and does not foreclose other cloud providers (Google uses CSAs just as Microsoft and Amazon do). Of Azure's **[CONFIDENTIAL]** "large" UK customers (with Azure spend of $20+ million), **[CONFIDENTIAL]** had CSAs and **[CONFIDENTIAL]**. On average, these customers consumed no less than **[CONFIDENTIAL]** services, each consumed **[CONFIDENTIAL]** common services, and spent $**[CONFIDENTIAL]** million on Azure in 2024.

#### 12.2 If not incorporated into CSAs, customers typically aggregate workloads into larger "projects" for which suppliers offer "deals"

(203) One customer project may require many jobs / IT needs – so competition regularly happens at a broader level than by workload i.e. by 'project'.

(204) The CMA notes on several occasions: "*A cloud provider also submitted that discounts are **typically at the deal level***".[210] A cloud provider like Google can, and does, include discounts and benefits relating to other business (e.g. advertising credits) to improve the overall deal. Competition for a "deal" could be for a single workload or for a customer project comprising multiple workloads aggregated into one procurement negotiation. This saves the customer transaction costs and total outlay due, in effect, to a "deal" or volume discount across 2+ workloads, even if not aggregated further into a CSA.

(205) For suppliers, discounts are at the deal level, but cash is fungible. The customer cares about the overall price, not the component pricing per se. The discount can notionally be applied to one input and not the other to achieve a workload price, or a deal price could apply across aggregated workloads.[211] For example, Microsoft could remove the AHB discount on Windows Server VMs usage and apply an equivalent discount to the 'storage' component. This would have no impact on the customers' overall price given that storage is always purchased with compute. Similarly, GCP and AWS can discount other services in the workload(s) that comprise the deal negotiation wholly unrelated to Microsoft software.

[210] PDR, paragraphs 6.253, 6.448(a) and 6.458(a)

[211] This is especially true in cloud where services often share the same costs. For example, infrastructure costs (data centres, servers, networking) incurred by a cloud provider often service multiple offerings. As such, a provider is likely to be indifferent between a discount applied to one product in a bundle or another

### 12.3 Competition for individual workloads is for the narrowest discrete unit of demand

(206) While customers will often deploy multiple workloads at one time (i.e. for one 'project' as set out above), the narrowest possible unit of demand, and in turn competition, is the cloud workload.

(207) The PDR defines a cloud 'workload' as a "general term which may mean a customer's application, service, capability or other task or activity".[212] A customer will consume a package of cloud services to carry out a particular job (i.e. to run one of the customer's own applications, services, capabilities, etc.). Or in AWS' words, *"Customers are typically looking to answer an IT need"* and *"assess their IT needs on a workload-by-workload basis."*[213]

(208) A workload will always comprise multiple services to solve a complex cloud IT business need. In principle, every workload is contestable, especially if the non-primary cloud provider has a differentiated offer / expertise (e.g. data, AI, etc).

(209) The PDR is therefore on point when it states that "*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*".[214]

(210) The number of cloud services in a 'workload' depends on the job.[215] However, because a customer needs these services together to carry out its job, a workload is the narrowest possible unit of demand. A customer who wants a car that they can drive, does not order an engine from one provider and a chassis from another. The PDR is therefore also correct when explaining "*it is more rare to split a single workload across different providers*"[216] and acknowledging in Chapter 3 that integrated multi-cloud is rare.[217] There is no evidence of customers segmenting a workload to deploy related Linux VMs in one cloud and Windows VMs in another.

### 12.4 Beyond the narrowest unit of demand, the "bare minimum" workload of compute, storage, and networking

(211) A 'default' or 'bare minimum' workload will always include compute, storage and networking. This is made clear throughout online marketing material – see Figure 3 below.[218] However, this is still hypothetical – workloads will generally include many more services. See Keystone Annex.[219]

(212) As the CMA explains: *"[t]he IaaS layer provides access to raw computing resources (compute, storage and network) for processing workloads and storing data*".[220] Compute, or

[212] PDR, footnote 19

[213] AWS, response to the CMA's working papers and updated issues statement dated 25 June 2024, paragraph 7. See also PDR, paragraph 3.78

[214] PDR, paragraph 6.243

[215] PDR, paragraph 6.243

[215] In the ordinary course, the "bare minimum" or "default" workload includes compute, storage and networking is considered the "bare minimum" or "default" – i.e. the IaaS components. In practice, many other services are added depending on the customer's specific job – for example, management and security. A workload may include PaaS or SaaS services mixed in with the IaaS components

[216] PDR, footnote 117

[217] PDR, paragraph 3.354(b)

[218] See https://learn.microsoft.com/en-us/azure/architecture/reference-architectures/n-tier/windows-vm and https://learn.microsoft.com/en-us/azure/architecture/virtual-machines/baseline

[219] See Section 4 of Keystone Annex. This shows that large customers, in particular, consume ever more services in addition to WS VMs, with **[CONFIDENTIAL]**

[220] PDR, paragraph 1.16(a)

VMs, is one of the "*raw*" computing resources, or IaaS components, alongside storage and networking. The CMA correctly finds that "*all IaaS providers supply standard compute, storage and networking* [and] *that competitive conditions are* [not] *different for each element*".[221] The CMA also refers to "*core services*" as the "*main IaaS and PaaS services that contribute to the key objectives of the customer's workload*"[222] (emphasis added) but does not confuse core IaaS services, as contributing components, with the workload itself.

(213) Google describes compute, storage and networking as "*basic services*",[223] while Oracle describes them as "*fundamental*" along with "*identity* [and] *security*".[224] Azure explains to customers that a VM "*requires additional components besides the VM itself, including networking and storage resources*" (see. Figure 3 below). Adding these is known as "provisioning" the VM.

(214) A 'workload' will, therefore, never be limited to the single infrastructure component that is a Windows Server VM. This does not carry out any conceivable customer job (much like a mere car engine does not allow a customer to drive), and so therefore no customer buys it alone.

**Figure 3: Running a Windows Server VM on Azure**

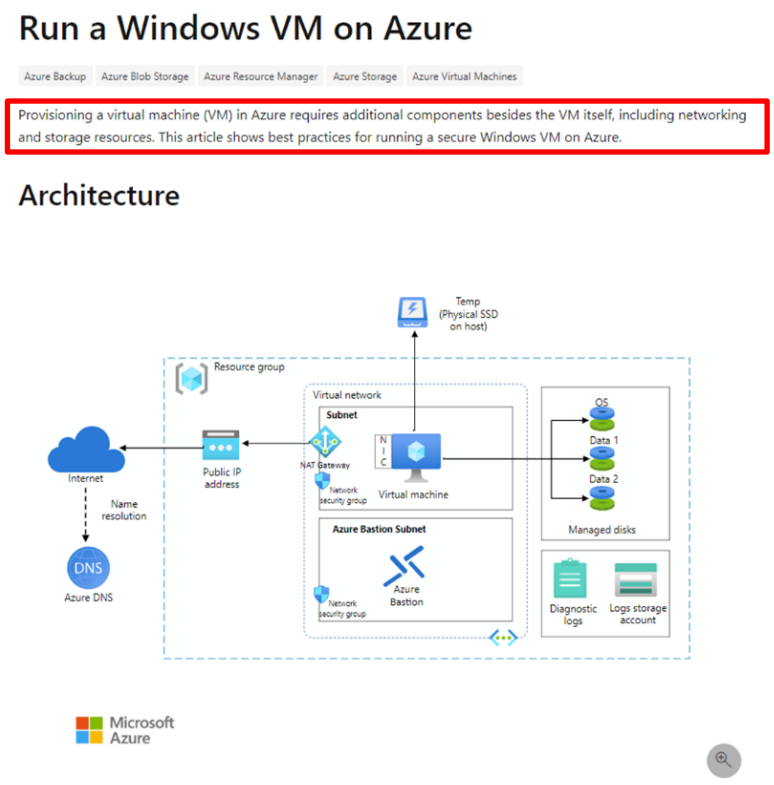


Source: Run a Windows VM on Azure - Azure Architecture Center | Microsoft Learn

(215) No Azure UK customer in 2023 or 2024 bought a [WINDOWS SERVER VM] on its own. This is also the case for multi-cloud customers (see Keystone Annex). No cloud supplier "competes" to supply a [WINDOWS SERVER VM] on its own - they compete for workloads and given the

[221] PDR, paragraphs 3.16-3.17

[222] PDR, footnote 915

[223] PDR, paragraph 4.65(c)

[224] PDR, paragraph 4.86(b)

"bare minimum" composition, Google and Amazon will know that, in addition to any margin on a compute component like [WINDOWS SERVER VM], they can make additional and high gross margin on at least the two essential and default IaaS complements: storage and networking. That is, their margins will be at the very least on these three core complementary and integral IaaS services.

## 13 Some distortive effects of the "Windows Server VM" component lens

(216) The CMA claims to understand and acknowledge that competition occurs for workloads at the narrowest, and more typically for broader customer "projects". Windows Server VMs (i.e. compute) are a mere component of a workload. Yet the CMA frequently relies on a narrow "Windows Server VM" lens to build its partial foreclosure case.

(217) This leads to distortions in Amazon and Google's favour. Set out below are select important examples.

### 13.1 Distortion #1: CMA suggests that Windows Server VMs are the narrowest unit of demand (for at least a subset of customers)

(218) As set out in Part C.I, the CMA analysis refers to the proposition: "*[A] Windows Server VM is the narrowest set of downstream products that Windows Server can serve as an input for.*"[225] However, serving as an input for a component does not make the component a workload any more than a fuel injection systems makes an engine a car. It is just an input for a component.

(219) This fallacy serves as a worst-case scenario or upper bound of the significance of Windows Server IP input costs to downstream competition for demand. This is purely hypothetical and misleading since, as explained above, no customer purchases only Windows Server VMs. Accordingly, this is not a legitimate sensitivity check and is distortive as to the materiality of Windows Server for customers. This distorts the results in Figure 6.1, and Tables 6.3 – 6.6.

**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]** |

(220) The CMA suggests that large multi-cloud customers might disaggregate decision making, and therefore may buy just Windows Server VMs. It is not disputed that "*customers with a high propensity to multi-cloud are less likely to decide which cloud to deploy all their workloads in a single decision*".[226] However, the CMA then says that "*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*". The logic appears to be that these customers demand (and cloud providers compete for) individual 'services' (such as Windows Server VMs).

(221) This is inconsistent with the CMA's findings about how cloud competition works. Customers who multi-cloud may deploy workloads on different clouds, but the workload remains the

[225] PDR, paragraph 6.267

[226] PDR, paragraph 6.252

narrowest unit of demand. It is intuitively impractical for a multi-cloud customer to disaggregate decision making level to the workload component level - e.g. deciding to purchase its storage from Azure, but its compute from AWS. This would be equivalent to a driver who wants to drive in three cars, buying three engines from BMW, and three chassis from VW. That driver is not able to drive, nor is the multi-cloud customer able to cloud. The only scenario where this analysis would hold is integrated multi-cloud. But as the CMA correctly notes, integrated multi-cloud is vanishingly rare.[227]

(222) Nor has the CMA presented any evidence that this component level decision-making happens. Instead, Microsoft's own positive evidence from its multi-clouding customers (see the Keystone Annex) shows that customers who multi-cloud purchase multiple services.

(223) This is also an important error in AWS and GCP's favour. The CMA appears to conclude that AWS and GCP are foreclosed from at least for a subset of customers – large users of multi-cloud. This distorts the conclusions drawn from Figure 6.1, and Tables 6.3 – 6.6.

### 13.2 Distortion #2: Sensational retail and wholesale 'margin squeeze' price comparison of Windows Server VMs

(224) The CMA takes the purely hypothetical demand for just the WS VM infrastructure component and compares the price between:

(i) the Azure retail price for WS VM, but with the maximum conceivable discount (AHB for 100% of usage – which itself distorts reality);

against

(ii) the SPLA wholesale IP price that AWS / GCP must pay for WS, in order to compete for that hypothetical purchase.

**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]** |

(225) There is nothing surprising in Table 6.7 of the PDR. Microsoft's discount leads to a price difference between what an AHB customer pays, and what Google and Amazon pay (as well as non-AHB customers). And, because Windows Server is naturally a very material input cost to Windows Server VMs, the difference looks very high.

(226) The analysis might be useful if competition were all about the component – i.e. if AWS and GCP were pure resellers of Window IP use rights. But it is misleading (and, at 1,000% deltas, radically so) as a guide to whether AWS or GCP are being foreclosed from what everyone agrees is the narrowest unit of demand, i.e. *workloads*. Again, it is a critical error in AWS and GCP's favour.

[227] PDR, paragraph 2.84

### 13.3 Distortion #3: Windows Server VM usage shares

(227) The CMA shows that Microsoft has the lion's share of Windows Server usage in the cloud. This is not disputed, and nor is it evidence of competitive harm. 'WS VM usage' is not a lens that measures competitiveness.

(228) First, as the CMA notes, there are multiple factors that drive such differences in usage.[228]

(229) Second, it is unsurprising that Microsoft is winning workloads that involve its software from customers who are on average bigger users of Microsoft software. These are entities with whom Microsoft will typically have close and long standing relationships. Nor should be it surprising that Microsoft succeeds with some customers for workloads that involve a component for which it is choosing to offer a discount. This is no bad thing – it is good evidence of price competition. AHB, as a discount, is used because workload and deal competition is strong and because that strong competition is not (merely) for the discounted component. And critically, Amazon and Google are capable (and indeed, do) compete with their own discounts – either on the [WINDOWS SERVER VM] component or any other component.

(230) Finally, the PDR ignores the fact that Amazon and Google are constantly moving customers off Windows Server. The cloud provides very little benefit to customers when they elect simply to replicate on premises Windows Server workloads. Rather, customers get the greatest benefit from a hyperscale cloud when they modernise and update their IT infrastructure to use more cloud native technologies. As Gartner notes, few strategic global hyperscalers combine the breadth and depth of capabilities necessary to meet any enterprise where it is, help it determine where it's headed, and provide the cloud services and transformation support to help it on its strategic cloud journey.[229] These cloud journeys involve technology modernisation, replacement, cloud-native adoption, cloud innovation, and business transformation[230] – but not mere rehosting of on-premises workloads. Amazon and Google both have high incentives to help accelerate customers' cloud journeys, moving them off Windows Server more quickly, including by modernising workloads before bringing them to the cloud. Both companies advertise these offers.[231] Even on Azure, Windows Server is becoming less relevant over time.

### 13.4 Distortion #4: CMA's margin analysis for "Microsoft Workloads"

(231) Microsoft agrees that an analysis of alleged foreclosure victims' margins is highly instructive for the AEC analysis.

(232) The CMA's margin analysis takes the following steps:

(i) Takes revenues and costs data from Google, "*incurred directly from licence sales of Windows Server and SQL Server*".[232]

(ii) Uses that data to calculate the estimated margin on "*Windows Server workloads*".[233]

---

[228] PDR, paragraph 6.475

[229] See Gartner 2024 Magic Quadrant for Strategic Cloud Platform Services, available here: https://www.gartner.com/doc/reprints?id=1-2J4WMCRK&ct=241021&st=sb

[230] *Ibid*

[231] See Pricing Overview | Google Cloud and AWS Windows Migration Accelerator - Amazon Web Services

[232] PDR, paragraph 6.459

[233] *Ibid*

(iii) Finds that the "*calculated margins and mark-ups do not appear to be particularly high*" and therefore "*AWS and Google are less likely to compete fiercely for Microsoft workloads*".[234]

(233) The change in the CMA's nomenclature at each of stage of the analysis is potentially misleading. The CMA does not indicate that it received Google's data on revenues and costs other than for the licence sales of Windows Server and SQL. Therefore, our assumption is that the CMA's margin calculation is only on the Windows Server or SQL component of workloads. However, the CMA refers to a margin for 'Windows Server workloads' and 'Microsoft workloads'.

(234) For Microsoft's own analysis of margins, that Amazon and Google could make over workloads with a [WINDOWS SERVER VM] component, see Section 14 below and the Keystone Annex.

(235) Labelling the putative class of "foreclosed workloads" where Amazon and Google allegedly cannot or will not compete hard as: workloads-with-high-Microsoft-IP-input-costs-for-Amazon-and-Google, would sound contrived – although it would in truth be the most straightforward account of the unit of demand considered.

(236) The labelling of these IP component costs as proxies for customer workloads makes them sound superficially more realistic. For example: 'Windows Server workloads', or 'SQL Server workloads', or simply 'Microsoft workloads'. To label a complex product based on one dimension (i.e. the relevant operating system or database infrastructure service) becomes dangerously misleading when it colours the assessment of the whole, as it has in the Licensing Chapter.

(237) Consider that at the automotive component level, if a vehicle is one-dimensionally called a "diesel vehicle" because it runs on diesel and has a diesel engine that includes some diesel IP (patents), it does not mean that the diesel engine component(s) are most of the cost, and most of content, of the car or the truck or the van. Nor does it mean that the diesel-technology-specific engine components (containing the IP input) are a large share of the cost of the car or truck or van.

(238) If a workload is one-dimensionally called a *Microsoft workload* because it includes *some* Microsoft software components, it does not follow that the workload is 60%, 70%, or 80% Microsoft IP cost. It only means that there is a Microsoft software component. And as noted, the CMA cites no evidence that customers structure their demand (consciously) in ways that would foreclose Amazon and Google.

## 14 Amazon and Google can profitably compete for workloads on any sensible worst-case lens

### 14.1 The meaning of 'profitably compete' in cloud

(239) The concept of "*profitably compete*" does not relate to the preferred margin Amazon or Google would make on one component in the workload, namely the rental of a VM. This is one infrastructure component of demand, but not a discrete one.[235]

[234] PDR, paragraph 6.462

[235] Most people do not seek to negotiate the price of infrastructure components, such as plumbing, when comparing builders' quotes for renovating their house, even if the quote has a line item for the hot water system, and even if they chose the type of boiler or heat pump

(240) Profitable competition instead relates to positive gross margin over discrete units of customer cloud demand *as those units of demand are actually structured by customers* and for which suppliers actually have an opportunity to compete. As described above, these are at an absolute minimum:

(i) a single cloud workload, comprising consumption of multiple component cloud services such as:

(a) compute, i.e. Windows Server VM rental or Linux VM rental;

(b) storage (**where Amazon and Google will make high gross margins**);

(c) networking (**high gross margin**); and

(d) other types of cloud services (**high gross margin**),

(ii) but in practice often: "*projects*" or "*deals*" or "*committed spend*" that comprise *multiple* workloads over time (aggregates of high gross margin component services).

(241) The question is then as follows: can Amazon and Google profitably compete with Azure (i) at the narrowest, for workloads or (ii) more appropriately, for deals?

### 14.2 Keystone analysis: an as-efficient-competitor <u>can</u> profitably compete for even the <u>narrowest</u> subset of cloud workload components necessary to use Windows Server at all: VM, storage, and networking

(242) Keystone have been able to analyse Azure cost data in the standard "as-efficient-competitor" or "margin squeeze" test, developed in the abuse of dominance context.[236]

(243) This analysis shows that Azure would have a positive gross margin if it paid the licensing cost of AWS / GCP on Windows Server, <u>and</u> applied the AHB discount, when competing to supply the following discrete units of demand at the narrow extreme – taking the "bare minimum services" of (1) compute (the WS VM), (2) storage, and (3) networking – Azure's margin would be **[CONFIDENTIAL]** in the UK.

(244) Table 9 in Keystone's Annex – the most conservative lens, since it accounts only for AHB usage – is worth repeating here in full:[237]

**Table 1: As-efficient competitor results for WS AHB UK customers, 2024**

| | | **[A] WS AHB, storage and virtual network (i.e. the "bare minimum" workload)** | **[B] WS AHB, storage, virtual network and bandwidth** |
|---|---|---|---|
| *Revenues* | Windows Server AHB | **[CONFIDENTIAL]** | **[CONFIDENTIAL]** |
| *Revenues* | Storage | **[CONFIDENTIAL]** | **[CONFIDENTIAL]** |

[236] It is probative to use Azure cost data as a proxy because it would be untenable to suggest that Amazon, the long-term #1, is not "as efficient" as Microsoft. Its economies of scale and scope surely exceed those of Microsoft. Google does not have the same position as Amazon, but is not realistically an "inefficient" new entrant or small competitor for whom the test is ill-suited. There can be no serious debate that there are high gross margins on infrastructure components such as storage and networking, and they will be high for Amazon, high for Google and high for Microsoft

[237] Table 9 in the Keystone Annex is arguably more conservative still since it only includes customers whose workloads include only WS VMs. However, as Table 9 shows, the margins are even greater for these customers – **[CONFIDENTIAL] for the "bare minimum" workload**

| | | | |
|---|---|---|---|
| *Revenues* | Virtual network | **[CONFIDENTIAL]** | **[CONFIDENTIAL]** |
| *Revenues* | Bandwidth | **[CONFIDENTIAL]** | **[CONFIDENTIAL]** |
| *Hypothetical Costs – IP* | Azure WS AHB vCore hours | **[CONFIDENTIAL]** | **[CONFIDENTIAL]** |
| ***Margin*** | Margin over licensing costs and other COGS | **[CONFIDENTIAL]** | **[CONFIDENTIAL]** |
| ***Margin*** | Margin over licensing costs and other COGS (%) | **[CONFIDENTIAL]** | **[CONFIDENTIAL]** |

Source: see Keystone Annex

(245) The as-efficient-competitor test asks whether the pricing practices of an allegedly dominant operator make it "*very difficult or practically impossible for that operator to offer its goods or services in the market at a profit*" – correctly cited by Google.[238] The margins above are clearly inconsistent with an as-efficient-competitor finding it "*very difficult*" or "*practically impossible*" to "*offer its goods or services on the market*". Amazon and Google may prefer to have more margin, but that preference does not merit regulatory intervention.

(246) Azure only fails the as-efficient-competitor test in the hypothetical scenario where the customer has discrete demand for one infrastructure component in isolation – a Windows Server VM – and is forced to match Microsoft's discount on that component. But as described above, there is no market, sub-market, discrete share of demand, or contestability, for this unit of demand. A VM is in fact commercially worthless without storage and networking. And therefore it cannot be used to assess whether Google or Amazon can "*offer their goods or services in the market at a profit*".

## 15 The CMA provides no representative evidence of customer detriment

(247) The standard economic and legal benchmark for an AEC conclusion based on foreclosure effects is not limited to financial harm to rivals but also whether it will harm customers.[239] Customer views on whether they are harmed, while not dispositive, are important to an AEC conclusion that foreclosure effects exist in the market that harms those customers. This is particularly so in the case of large enterprise customers who can be expected to have – and voice – their views. In this investigation, potentially affected customers include the UK's largest public and private companies, and government departments, spending millions on IT, and who are certainly *capable* of expressing concerns.

(248) The CMA has carried out detailed market testing and sought to engage with a number of third parties. The CMA has engaged with a large number of enterprise customers over the course of the MIR.[240] With respect to licensing practices alone, the CMA has told Microsoft

[238] Court of Appeal, Royal Mail v Ofcom (2021), EWCA Civ 669, paragraph 73. See also Google Hearing Follow-Up (23 August 2024), footnote 34

[239] (*emphasis added*) CMA's Market Investigation Guidelines, paragraph 270. Financial "*harm*" to rivals without harm to the process of competition that adversely impact customers is not actionable anti-competitive "*foreclosure*" for AEC (or SLC) purposes because it does not "*weaken the constraints that the [firm in question] faces and as a result harm competition and therefore customers*"; see CMA Merger Assessment Guidelines 2021, paragraph 7.2

[240] The PDR sets out that the CMA has made over 300 information requests to over 100 stakeholders, held more than 70 calls and received over 50 responses from 3 consultations

that from November 2023 to July 2024, it sent six sets of customer questionnaires, amounting to 175 requests for information that went to 118 customers, of which 12 were public sector bodies.

(249) However, responses from only a small portion of these 118 customers are cited as referring to the discounts available on Azure when using services that incorporate Microsoft software, or licensing as a relevant factor influencing their choice.

(250) The CMA cites certain **[CONFIDENTIAL]** customers that refer to discounts available on Azure when using services that incorporate Microsoft software, of which the majority **[CONFIDENTIAL]** refer to the AHB discount.[241] However, competitive pricing is not indicative of harm to consumers. And it is far from clear that any of these consumers have expressed the view that Microsoft's discounts cause them harm. There is only one large customer (out of the 118) that is cited as saying that it is expensive and difficult to use Microsoft products on a competitor cloud.[242] Most competitors did not identify any quality differences in using Microsoft products across different clouds.[243]

(251) The CMA claims that customer evidence is "largely consistent" with Microsoft's licensing practices having an effect on their choice of cloud provider.[244] However, neither the number of respondents cited, nor what they've been cited as saying, suggest licensing was a choice driver – let alone the decisive one. Of the 118 customers contacted:

(i) Only **[CONFIDENTIAL]** customers are cited in support of the CMA's statement that "*Many customers said that licensing terms affected their choice of Microsoft workloads*";[245]

(ii) Only **[CONFIDENTIAL]** customers are cited to support the statement that "most customers" identified licensing as a causal factor affecting their choice;[246] and

(iii) Only **[CONFIDENTIAL]** customers are cited as explicitly mentioning the ability to bring on-premises licences to the public cloud was important.[247]

(252) Given overlaps in citations, there are **[CONFIDENTIAL]** distinct customers cited in this section. This leaves **[CONFIDENTIAL]** customers that the CMA contacted who either said licensing was not a (material factor) or did not respond to this question.

(253) The CMA has (correctly) dismissed evidence from Amazon and Google as being anecdotal and not representative of customer choice, choosing to place greater weight on the customer evidence gathered directly.[248]

(254) However, it is far from clear that the examples cited from the CMA's own evidence are representative either of material customer concerns or a material adverse impact on customer choice as a result of the SPLA licensing terms on offer to Amazon and Google.

[241] PDR, paragraph 6.480

[242] PDR, paragraph 6.481

[243] PDR, paragraph 6.482

[244] PDR, paragraph 6.483

[245] PDR, paragraph 6.483(c)

[246] PDR, paragraph 6.483(a)

[247] PDR, paragraph 6.483(a)

[248] PDR, paragraph 6.503

## 16 Adding SQL Server does not change the foreclosure analysis

(255) The CMA argues that the Windows Server VM focus of Microsoft misses the fact that many customers use both, and that the SQL Server database product is only really used in conjunction with a WS VM; it is not used with Linux.

(256) For customers *with* SQL Server on-premises licences, they have license mobility and can bring their pre-existing SQL Server license to Amazon or Google's clouds free of charge. Therefore, while the SQL data may be relevant to "Microsoft-IP-relevant-Azure spend" (which should look at both) but it is *irrelevant* to the analysis of AWS / GCP foreclosure based on Azure spend. AWS and GCP incur no incremental SPLA cost over and above the WS SPLA cost to serve the SQL component when there is BYOL – and therefore no cost differential, and therefore no raising of rivals' costs.[249]

(257) For customers without SQL Server on-premises licences and who consume SQL on a PAYG basis, those customers are typically choosing between many different database options in the cloud of which SQL Server is not the most selected. Even for those customers that choose SQL, Keystone's analysis shows that the as-efficient-competitor test is satisfied even on a purely hypothetical 'worst-case' workload of Windows Server VM, storage, networking, and SQL Server PAYG, i.e. Amazon and Google can clearly and easily "compete effectively" for even a workload which includes both a Windows Server VM and SQL Server PAYG. This analysis is reflected in the sections above.

## 17 Adding select Microsoft SaaS does not change the foreclosure analysis

(258) Equally surprising are complainant assertions that Microsoft SaaS offerings (Windows, Office and Visual Studio) wholly unrelated to cloud infrastructure (IaaS services) further worsen the foreclosure in IaaS.

(259) The PDR recognises that the Microsoft products Windows 10/11, Microsoft productivity suites and Visual Studio "*do not constitute important inputs on their own*" into cloud services but contribute to the foreclosure concerns. As they are not important inputs, there is no stand-alone competition issue of foreclosure. As the main foreclosure case is flawed, their addition makes no difference to the analysis and does not compensate for the flaws.

(260) More broadly, each of these are SaaS offerings that do not impact the choice of IaaS provider. Just like a user could license Gmail or Google Workspaces, these could be used to provide a user with a desktop. The same is true of developer tools which do not in any way prohibit and distort customers' choice of IaaS provider.

(261) While it may be true that Google would like to offer customers a VDI solution, they can and should do so using their own operating system. The same way that Google does not share its internet search SaaS service to power Bing, Microsoft should not be forced to power Google's VDI solution for Google in the name of solving an alleged foreclosure problem in a cloud infrastructure services market investigation.

[249] The CMA suggest that Microsoft's analysis "*does not consider that the importance of some of the Microsoft software may increase in the future as customers migrate on-premises workloads to the cloud*" (PDR paragraph 6.259) but the CMA ignores this fact in its consideration of SQL Server. For these potential customers, there is no foreclosure effect on AWS / GCP

keystone

# Memo

To: Competition and Markets Authority

From: Keystone

Date: February 23, 2025

Subject: Keystone analysis for response to licensing PDR (**contains confidentiality ring material**)

---

*Confidentiality ring material is redacted in blue.*

## 1.0 Introduction

1. This note provides additional economic evidence supporting Microsoft's response to the licensing chapter of the CMA's Provisional Decision Report ("PDR"). It should be read in conjunction with Microsoft's response, and it focuses on the shortcomings in the CMA's assessment of alleged partial foreclosure resulting from the terms under which Microsoft licenses its software to AWS and GCP.
2. As set out in Microsoft's response to the PDR, the CMA has incorrectly focused on a single IaaS "compute" component of workloads: Virtual Machines ("VMs").[1] It does not account for the fact that compute is just one of the components involved in any workload.[2] Within this narrow lens, the CMA has focused on VMs that run Windows Server ("WS"). In the UK in 2024, 55%[3] of VM vCore hours on Azure ran on Linux OS rather than WS OS. These are contestable without any licensing of Microsoft software by AWS and GCP, who have been growing their cloud business.[4]
3. We show that Microsoft's licensing practices, particularly with regards to WS and SQL Server, do not foreclose AWS and GCP from competing in the cloud.
4. **First**, contrary to the CMA's hypothesis,[5] the evidence shows that [redacted] WS-using customers.[6] That is, WS customers [redacted].

---

[1] Although SQL can also be provided as a PaaS service, the CMA has focused on SQL IaaS.

[2] VMs and VM Licenses accounted for 27% of Azure revenue in the UK in 2024 and 29% in 2023.

[3] 45% of compute vCore hours in the UK in 2024 had WS as an OS. The remaining had Linux OS. Similarly, in 2023, 50% of vCore hours were with WS OS and the remaining had Linux OS. In the period May-December 2022, 51% of vCore hours were with WS OS and the remaining had Linux OS.

[4] *See* Microsoft's response to the PDR.

[5] CMA PDR, paras 6.250-254.

[6] These are UK Azure customers with positive WS vCore hours in a given year. They are referred to as "WS customers" in the rest of this document.

5. **Second,** *even if* the price comparison analysis conducted by the CMA[7] accurately estimates the discounts that AWS and GCP need to offer to compete with Azure's discounts, the PDR does not adequately explain **why AWS and GCP could not provide these discounts to win WS workloads**. The PDR fails to properly consider that customers (including multicloud customers) consume many more services than just WS VMs, despite other portions of the PDR on committed spend discounts recognizing that competition is for total spend by customers and not separated on a service by service basis.[8] As a result, even if competitors accept low margins on single components (WS IP or SQL Server IP), this does not impair their ability to compete for the overall workload.[9] This is especially true when multiple workloads make up a single deal (such as Committed Spend Agreements ("CSAs"))[10] and the overall deal is very profitable.[11]
6. **Third,** for the same reason as above, it is likely that the **input cost share of WS IP in the overall workload (or deal) will be closer to the CMA's lower bound estimates in the vast majority of cases.** Even if the input cost share may be considered "high", cloud service providers may not have an incentive to pass it on unless the input cost is high enough for a workload. That is to say, it is the costs of the workload, rather than a component, that are relevant.
7. **Fourth,** we find that the as-efficient competitor test is passed even for conservatively designed workloads (which we note are still narrower than what happens in practice). According to the PDR, positive margins are not inconsistent with foreclosure.[12] However, the PDR does not indicate a reasonable margin for Microsoft to leave for its main rivals to be able to compete effectively. Any margin estimates on WS VMs should be considered along with (high) margins on other cloud services that are closely related to Windows VMs[13] and generally the lack of evidence of foreclosure for AWS and GCP.
8. **Fifth,** we show that no foreclosure concerns arise for SQL IaaS PAYG customers.[14] We show, with an as-efficient competitor test incorporating SQL IaaS PAYG and WS, that an as-efficient competitor is able to compete with Azure even after paying the Hyperscale Service Provider Licence Agreement ("SPLA") costs for both SQL Server and WS.

---

[7] CMA PDR, from para 6.351 onwards.

[8] CMA PDR, para 7.87 for example.

[9] This is the same approach the CMA used to quantitatively assess the effect of Committed Spend Discounts (CSDs) on rivals in Appendix U.

[10] CMA PDR, para 7.40. The CMA shows that AWS' and Azure CSAs cover a significant portion of cloud spend.

[11] As mentioned in the main response, all $20m+ Azure customers have active CSAs or large agreements with Azure, as of 2024. These customers negotiate based on the entire deal rather than specific workloads.

[12] CMA PDR, para 6.439.

[13] For example, the CMA estimates that [redacted] (CMA PDR, Appendix Q, para Q.86).

[14] We focus on SQL IaaS PAYG usage as customers can BYOL their existing SQL Server licenses to AWS and Google through license mobility (CMA PDR, footnote 1740).

## 2.0 AWS and GCP can and do win significant workloads of multicloud Azure customers

9. The CMA argues that multicloud customers (more likely to be larger customers) are less likely to choose a single cloud provider for their entire cloud spend. Instead, they may make sequential cloud decisions on a subset of their cloud spend.[15]

10. In the following section, evidence on WS spend shares of multicloud customers shows that the majority of workloads for these customers are still contestable. AWS and GCP can still compete for these customers who may make cloud provider decisions on a subset of cloud services. Further, the WS spend share of multicloud customers indicates that WS VMs is just one component of WS-related workloads.

### 2.1 EVIDENCE SHOWS AZURE WS CUSTOMERS SPEND A MATERIAL SHARE OF THEIR WALLET ON RIVALS

11. The evidence from [redacted].

12. The figure below shows the share of wallet spent on rival hyperscalers and how it varies by WS spend share (i.e. spend on WS VMs as a share of total spend on Azure). Customers [redacted].

**Figure 1: [redacted]**

## 3.0 Cloud providers cross-sell even within WS workloads: no plausible use case or demand for WS-VM only workloads

13. The PDR overlooks the fact that sophisticated enterprise customers always consume more than just WS VMs.[16] As a result, rivals would be able to and have an incentive to accept low margins (and even losses if that were necessary) on a single component (WS IP) if the overall workload is profitable. This is especially true when multiple workloads make up a single deal (CSA), and the overall deal is profitable.

14. We agree with the CMA that competition in the cloud occurs for deals that may comprise multiple workloads.[17]

15. Cloud providers compete to win customers' workloads and offer many different services from which the customer can choose to meet the workload's requirements. Customers create bespoke combinations of different services that they consume. WS VMs are just one of many services that customers consume for a workload.

16. We provide evidence below showing that:

    a. No cloud customers use WS VMs in isolation.

---

[15] CMA PDR, paras 6.250-254.

[16] Aside from the fact that additional services must always be consumed to be used with a VM, customers also have an incentive to use multiple services together due to volume discounts in the form of CSAs (CMA PDR, para 7.124a).

[17] CMA PDR, para 6.253.

b. Even for multicloud customers, WS is just one component of a wider use of cloud services.

c. Even for plausibly WS-intensive workloads on Azure, such as those that only use WS compute (and no Linux), WS VMs are still only a component of a wider consumption of Azure cloud services. Considering customers who use only WS VMs and no Linux VMs on Azure, they consumed, at a minimum, 3 Azure services and, on average, 12 services in both 2023 and 2024.[18] And all of them consumed storage in addition to WS VMs.

17. Taken together, the results of the analysis in this section show that cloud providers compete to provide many services to cloud customers. There is no segment of cloud customers, as explored in the Azure data and shown below, where this is not the case. WS is but a subcomponent of a customer's demand which cloud providers compete for.

### 3.1 CLOUD CUSTOMERS DO NOT USE WS VMS ALONE

18. **In both 2023 and 2024, none of Azure's UK (active) customers paid for the use of only VMs (whether using the Windows or Linux OS).**[19] Even when focusing only on WS customers, we observe that they consume multiple services.

19. Restricting the analysis to WS Azure customers shows that they consume at least the same number of services, if not more, than the average Azure customer.[20] This is shown in Figure 2 for all WS customers and Figure 3 disaggregating the same data by revenue categories.

20. Figure 2 shows that when UK customers deploy WS on Azure, they typically consume an average of ~13 services in 2023 and ~14 in 2024. And Figure 3 shows that this is even higher for larger customers (who are more likely to multicloud): **WS customers who spent over $1 million on cloud services in 2024 consumed between 46 and 79 services on average.**

[18] Note that very few customers consumed only 3 services; in 2023, it was 3 customers and in 2024, 2 customers (0.08% of customers who use only WS VMs in 2023 and 0.05% in 2024; this comprised 0.03% and 0.02% of the revenues of customers who only use WS VMs in 2023 and 2024 respectively).

[19] This also includes the Service Level 2 "Virtual Machines Licenses" which covers IP payments for VMs. Hence, we have included customers who either: only pay for the use of VMs, only pay for the use of VM Licenses, or only pay for the use of VM and VM Licenses (and no other Service Level 2). The field "Service Level 2" refers to the level of granularity at which Microsoft aggregates various services in its internal financial systems. Service Level 2 categories broadly align with how services are listed on Azure's website. Active customers are defined as those with at least $1,000 of Azure spend in each year; customers with annual spend below $1,000 are likely to be trials rather than deployments.

[20] In 2024, for example, the average customer consumed ~11 services.

**Figure 2: Average number of services consumed for WS UK customers, 2020-24**

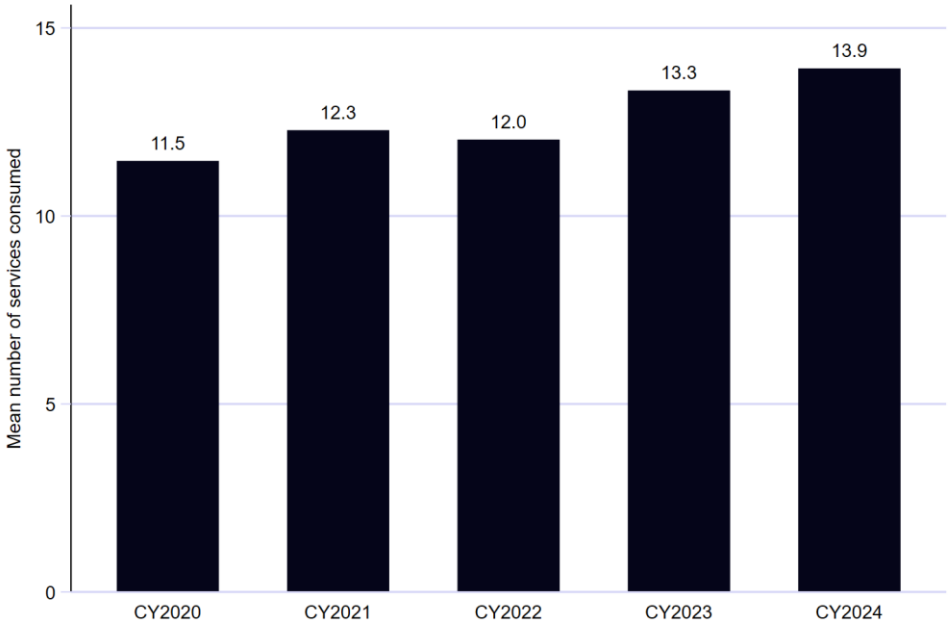


*Source: Keystone analysis of Microsoft data. Notes: This chart shows the mean number of Service Level 2 (SL2) services consumed by active Azure customers in the UK between years 2020 and 2024. Active customers are defined as those with at least $1,000 of Azure spend in each year; customers with annual spend below $1,000 are likely to be trials rather than deployments. Customers are classified as WS users based on having positive WS vCore hours in the given year. This overstates the usage of WS prior to May 2022 since WS and Linux hours are commingled for reserved instances in these years. These hours have been fully attributed to WS on a conservative basis. Further, prior to March 2021, WS and Windows Client hours are commingled. A small number of customers (<10) are excluded if their WS spend exceeds total Azure Consumed Revenues ("ACR") due to an anomaly in the data. The prefix "CY" indicates "Calendar Year" for emphasis.*

**Figure 3: Average number of services consumed by WS UK customers, by revenue category, 2020-24**

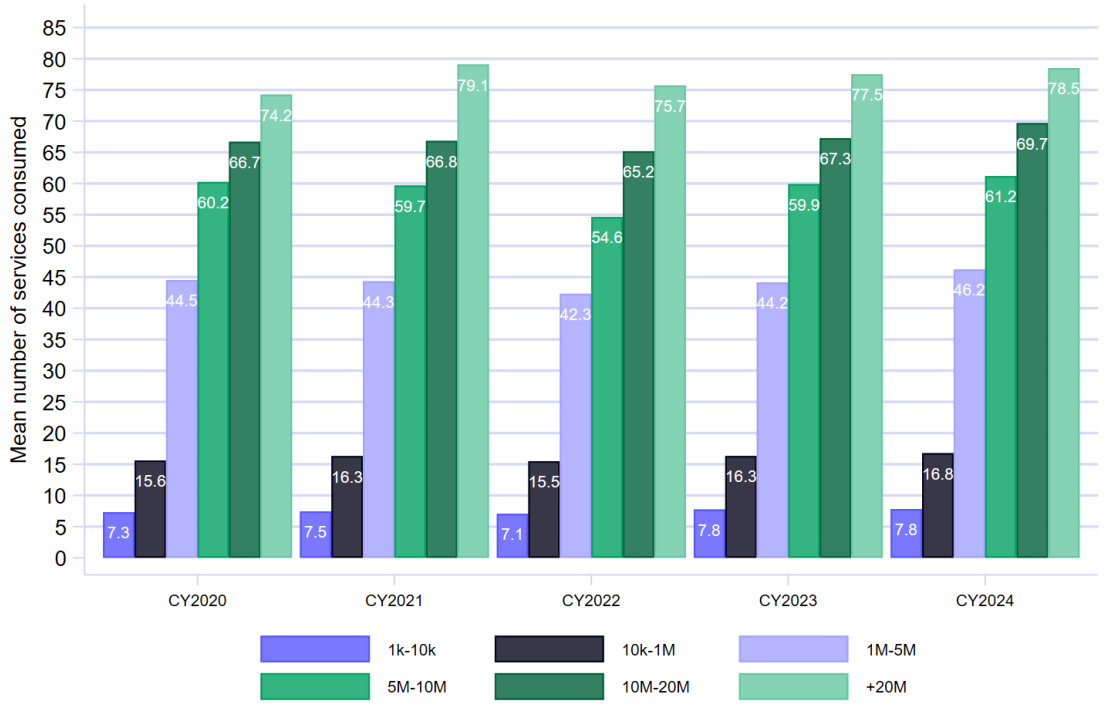


*Source: Keystone analysis of Microsoft data. Notes: This chart shows the mean number of Service Level 2 (SL2) services consumed by active Azure customers in the UK between years 2020 and 2024 by revenue category. Service Level 2 is the level of granularity with which Microsoft records data in its financial systems. Revenue categories are determined by the total cloud spend of each customer in each year. Active customers are defined as those with at least $1,000 of Azure spend in each year; customers with annual spend below $1,000 are likely to be trials rather than deployments. Customers are classified as WS users based on having positive WS vCore hours in the given year. This overstates the usage of WS prior to May 2022 since WS and Linux hours are commingled for reserved instances in these years. These hours have been fully attributed to WS on a conservative basis. Further, prior to March 2021, WS and Windows Client hours are commingled. A small number of customers (<10) are excluded if WS spend exceeds total Azure Consumed Revenues ("ACR") due to an anomaly in the data. The prefix "CY" indicates "Calendar Year" for emphasis.*

21. Following on from the fact that WS customers consume many more services than just WS VMs, the spend on WS for these customers represents a limited share of total spend on Azure, 12% in 2024, which has been declining from 14% in 2023 and 16% in 2022,[21] as shown below.[22]

---

[21] The CMA's estimate of the proportion of total revenues that WS VM revenues account for when looking at WS-using customers is 28% in 2022, as stated in paragraph 6.303 of the CMA PDR. However, the CMA takes an unweighted mean of the proportion of revenues. When taking the revenue-weighted proportion of the same data, this figure falls to 16.3%. Moreover, the CMA uses now-outdated WS spend data to perform the calculations in the PDR, for which Linux and WS hours are commingled in 2022. When using the updated data (which is accurate for the period May-Dec 2022 onwards), this figure falls to 16.2%, as can be seen in Figure 4 below.

[22] Note that this is a revenue-weighted share. We note that although WS spend is likely to be relatively higher on Microsoft by virtue of being its proprietary product, the Azure WS spend share may also understate usage on rival clouds as the revenues include the spend from AHB customers who receive a discount. Therefore, we perform a sensitivity to approximate the WS spend share assuming AHB customers pay the same effective price as non-AHB WS customers. Comparing the WS AHB and non AHB spend to their respective vCore hours, we can approximate the average per unit cost of one vCore hour for AHB and non-AHB usage. Then, assuming AHB costs as much as non-AHB usage, the share of WS spend in 2024 is 14.7%. Alternatively, if we assume that AHB is sold at a 40% discount,

**Figure 4: WS VM share of ACR (Azure spend) for WS UK customers, May-Dec 2022 and 2023-24**

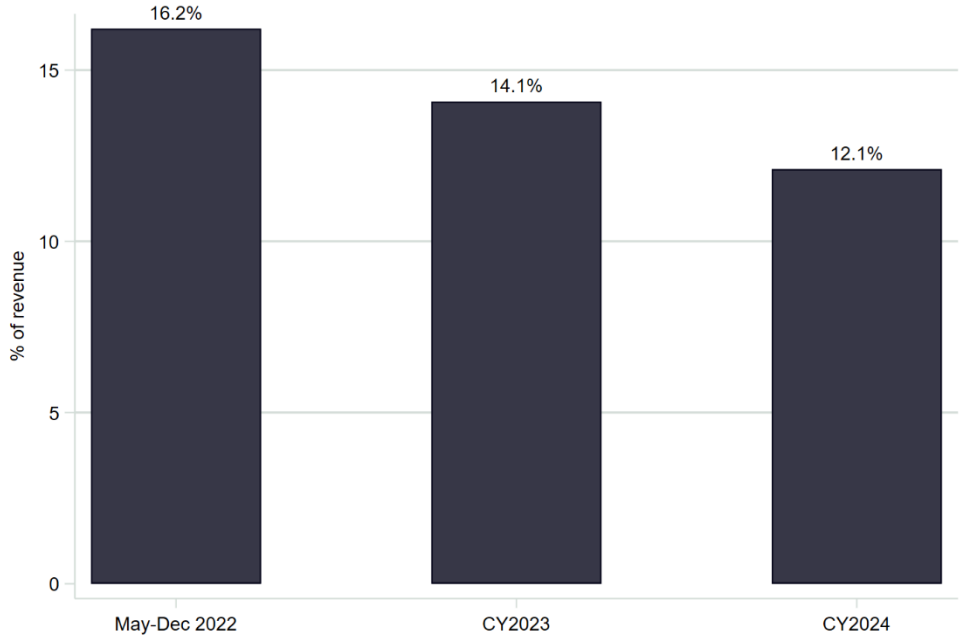


*Source: Keystone analysis of Microsoft data. Notes: Azure Consumed Revenue ("ACR") is the revenue metric tracked by Microsoft for Azure customers. Note: This chart shows the share of total cloud spend on WS for WS users in the UK in the periods May-Dec 2022, 2023, and 2024. Customers are classified as WS users based on having positive WS vCore hours in a given period. The data for May-Dec 2022 includes some customers who have >$1,000 cloud spend in the full year 2022, but over the 8-month period may not meet this threshold. The ACR and WS spend presented for May-Dec 2022 is based on this 8-month period alone. The prefix "CY" indicates "Calendar Year" for emphasis.*

22. The WS spend share is relatively lower for larger customers (who are more likely to multicloud), with the WS spend share being lower than the average for all customers who spend above $5m. This is shown below. This is in direct contrast with the CMA's hypothesis that larger customers who are more likely to multicloud are also more likely to have a higher share of WS in their workloads.

we can calculate the pre-discount cost, and this leads to a WS spend share of 14.5%. We are conservatively assuming a round 40% discount given that Microsoft states that "[w]ith the Azure Hybrid Benefit discount, WS customers can save an average of 36% compared to the leading cloud provider" on their AHB cost calculator webpage: Azure Hybrid Benefit - Hybrid Cost Calculator | Microsoft Azure (accessed 20th February, 2024).

**Figure 5: WS VM share of ACR (Azure spend) for WS UK customers, by revenue category, May-Dec 2022 and 2023-24**

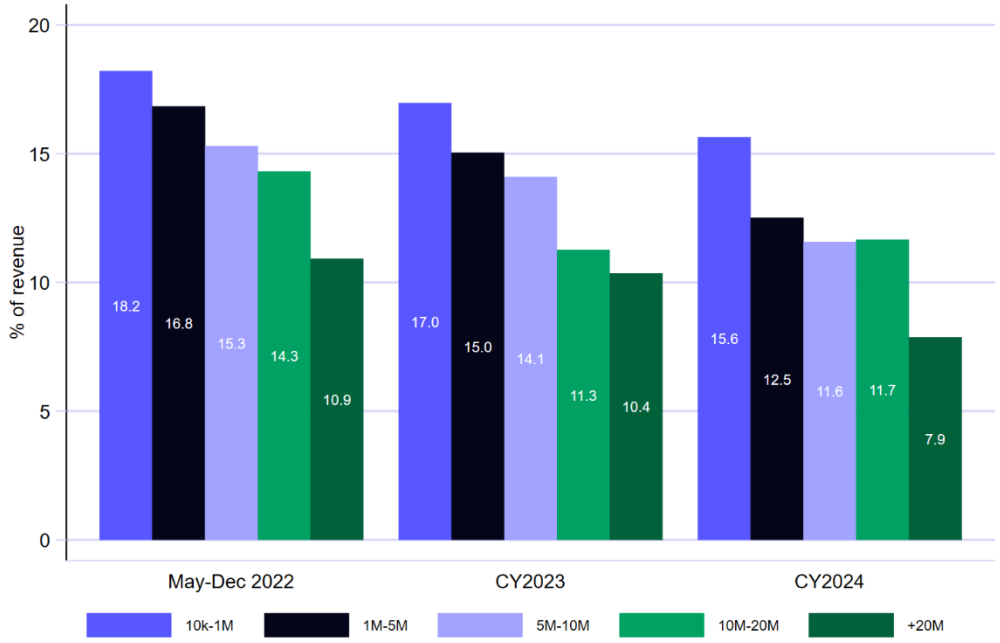


*Source: Keystone analysis of Microsoft data. Notes: Azure Consumed Revenue ("ACR") is the revenue metric tracked by Microsoft for Azure customers. Note: This chart shows the share of total cloud spend on WS for WS users in the UK in the periods May-Dec 2022, 2023, and 2024, by revenue category. Customers are classified as WS users based on having positive WS vCore hours in a given period. Customers who spend less than $10,000 in a given period have been omitted. The data for May-Dec 2022 includes some customers who have >$1,000 cloud spend in the full year 2022, but over the 8-month period may not meet this threshold. The ACR and WS spend presented for May-Dec 2022 is based on this 8-month period alone. The prefix "CY" indicates "Calendar Year" for emphasis.*

23. In response to the CMA's concerns with WS spend share results,[23] we acknowledge that these shares represent a lower bound estimate of the importance of WS to customer decision making and in the cost base of rival providers. Based on the evidence provided above, we believe this lower bound estimate will be accurate for the majority of customers as competition in the cloud is for multiple services which together make a workload, or even wider, a group of workloads making a deal.

24. However, we understand the CMA's approach in the PDR to stress-test this lower bound through exploring a plausible upper bound. In our view, the CMA has, however, chosen a narrow upper bound that is **inconsistent with the business reality of how cloud customers consume cloud services.** Regardless, we address the CMA's concerns for a suitable upper bound in the analysis below.

25. In Section 3.2, we consider multicloud customers, who likely disaggregate their decision making and we show that for these customers, the number of services consumed as well as WS spend share are similar to the figures for the average Azure customer. Further, In Section 3.3, we consider customers who consume WS compute and no Linux (a plausible segment of customers who are reliant on WS IP) and show that these customers also consume multiple services despite having

[23] CMA PDR, para 6.298.

slightly higher than average WS spend share. Together, the analysis in these sections shows that even for these two sets of customers, the WS share of spend is close to the lower bound estimate.

26. In Section 4.0, we consider plausible upper bounds for this analysis, based on a combination of services most commonly used with WS VMs (as no customer consumes VMs alone). We recognize that, given the fact that customers use many services on the cloud, our upper bound is still hypothetical but we present this as a suitable comparison point for the CMA's framework.

27. Finally, we address the CMA's concerns about rival margins in Section 5.0 and present an analysis including SQL Server in Section 6.0.

### 3.2 EVEN FOR MULTICLOUD CUSTOMERS, WS IS A COMPONENT OF A WIDER BUNDLE OF CLOUD SERVICES

28. The evidence shows that even for customers who multicloud across Azure, AWS and GCP, WS is just one component of a wider bundle of services they consume on Azure.

29. As shown in [redacted].

**Table 1: [redacted]**

30. Further, [redacted].[24]

**Table 2: [redacted]**

31. Finally, [redacted].

**Figure 6: [redacted]**

### 3.3 EVEN FOR THE MOST WINDOWS-INTENSIVE WORKLOADS, WS IS A COMPONENT OF A WIDER BUNDLE OF CLOUD SERVICES

32. In this section, we further stress-test the CMA's concerns and show that even customers with plausible WS-centric workloads consume a number of Azure services, despite a larger-than-average WS spend share.[25]

33. We consider three categories of customers who are likely to have workloads with WS as a significant component (evidenced by higher-than-average WS spend share):

    a. **WS compute and no Linux:** these are customers who use WS VMs but do not use Linux VMs (i.e. customers with positive WS vCore hours and zero Linux vCore hours in

[24] There are a few exceptions where multicloud customers spend a higher share of their cloud spend on WS. For these customers, the difference in WS spend share is still not large.

[25] We are presenting this analysis in response to the CMA's concern that there may be other customer groups where WS is a significant share of spend (CMA PDR, para 6.301). We are presenting three plausible categories for which this may be a concern.

a given year). These customers comprise 5% of total WS vCore hours in 2024 in the UK; they had 20% WS spend share in 2024 (compared to the average of 12% in 2024 across all WS UK customers).

b. **WS AHB:** these are customers who use (at least in part) WS AHB (i.e., positive WS AHB vCore hours in a given year), comprising 48% of total WS vCore hours in 2024. These customers had 12% WS spend share in 2024, similar to the average of 12% in 2024 across all WS UK customers.

c. **WS AHB with 100% WS compute:** this is the intersection of the two aforementioned categories, i.e., WS users with zero Linux vCore hours and who also use WS AHB (at least in part) in a given year. This segment is quite small as the usage of these customers comprises less than 1% of total WS vCore hours in 2024.These customers had 19% WS spend share in 2024.

34. A more detailed distribution of the number of services consumed by each customer category is provided in the Annex.

**Table 3: Distribution of services and WS spend share for WS-intensive workloads, 2023-24**

| Category | Year | Average services | Median services | Minimum services | Maximum services | WS spend share (%) |
|---|---|---|---|---|---|---|
| **WS compute only** | 2023 | 12 | 10 | 3 | 39 | 22% |
| | 2024 | 12 | 10 | 3 | 37 | 20% |
| **WS AHB** | 2023 | 23 | 19 | 4 | 102 | 14% |
| | 2024 | 24 | 20 | 4 | 103 | 12% |
| **WS compute only *and* WS AHB** | 2023 | 12 | 11 | 4 | 39 | 22% |
| | 2024 | 12 | 11 | 4 | 34 | 19% |

*Notes: Based on Microsoft UK Azure data, this analysis is restricted to customers with at least $10,000 spend in 2024. The services are defined based on Service Level 2 (SL2) data; Service Level 2 is the level of granularity with which Microsoft records data in its financial systems.*

## 4.0 Input cost shares considering services most commonly consumed with WS VMs

35. We do not dispute that some of the services a customer consumes may be unrelated to WS VMs as argued in the PDR. As such, the average number of services analysed above may be an overestimate for a specific workload that includes WS as an input.[26] In that sense, the CMA is right to use total cloud spend as the *lower* bound for the denominator in the analysis as not all cloud spend may be related to the WS input. However, as we show in this section, any upper bound (in

[26] CMA PDR, para 6.254

the CMA's upper bound/lower bound framework) of the denominator cannot be limited to just WS VMs.

36. Most uses of a WS VM necessitate using other cloud services. In the pricing calculators for AWS, Azure and GCP, the configuration pages for WS VMs include ancillary services required for a VM, namely, storage and networking. Azure's pricing calculator includes "managed disks",[27] "storage transactions" and "bandwidth" as additional services for a VM.[28] AWS' pricing calculator includes "Amazon Elastic Block Store (EBS)"[29], "Detailed Monitoring"[30] and "Data transfer".[31] GCP's pricing calculator includes "Boot disk"[32] and "Local SSD".[33, 34]

37. Each customer consumes a diverse combination of services depending on the workload. The table below shows the share of WS customers who use a given service or combination, both in terms of number of customers and revenue. As shown in the below table, all customers consume storage and networking with WS VMs as otherwise the VM would be of no use. [35] Given the results below, for purposes of addressing the PDR hypothetical scenarios, we consider VMs plus storage and virtual networking as the narrowest plausible combination of services consumed by any WS customer, followed by VMs, storage, virtual network and bandwidth. However, this is still hypothetical as most customers use many more services.

**Table 4: Service combinations by WS UK customers, 2023-24**

| Service(s) in addition to WS VMs | 2023 | | 2024 | |
|---|---|---|---|---|
| | Share of WS customers (%) | Share of WS customer Azure revenue (%) | Share of WS customers (%) | Share of WS customer Azure revenue (%) |
| **Storage, Virtual Network** | 99.30% | 99.94% | 99.73% | 99.97% |
| **Storage, Virtual Network, Bandwidth** | 98.82% | 99.90% | 99.27% | 99.94% |

[27] Azure managed disks are "block-level storage volumes that are managed by Azure and used with Azure Virtual Machines." *See* Overview of Azure Disk Storage - Azure Virtual Machines | Microsoft Learn.

[28] *See* Pricing Calculator | Microsoft Azure.

[29] Amazon Elastic Block Store (Amazon EBS) is "an easy-to-use, scalable, high-performance block-storage service designed for Amazon Elastic Compute Cloud (Amazon EC2)." *See* Cloud Block Storage - Amazon EBS - AWS.

[30] Monitoring of key metrics on EC2 instances. *See* Manage detailed monitoring for your EC2 instances - Amazon Elastic Compute Cloud.

[31] *See* Create estimate: Configure Amazon EC2.

[32] "When you create a Compute Engine instance, you must also create a boot disk for the instance." *See* Create a customized boot disk | Compute Engine Documentation | Google Cloud.

[33] Local SSD is a form of "high performance, low latency, temporary storage". *See* About Local SSD disks | Compute Engine Documentation | Google Cloud.

[34] *See* Google Cloud Pricing Calculator.

[35] We note that there were only 2 active UK customers in 2024 who use WS VMs but have negative storage ACR. The SL2s with negative ACRs are not included in the dataset, hence the customers are recorded as having not used storage for the analysis. This is the case for some observations for the other services as well, meaning that the table may slightly underestimate the share of customers for various service combinations.

| | | | | |
|---|---|---|---|---|
| **Storage, Virtual Network, Bandwidth, Backup** | 82.30% | 96.64% | 83.84% | 97.32% |

*Notes: Based on Microsoft UK Azure data, this analysis is restricted to customers with positive WS vCore hours in a given year and at least $10,000 annual total cloud spend in a given year. The services are defined based on Service Level 2 (SL2) data. Service Level 2 is the level of granularity with which Microsoft records data in its financial systems.*

38. Considering the above common combinations, we present the WS Hyperscaler SPLA input cost as a share of total revenues for the given services.[36] The SPLA input cost share as a share of all Azure revenues for WS customers is presented in the last row – this is the CMA's lower bound. We show that considering the additional services consumed in combination with VMs, the input cost shares are closer to the CMA's lower bound estimates. Although SPLA input cost shares are approximately 70% when one considers an unrealistic bundle, WS VMs alone (which is the CMA's upper bound and on the top row of the below table), this share drops significantly when factoring in the storage and virtual networking resources consumed by these customers.[37]

39. **However, note that these shares are still conservative as they do not consider that WS customers consume additional services.** In 2024, only 45 WS customers with spend above $10,000 consumed the combination of VMs, Storage, Virtual Network, and Bandwidth alone; this represents 0.38% of customers and 0.02% of revenues. Moreover, if we consider larger customers, which the CMA suggests lie closer to the (hypothetical) upper bound, they consume even more services in addition to WS VMs, with at least 95% of customers who spend above $1m[38] consuming Storage, Virtual Network, Bandwidth, Log Analytics, Azure DNS, Azure Monitor, Key Vault, Microsoft Defender for Cloud, SQL Database and Backup in addition to WS VMs, as shown in Table 6. Further, if we consider customers who spend above $20m on Azure, all customers in this bucket used 24 services in common, in addition to WS VMs in 2024.[39] This further shows that there are very few cases, if any, where the input share will be close to the CMA's upper bound or even the upper bounds explored in this note.

**Table 5: SPLA input cost shares of revenues for WS customers, 2023-24**

| Service(s) in addition to WS VMs | 2023 (%) | 2024 (%) |
|---|---|---|
| **WS VMs only** | 68.8% | 70.6% |

[36] We acknowledge, as the CMA does, that costs are a more appropriate base for analysis. However, due to lack of accurate information on the costs for each service combination, we have used revenues, which is in line with the CMA's approach. This means that our numbers can be compared to the upper bound shares in the CMA's analysis. Further, although we approximate costs in our as-efficient competitor test below, we note that this is only an approximation based on Azure-wide COGS, and is not service-specific.

[37] Note that we only include the WS share of storage, virtual network, and bandwidth revenues, estimated as a customer's WS share of total compute (WS vCore hours divided by WS + Linux vCore hours) multiplied by the relevant revenues.

[38] These customers represented 71% of the total WS-using customer UK Azure revenue in 2024.

[39] These services are Azure Data Factory v2, Bandwidth, VPN Gateway, Automation, Logic Apps, Azure Cosmos DB, Cognitive Services, Microsoft Defender for Cloud, Storage, Functions, Virtual Network, Load Balancer, Container Registry, Network Watcher, Backup, Service Bus, Key Vault, Redis Cache, Azure Monitor, Log Analytics, Container Instances, Event Grid, Azure DNS, Azure Database for PostgreSQL. Note that these customers still consumed, on average, 79 services in 2024.

| | | |
|---|---|---|
| **WS VMs, Storage, Virtual Network** | 36.2% | 35.4% |
| **WS VMs, Storage, Virtual Network, Bandwidth** | 35.6% | 34.7% |
| **All cloud services** | 9.6% | 8.4% |

*Notes: Based on analysis of Microsoft data, this analysis is restricted to customers with annual total cloud spend of at least $10,000. The SPLA input cost is calculated based on the hyperscaler cost of $0.0391 per vCore hour applied to total WS vCore hours, and the input cost share is calculated in relation to the pro-rated ACR of the services included in each row. These input shares are weighted input shares.*

**Table 6: Service combinations by WS UK customers who spend above $1m, 2023-24**

| Service(s) in addition to WS VMs | 2023 | | 2024 | |
|---|---|---|---|---|
| | Share of WS customers (%) | Share of WS customer Azure revenue (%) | Share of WS customers (%) | Share of WS customer Azure revenue (%) |
| **Storage, Virtual Network, Bandwidth** | 100.00% | 100.00% | 100.00% | 100.00% |
| **Storage, Virtual Network, Bandwidth, Log Analytics** | 99.61% | 99.89% | 99.51% | 99.61% |
| **Storage, Virtual Network, Bandwidth, Log Analytics, Azure DNS** | 99.03% | 99.65% | 98.86% | 99.42% |
| **Storage, Virtual Network, Bandwidth, Log Analytics, Azure DNS, Azure Monitor** | 98.64% | 99.51% | 98.20% | 99.10% |
| **Storage, Virtual Network, Bandwidth, Log Analytics, Azure DNS, Azure Monitor, Key Vault** | 98.25% | 99.42% | 98.20% | 99.10% |
| **Storage, Virtual Network, Bandwidth, Log Analytics, Azure DNS, Azure Monitor, Key Vault, Microsoft Defender for Cloud** | 96.89% | 98.70% | 97.06% | 98.78% |
| **Storage, Virtual Network, Bandwidth, Log Analytics, Azure DNS, Azure Monitor, Key Vault, Microsoft Defender for Cloud, SQL Database** | 95.92% | 98.21% | 96.08% | 97.31% |
| **Storage, Virtual Network, Bandwidth, Log Analytics, Azure DNS, Azure Monitor, Key Vault, Microsoft Defender for Cloud, SQL Database, Backup** | 94.56% | 97.58% | 95.42% | 97.11% |

*Notes: Based on Microsoft UK Azure data, this table is restricted to customers with positive WS vCore hours in a given year and at least $1m annual total cloud spend in a given year. The top services are defined based on 2023 Service Level 2 (SL2) data. Service Level 2 is the level of granularity with which Microsoft records data in its financial systems.*

40. Moreover, even if the input cost share for WS VMs as the downstream product is high, cloud service providers may not have an incentive to pass it on unless the input cost is high enough for a workload, i.e., the margins cloud providers make on the workload, rather than a component, are relevant.[40] We turn to this next.

## 5.0 An As-Efficient competitor (AEC) can compete with Microsoft

41. In this section, we present the results applying an 'as-efficient competitor' test. The margins are positive even when including just the services that are almost always consumed by customers deploying WS VMs, although customers tend to consume many more services on average. We present the summary results in this section and the more detailed tables (as well as a detailed methodology) in the Annex.

42. Further, we find that the margins are positive even for customers with the most WS-centric workloads.[41]

43. Even for a narrow hypothetical workload, an as-efficient competitor can compete with Microsoft. When including just storage, network and bandwidth revenues allocated to WS VMs,[42] the margins are positive and above 25% if Microsoft had to bear the Hyperscale SPLA costs. This is still conservative as it does not consider that WS customers consume additional services - WS customers consumed on average ~14 services in 2024 (as shown in Figure 2).

**Table 7: As-efficient competitor results for WS UK customers, 2024**

| | | [A] WS VMs, storage and virtual network | [B] WS VMs, storage, virtual network and bandwidth |
|---|---|---|---|
| *Revenues* | WS VMs | $500M | $500M |
| *Revenues* | Storage | $463M | $463M |
| *Revenues* | Virtual Network | $33M | $33M |
| *Revenues* | Bandwidth | - | $20M |
| *Hypothetical Costs – IP* | *Azure WS vCore hours* | *9,036M* | *9,036M* |
| ***Margin*** | **Margin over licensing costs and other COGS** | **$268M** | **$280M** |
| ***Margin*** | **Margin over licensing costs and other COGS (%)** | **27%** | **28%** |

*Source: Keystone analysis of Azure data. Notes: Based on Microsoft UK Azure data, this table is restricted to customers with positive WS vCore hours in a given year and at least $10,000 annual total cloud spend in a given year. Additional*

[40] As previously stated, the margins cloud providers make on related services within a workload are also relevant. The CMA estimates that [redacted] (CMA PDR, Appendix Q, para Q.86), indicating that cloud providers are able to make significant margins on WS-related services in the workload.

[41] All analysis in this section is restricted to customers with >$10,000 spend in 2024 following the CMA's methodology (CMA PDR, para 6.280).

[42] The revenues for storage, virtual network and bandwidth are pro-rated based on the WS share of total compute (WS + Linux) vCore hours. The pro-rating exercise is conducted for each customer individually and then aggregated. This is true for all the AEC analyses in this note.

*Service Level 2 (SL2) service revenues are prorated based on a customer's WS vCore share of total compute vCore hours. Service Level 2 is the level of granularity with which Microsoft records data in its financial systems.*

### 5.1 EVEN FOR WS-CENTRIC WORKLOADS, AN AS-EFFICIENT COMPETITOR CAN COMPETE WITH MICROSOFT

44. In this section, we look at the as-efficient competitor test results for the same narrow subsets of plausible WS-centric workloads described in Section 3.3 above:[43]

    a. **WS users with 100% compute** – customers with positive WS vCore hours and zero Linux hours, comprising 5% of WS vCore hours in 2024.

    b. **WS AHB** - customers who use WS AHB (i.e., positive WS AHB vCore hours in a given year), comprising 48% of WS vCore hours in 2024.[44]

    c. **WS users with 100% compute (AHB)** – customers with positive WS AHB vCore hours and zero Linux hours, comprising <1% of WS vCore hours in 2024.

45. The table below provides the results for the as-efficient competitor test when considering the first group - WS-users who do not use Linux (100% WS compute). For this subset of customers, the margins are positive and above 30%, even if Microsoft had to pay the SPLA costs for WS IP.

**Table 8: As-efficient competitor results for WS UK customers with 100% WS Compute, 2024**

| | | [A] WS, storage and virtual network | [B] WS, storage, virtual network and bandwidth |
|---|---|---|---|
| *Revenues* | WS | $33M | $33M |
| *Revenues* | Storage | $32M | $32M |
| *Revenues* | Virtual Network | $1M | $1M |
| *Revenues* | Bandwidth | - | $2M |
| *Hypothetical Costs – IP* | Azure WS vCore hours | *477M* | *477M* |
| ***Margin*** | **Margin over licensing costs and other COGS** | **$23M** | **$24M** |
| ***Margin*** | **Margin over licensing costs and other COGS (%)** | **34%** | **35%** |

*Source: Keystone analysis of Azure data. Notes: Based on Microsoft UK Azure data, this table is restricted to customers with positive WS vCore hours and zero Linux vCore hours in a given year and at least $10,000 annual total cloud spend in a given year. Additional Service Level 2 (SL2) service revenues are prorated based on a customer's WS vCore share of total compute vCore hours. Service Level 2 is the level of granularity with which Microsoft records data in its financial systems.*

[43] *See* the Annex for additional sensitivities considering WS non-AHB customers.

[44] AHB customers are defined based on positive AHB vCore hours and their AHB revenues are taken into consideration for the AEC test, regardless of any non-AHB usage/spend by the same customers. Similar caveats apply for non-AHB users. Hence, the sum of customer shares will exceed 100%, as customers can consume both.

46. For the second set of customers described above, i.e. WS AHB customers, the margins are positive and at least 20%.

**Table 9: As-efficient competitor results for WS AHB UK customers, 2024**

| | | [A] WS AHB, storage and virtual network | [B] WS AHB, storage, virtual network and bandwidth |
|---|---|---|---|
| *Revenues* | WS AHB | $174M | $174M |
| *Revenues* | Storage | $215M | $215M |
| *Revenues* | Virtual Network | $17M | $17M |
| *Revenues* | Bandwidth | - | $8M |
| *Hypothetical Costs – IP* | *Azure WS AHB vCore hours* | *4,343M* | *4,343M* |
| ***Margin*** | **Margin over licensing costs and other COGS** | **$81M** | **$86M** |
| ***Margin*** | **Margin over licensing costs and other COGS (%)** | **20%** | **21%** |

*Source: Keystone analysis of Azure data. Notes: Based on Microsoft UK Azure data, this table is restricted to customers with positive WS AHB vCore hours in a given year and at least $10,000 annual total cloud spend in a given year. Additional Service Level 2 (SL2) service revenues are prorated based on a customer's WS vCore share of total compute vCore hours. Service Level 2 is the level of granularity with which Microsoft records data in its financial systems.*

47. Finally, we also consider the intersection of these two segments i.e., WS AHB using customers with 100% of their compute on WS. As shown in the table below, margins are positive and above 25% for this small customer segment.

**Table 10: As-efficient competitor results for WS AHB UK customers with 100% WS Compute, 2024**

| | | [A] WS AHB, storage and virtual network | [B] WS AHB, storage, virtual network and bandwidth |
|---|---|---|---|
| *Revenues* | WS AHB | $4M | $4M |
| *Revenues* | Storage | $5M | $5M |
| *Revenues* | Virtual Network | $0.2M | $0.2M |
| *Revenues* | Bandwidth | - | $0.3M |
| *Hypothetical Costs – IP* | Azure WS AHB vCore hours | *80M* | *80M* |
| ***Margin*** | **Margin over licensing costs and other COGS** | **$2M** | **$3M** |

| | | | |
|---|---|---|---|
| ***Margin*** | **Margin over licensing costs and other COGS (%)** | **27%** | **28%** |

*Source: Keystone analysis of Azure data. Notes: Based on Microsoft UK Azure data, this table is restricted to customers with positive WS AHB vCore hours and zero Linux vCore hours in a given year and at least $10,000 annual total cloud spend in a given year. Additional Service Level 2 (SL2) service revenues are prorated based on a customer's WS vCore share of total compute vCore hours. Service Level 2 is the level of granularity with which Microsoft records data in its financial systems.*

### 5.2 RESTRICTING THE AS-EFFICIENT COMPETITOR TEST TO NON-AHB WS USAGE YIELDS POSITIVE MARGINS

48. We have also looked at the as-efficient competitor test results for WS non-AHB customers, and specifically the following two groups:

    a. **WS non-AHB** – customers who use WS without AHB (i.e., positive WS non-AHB vCore hours in a given year), comprising 52% of WS vCore hours in 2024.

    b. **WS users with 100% compute (non-AHB)** – customers with positive WS non-AHB vCore hours and zero Linux hours, comprising 4% of WS vCore hours in 2024.

49. For non-AHB WS customers, as shown in the table below, the margins are above 30% with the inclusion of storage and other services.

**Table 11: As-efficient competitor results for WS non-AHB UK customers, 2024**

| | | [A] WS non-AHB, storage and virtual network | [B] WS non-AHB, storage, virtual network and bandwidth |
|---|---|---|---|
| *Revenues* | WS non-AHB | $327M | $327M |
| *Revenues* | Storage | $248M | $248M |
| *Revenues* | Virtual Network | $17M | $17M |
| *Revenues* | Bandwidth | - | $12M |
| *Hypothetical Costs - IP* | Azure WS non-AHB vCore hours | *4,693M* | *4,693M* |
| ***Margin*** | **Margin over licensing costs and other COGS** | **$187M** | **$194M** |
| ***Margin*** | **Margin over licensing costs and other COGS (%)** | **32%** | **32%** |

*Source: Keystone analysis of Azure data. Notes: Based on Microsoft UK Azure data, this table is restricted to customers with positive non-AHB WS vCore hours in a given year and at least $10,000 annual total cloud spend in a given year. Additional Service Level 2 (SL2) service revenues are prorated based on a customer's WS vCore share of total compute vCore hours. Service Level 2 is the level of granularity with which Microsoft records data in its financial systems.*

50. For the second group i.e. non-AHB customers with WS 100% compute, the margins are above 35% with the inclusion of storage and other services.

**Table 12: As-efficient competitor results for WS non-AHB UK customers with 100% WS Compute, 2024**

| | | [A] WS non-AHB, storage and virtual network | [B] WS non-AHB, storage, virtual network and bandwidth |
|---|---|---|---|
| *Revenues* | WS non-AHB | $29M | $29M |
| *Revenues* | Storage | $27M | $27M |
| *Revenues* | Virtual Network | $1M | $1M |
| *Revenues* | Bandwidth | - | $2M |
| *Hypothetical Costs - IP* | Azure WS non-AHB vCore hours | *397M* | *397M* |
| ***Margin*** | **Margin over licensing costs and other COGS** | **$20M** | **$21M** |
| ***Margin*** | **Margin over licensing costs and other COGS (%)** | **35%** | **36%** |

*Source: Keystone analysis of Azure data. Notes: Based on Microsoft UK Azure data, this table is restricted to customers with positive non-AHB WS vCore hours and zero Linux vCore hours in a given year and at least $10,000 annual total cloud spend in a given year. Additional Service Level 2 (SL2) service revenues are prorated based on a customer's WS vCore share of total compute vCore hours. Service Level 2 is the level of granularity with which Microsoft records data in its financial systems.*

## 6.0 No foreclosure concerns arise for SQL customers

51. As the CMA notes, SQL Server is an RDBMS software product that can be run as an IaaS (i.e., on a VM) or PaaS (i.e., as a managed service) product.[45]

52. SQL IaaS makes up the majority of SQL Server usage on Azure; in 2024, SQL PaaS vCore hours made up 30% of total SQL Server vCore hours (IaaS and PaaS) consumed by active UK Azure customers. Although the CMA notes that SQL PaaS products are one of the most popular cloud services, with SQL Database representing 8% of total UK Azure revenues in 2022,[46] we note that this is [redacted] AWS' RDS share of total UK revenues, [redacted].[47]

53. As PaaS is a minority of SQL usage on Azure and the PDR only addresses IaaS, the rest of this analysis focuses on SQL IaaS, specifically on SQL IaaS PAYG usage as existing customers can BYOL to AWS and Google.[48] In 2024, SQL IaaS PAYG accounted for the minority of usage of SQL IaaS on Azure (compared to AHB), making up 40% of total SQL IaaS vCore hours by active UK Azure customers.

[45] CMA PDR, para 6.85

[46] CMA PDR, para 6.274

[47] Keystone analysis of [redacted].

[48] CMA PDR, footnote 1740.

### 6.1 THE AEC TEST ON WS AND SQL PAYG CUSTOMERS DEMONSTRATES A LACK OF FORECLOSURE CONCERNS

54. In this section, we show that an as-efficient competitor can compete for Azure SQL IaaS PAYG customers. Similarly to Section 5.0, we show that an as-efficient competitor can compete for WS and SQL Server IaaS PAYG workloads despite paying the SPLA cost.

55. As the CMA notes, most SQL Server workloads are run on WS.[49] In 2024, 99.9% of active UK Azure customers who used SQL Server IaaS PAYG also consumed WS. As such, we extend the as-efficient competitor test conducted for WS users to also include revenues and costs from SQL IP.[50] Specifically, we add SQL IP revenues to the total revenues and SQL SPLA cost to the total costs considered.

56. The CMA's analysis of SQL price differences and input cost shares considers hypothetical splits of SQL server usage by Enterprise and Standard editions (25:75, 50:50 and 40:60). We find that these likely overestimate the usage of Enterprise edition, and do not account for the Web edition, which is comparatively lower in SPLA price. As such, we update the CMA's analysis presenting 4 scenarios of hypothetical SPLA costs to reflect the relative usage between Standard, Enterprise, and Web editions on Azure, AWS and GCP, as well as the variation in SQL SPLA prices between AWS and GCP.[51] We consider the following scenarios:

    a. Relative SQL usage on AWS using AWS SPLA prices.

    b. Relative SQL usage on GCP using GCP SPLA prices.

    c. Relative SQL usage on Microsoft using AWS SPLA prices.

    d. Relative SQL usage on Microsoft using GCP SPLA prices.

57. While previous CMA requests have been for all SQL Server vCore hours split by edition[52], we rely on just SQL PAYG usage data by edition.[53] The methodology for this process is detailed in the Annex.

58. The estimated edition usage ratios are presented in the below table, alongside AWS and GCP ratios. The table shows that compared to the CMA's assumptions, customers spend relatively more on Standard and Web editions across [redacted], which are lower in price than the Enterprise edition.

---

[49] CMA PDR, para 6.272

[50] The analysis is restricted to customers who use **both** WS and SQL PAYG.

[51] AWS and GCP's SPLA costs per SQL edition are i) Enterprise: $0.352 per vCore hour and $0.388 per vCore hour for AWS and GCP respectively ii) Standard: $0.099 per vCore hour and $0.101 per vCore hour for AWS and GCP respectively iii) Web: $0.006 per vCore hour for both AWS and GCP. AWS' SPLA contract notes a HSPLA price for SQL Server Enterprise edition only. General SPLA prices have been used for AWS for SQL Server Standard and Web edition. GCP's SPLA contract notes a HSPLA price for SQL Server Standard and Enterprise editions. General SPLA prices have been used for GCP for SQL Server Web edition.

[52] For example, question 2 of the CMA's follow-up email request March 2024.

[53] Azure's relative usage data has not previously been provided to the CMA as this is not recorded in the ordinary course of business; however, for the purposes of this exercise, Microsoft has estimated this by conducting a normalisation exercise of VM hours by number of cores (based on Service Level 5 (SL5) data, with SL5 being a level of granularity in Microsoft's financial data).

**Table 13: Relative usage of each edition of SQL Server as part of license-included services on Azure, AWS and GCP across customer revenue brackets, 2022-24**

| Revenue bracket ($) | Azure (2023) | | | Azure (2024) | | | AWS (2022) | | | GCP (2022) | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | Enterprise % | Standard % | Web % | Enterprise % | Standard % | Web % | Enterprise % | Standard % | Web % | Enterprise % | Standard % | Web % |
| 10k-1M | 10.3% | 71.6% | 18.2% | 9.4% | 73.6% | 17.0% | | | | | | |
| 1M-5M | 28.0% | 64.9% | 7.1% | 26.2% | 65.3% | 8.5% | | | | | | |
| 5M-10M | 52.3% | 46.5% | 1.2% | 43.5% | 55.0% | 1.5% | | | | | | |
| 10M-20M | 23.5% | 65.1% | 11.4% | 38.3% | 52.4% | 9.3% | | | | | | |
| +20M | 42.8% | 55.7% | 1.5% | 48.3% | 50.1% | 1.6% | | | | | | |

*Source: Keystone analysis of Azure data. [redacted]*

59. The as-efficient competitor test results below are presented in ranges, depending on the SQL SPLA cost scenario. Detailed results are presented in the Annex.

60. Similarly to the as-efficient competitor tests in Section 5.0 above, WS-prorated storage, virtual network and bandwidth revenues are included in the test, reflecting the fact that customers do not consume VMs alone, even with SQL Server IaaS.[54]

61. As seen in the table below, the as-efficient test results in positive margins (around 20%) for WS and SQL using customers including storage and other commonly consumed services.

**Table 14: As-efficient competitor results for WS and SQL UK customers, 2024**

| | | [A] WS, SQL, storage and virtual network | [B] WS, SQL, storage, virtual network and bandwidth |
|---|---|---|---|
| *Revenues* | WS | $381M | $381M |
| *Revenues* | SQL PAYG | $69M | $69M |
| *Revenues* | Storage | $346M | $346M |
| *Revenues* | Virtual Network | $25M | $25M |
| *Revenues* | Bandwidth | - | $13M |
| *Hypothetical Costs – IP* | *Azure WS vCore hours* | *6,958M* | *6,958M* |
| *Hypothetical Costs – IP* | *Azure SQL PAYG hours* | *645M* | *645M* |
| ***Margin*** | **Margin over licensing costs and other COGS** | **$159M - $187M** | **$167M - $195M** |
| ***Margin*** | **Margin over licensing costs and other COGS (%)** | **19% - 23%** | **20% - 23%** |

*Source: Keystone analysis of Azure data. Notes: Based on Microsoft UK Azure data, this table is restricted to customers with positive WS vCore hours in a given year, positive SQL Server IaaS PAYG vCore hours in a given year and at least $10,000 annual total cloud spend in a given year. Additional Service Level 2 (SL2) service revenues are prorated based on a*

[54] To see this, note that even SQL PaaS offerings include storage as part of the managed service package.

*customer's WS vCore share of total compute vCore hours. Service Level 2 is the level of granularity with which Microsoft records data in its financial systems.*

62. In addition, we conduct a similar exercise for the same categories of customers as the as-efficient competitor test for WS customers. Namely, we present results for WS users with 100% compute who also use SQL PAYG, WS AHB and SQL PAYG users, and WS users with 100% compute (AHB) who also use SQL PAYG.

63. The table below presents margins for WS users with 100% compute who also use SQL PAYG. This shows that margins are positive and substantial (>20%) with the inclusion of storage, virtual network and bandwidth.

**Table 15: As-efficient competitor results for WS (100% compute) and SQL UK customers, 2024**

| | | [A] WS, SQL, storage and virtual network | [B] WS, SQL, storage, virtual network and bandwidth |
|---|---|---|---|
| *Revenues* | WS | $14M | $14M |
| *Revenues* | SQL PAYG | $6M | $6M |
| *Revenues* | Storage | $11M | $11M |
| *Revenues* | Virtual Network | $0.4M | $0.4M |
| *Revenues* | Bandwidth | - | $0.5M |
| *Hypothetical Costs – IP* | Azure WS vCore hours | *190M* | *190M* |
| *Hypothetical Costs – IP* | Azure SQL PAYG hours | *65M* | *65M* |
| ***Margin*** | **Margin over licensing costs and other COGS** | **$7M - $8M** | **$7M - $9M** |
| ***Margin*** | **Margin over licensing costs and other COGS (%)** | **23% - 26%** | **23% - 27%** |

*Source: Keystone analysis of Azure data. Notes: Based on Microsoft UK Azure data, this table is restricted to customers with positive WS vCore hours and zero Linux vCore hours in a given year, positive SQL Server IaaS PAYG vCore hours in a given year and at least $10,000 annual total cloud spend in a given year. Additional Service Level (SL2) service revenues are prorated based on a customer's WS vCore share of total compute vCore hours. Service Level 2 is the level of granularity with which Microsoft records data in its financial systems.*

64. For WS AHB customers, the inclusion of SQL PAYG revenues and costs still shows positive margins, as shown in the table below.

**Table 16: As-efficient competitor results for WS AHB and SQL UK customers, 2024**

| | | [A] WS AHB, SQL, storage and virtual network | [B] WS AHB, SQL, storage, virtual network and bandwidth |
|---|---|---|---|
| *Revenues* | WS AHB | $139M | $139M |
| *Revenues* | SQL PAYG | $51M | $51M |
| *Revenues* | Storage | $169M | $169M |

| | | | |
|---|---|---|---|
| *Revenues* | Virtual Network | $13M | $13M |
| *Revenues* | Bandwidth | - | $6M |
| *Hypothetical Costs – IP* | *Azure WS AHB vCore hours* | *3,492M* | *3,492M* |
| *Hypothetical Costs – IP* | *Azure SQL PAYG hours* | *475M* | *475M* |
| ***Margin*** | **Margin over licensing costs and other COGS** | **$27M - $52M** | **$31M - $56M** |
| ***Margin*** | **Margin over licensing costs and other COGS (%)** | **7% - 14%** | **8% - 15%** |

*Source: Keystone analysis of Azure data. Notes: Based on Microsoft UK Azure data, this table is restricted to customers with positive WS AHB vCore hours in a given year, positive SQL Server IaaS PAYG vCore hours in a given year and at least $10,000 annual total cloud spend in a given year. Additional Service Level 2 (SL2) service revenues are prorated based on a customer's WS vCore share of total compute vCore hours. Service Level 2 is the level of granularity with which Microsoft records data in its financial systems.*

65. Looking at the intersection of these two categories i.e., WS AHB using customers with 100% of their compute used on WS, the margins are positive and substantial with the inclusion of storage, virtual network and bandwidth.

**Table 17: As-efficient competitor results for WS AHB and SQL UK customers (WS 100% Compute), 2024**

| | | [A] WS AHB, SQL, storage and virtual network | [B] WS AHB, SQL, storage, virtual network and bandwidth |
|---|---|---|---|
| *Revenues* | WS AHB | $1M | $1M |
| *Revenues* | SQL PAYG | $1M | $1M |
| *Revenues* | Storage | $1M | $1M |
| *Revenues* | Virtual Network | $0.04M | $0.04M |
| *Revenues* | Bandwidth | - | $0.1M |
| *Hypothetical Costs – IP* | *Azure WS AHB vCore hours* | *21M* | *21M* |
| *Hypothetical Costs – IP* | *Azure SQL PAYG hours* | *13M* | *13M* |
| ***Margin*** | **Margin over licensing costs and other COGS** | **$0M - $1M** | **$0M - $1M** |
| ***Margin*** | **Margin over licensing costs and other COGS (%)** | **9% - 15%** | **10% - 16%** |

*Source: Keystone analysis of Azure data. Notes: Based on Microsoft UK Azure data, this table is restricted to customers with positive WS AHB vCore hours and zero Linux vCore hours in a given year, positive SQL Server IaaS PAYG vCore hours in a given year and at least $10,000 annual total cloud spend in a given year. Additional Service Level 2 (SL2) service revenues are prorated based on a customer's WS vCore share of total compute vCore hours. Service Level 2 is the level of granularity with which Microsoft records data in its financial systems.*

66. Further, similar to the AEC analysis presented in Section 5.2, we can extend the analysis of WS non-AHB users to include SQL PAYG revenues and costs. For WS non-AHB users, the table below

shows that margins are positive (at least 19%) with the inclusion of SQL and storage, virtual network and bandwidth.

**Table 18: As-efficient competitor results for WS non-AHB and SQL UK customers, 2024**

| | | [A] WS non-AHB, SQL, storage and virtual network | [B] WS non-AHB, SQL, storage, virtual network and bandwidth |
|---|---|---|---|
| *Revenues* | WS non-AHB | $242M | $242M |
| *Revenues* | SQL PAYG | $68M | $68M |
| *Revenues* | Storage | $177M | $177M |
| *Revenues* | Virtual Network | $11M | $11M |
| *Revenues* | Bandwidth | - | $7M |
| *Hypothetical Costs – IP* | Azure WS non-AHB vCore hours | *3,466M* | *3,466M* |
| *Hypothetical Costs – IP* | Azure SQL PAYG hours | *641M* | *641M* |
| ***Margin*** | **Margin over licensing costs and other COGS** | **$97M - $125M** | **$101M - $129M** |
| ***Margin*** | **Margin over licensing costs and other COGS (%)** | **19% - 25%** | **20% - 25%** |

*Source: Keystone analysis of Azure data. Notes: Based on Microsoft UK Azure data, this table is restricted to customers with positive non-AHB WS vCore hours in a given year, positive SQL Server IaaS PAYG vCore hours in a given year and at least $10,000 annual total cloud spend in a given year. Additional Service Level 2 (SL2) service revenues are prorated based on a customer's WS vCore share of total compute vCore hours. Service Level 2 is the level of granularity with which Microsoft records data in its financial systems.*

67. We also conduct this exercise restricted to customers who consume WS but not Linux. For this category of customers, the margins continue to be positive.

**Table 19: As-efficient competitor results for WS non-AHB and SQL UK customers (WS 100% Compute), 2024**

| | | [A] WS non-AHB, SQL, storage and virtual network | [B] WS non-AHB, SQL, storage, virtual network and bandwidth |
|---|---|---|---|
| *Revenues* | WS non-AHB | $13M | $13M |
| *Revenues* | SQL PAYG | $6M | $6M |
| *Revenues* | Storage | $10M | $10M |
| *Revenues* | Virtual Network | $0.4M | $0.4M |
| *Revenues* | Bandwidth | - | $0.5M |
| *Hypothetical Costs – IP* | *Azure WS non-AHB vCore hours* | *169M* | *169M* |

| | | | |
|---|---|---|---|
| *Hypothetical Costs – IP* | *Azure SQL PAYG hours* | *63M* | *63M* |
| ***Margin*** | **Margin over licensing costs and other COGS** | **$7M - $8M** | **$7M - $8M** |
| ***Margin*** | **Margin over licensing costs and other COGS (%)** | **23% - 27%** | **24% - 27%** |

*Source: Keystone analysis of Azure data. Notes: Based on Microsoft UK Azure data, this table is restricted to customers with positive non-AHB WS vCore hours and zero Linux vCore hours in a given year, positive SQL Server IaaS PAYG vCore hours in a given year and at least $10,000 annual total cloud spend in a given year. Additional Service Level 2 (SL2) service revenues are prorated based on a customer's WS vCore share of total compute vCore hours. Service Level 2 is the level of granularity with which Microsoft records data in its financial systems.*

# Annex

# Annex 1 – WS distribution of services

## WS COMPUTE ONLY[55]

68. As discussed above, WS-using customers consume a variety of services, and this holds true even for customers with 100% compute used on WS, with an average of around 12 services.[56]

69. Figure 7 below shows the distribution of the number of services consumed in 2024 for UK Azure customers with 100% WS compute (i.e., no Linux) and who spent more than $10,000, who made up 4% of total UK cloud spend in 2024. It shows that these customers use multiple services: on average, about 12 services (10 for the median).

**Figure 7: Distribution of the number of services for WS 100% compute (2024)**

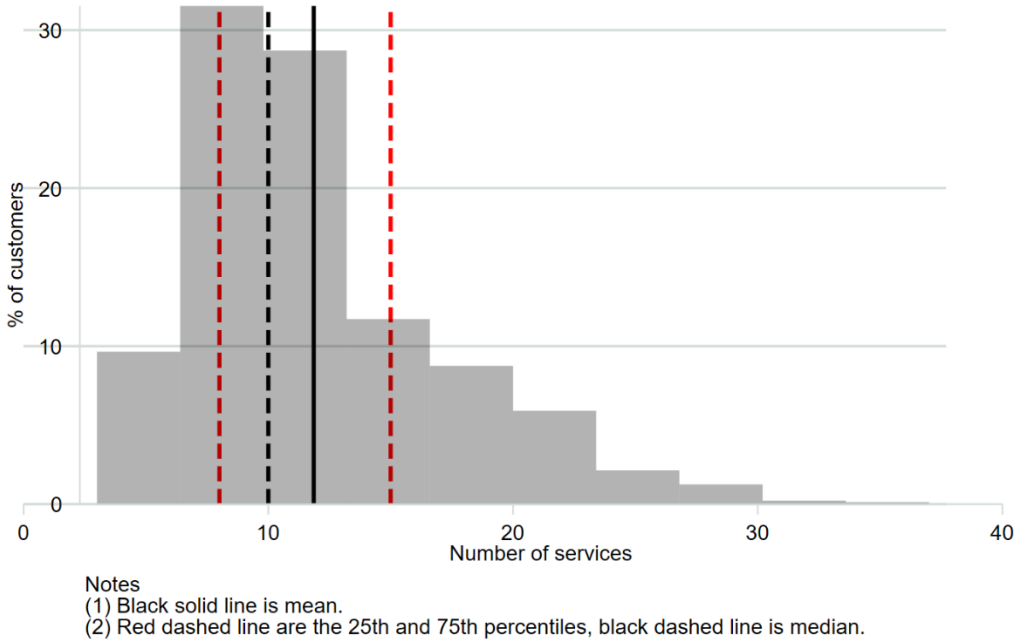


*Notes: Based on Microsoft UK Azure data, this chart plots the distribution of services for customers with positive WS vCore hours and zero Linux hours in 2024 (100% of compute attributed to WS) and above $10,000 spend. The services are defined based on SL2 data; Service Level 2 (SL2) is the level of granularity with which Microsoft records data in its financial systems.*

70. Looking at the services most commonly consumed by these customers, almost all of them consumed virtual network and bandwidth in addition to storage.

**Table 20: Service combinations by WS UK customers (WS compute only), 2023-2024**

| Service(s) in addition to WS VMs | 2023 | 2024 |
|---|---|---|

[55] All analyses in this section of the Annex are restricted to customers with >$10,000 spend.

[56] These are customers with positive WS vCore hours and zero Linux vCore hours, meaning that all of their compute is WS.

| | Share of WS customers (%) | Share of WS customer Azure revenue (%) | Share of WS customers (%) | Share of WS customer Azure revenue (%) |
|---|---|---|---|---|
| **Storage, Virtual Network** | 98.45% | 98.98% | 99.45% | 99.40% |
| **Storage, Virtual Network, Bandwidth** | 97.46% | 98.37% | 98.48% | 98.88% |
| **Storage, Virtual Network, Bandwidth, Backup** | 75.57% | 78.61% | 78.12% | 80.74% |

*Notes: Based on Microsoft UK Azure data, this analysis is restricted to customers with positive WS vCore hours and zero Linux hours in 2024 (100% of compute attributed to WS) as well as above $10,000 spend. The services are defined based on Service Level 2 (SL2) data; Service Level 2 (SL2) is the level of granularity with which Microsoft records data in its financial systems.*

**WS AHB USAGE ONLY**

71. Further, the distribution of services for WS AHB users shows that, in 2024, customers consumed on average 24 services. These customers make up 81% of total cloud spend in 2024.

**Figure 8: Distribution of the number of services for WS AHB users (2024)**

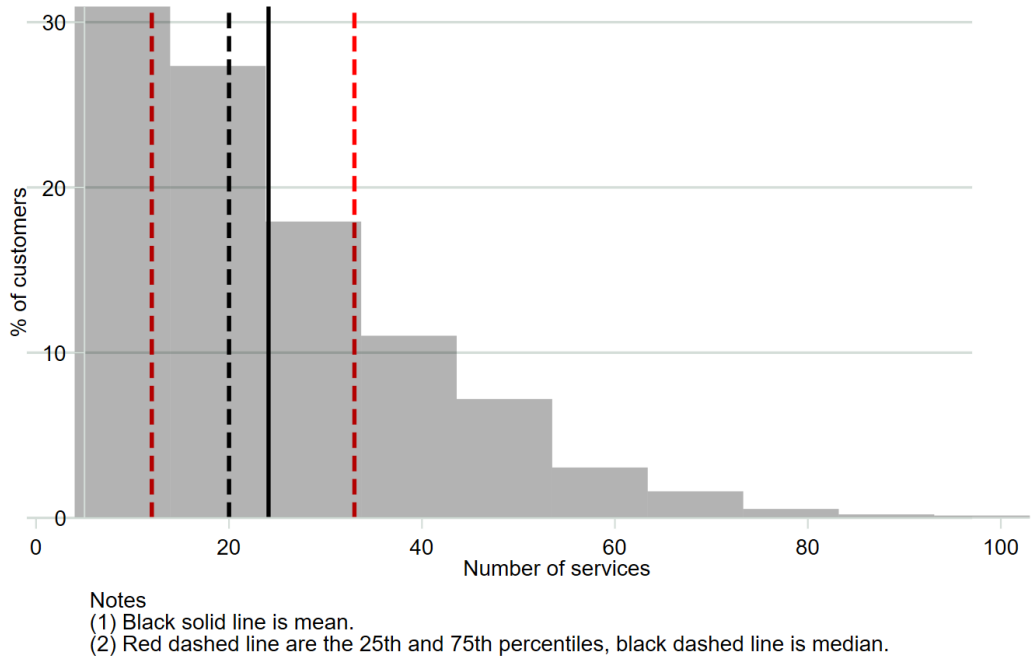


*Notes: Based on Microsoft UK Azure data, this chart plots the distribution of services for customers using WS AHB, i.e., customers with positive WS AHB vCore hours in 2024, and above $10,000 spend. The services are defined based on SL2 data; Service Level 2 (SL2) is the level of granularity with which Microsoft records data in its financial systems.*

72. In 2024, most of these customers also consumed storage and virtual network and over 99% of their ACR came from customers who consumed storage, virtual network, bandwidth and backup.

**Table 21: Service combinations by WS AHB UK customers, 2023-2024**

| Service(s) in addition to WS VMs | 2023 | | 2024 | |
|---|---|---|---|---|
| | Share of WS customers (%) | Share of WS customer Azure revenue (%) | Share of WS customers (%) | Share of WS customer Azure revenue (%) |
| **Storage, Virtual Network** | 99.63% | 99.99% | 99.93% | 99.999% |
| **Storage, Virtual Network, Bandwidth** | 99.32% | 99.98% | 99.69% | 99.99 % |
| **Storage, Virtual Network, Bandwidth, Backup** | 90.05% | 98.93% | 91.03% | 99.21% |

*Notes: Based on Microsoft UK Azure data, this analysis is restricted to customers with positive WS AHB vCore hours and above $10,000 spend. The services are defined based on Service Level 2 (SL2) data; Service Level 2 is the level of granularity with which Microsoft records data in its financial systems.*

## WS AHB USAGE AND WS COMPUTE ONLY

73. Similarly, the distribution of the number of services for AHB users with 100% WS compute shows a lower average number of services than for AHB users but still shows that these customers use a range of services. These customers made up 1% of total cloud spend in 2024.

**Figure 9: Distribution of the number of services for WS AHB users with 100% WS Compute (2024)**

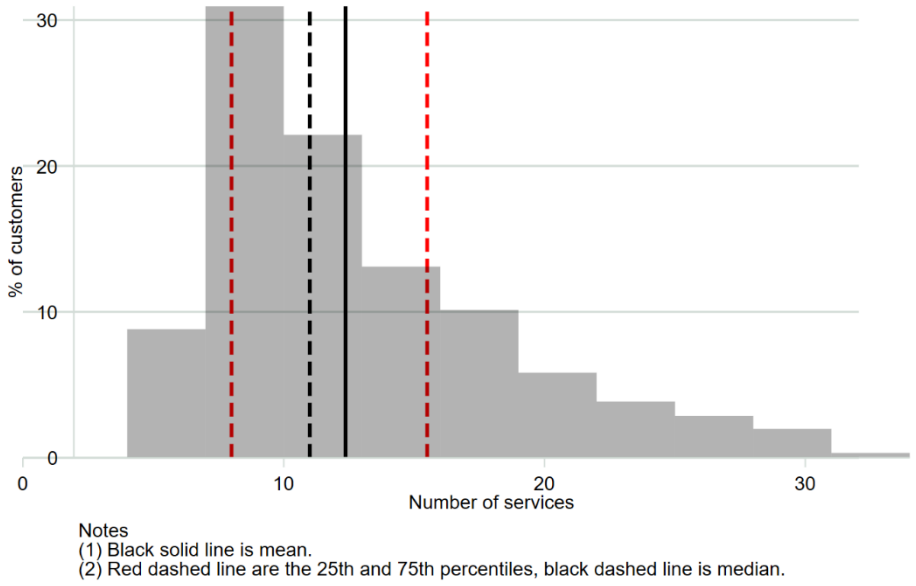


*Notes: Based on Microsoft UK Azure data, this chart plots the distribution of services for customers using WS AHB, i.e., customers with positive WS AHB vCore hours and zero Linux vCore hours in 2024, and above $10,000 spend. The services are defined based on Service Level 2 (SL2) data; Service Level 2 is the level of granularity with which Microsoft records data in its financial systems.*

74. In addition, a majority of the ACR in 2024 for this category came from customers who consumed storage, virtual network, bandwidth and backup.

**Table 22: Service combinations by WS AHB UK customers with 100% WS Compute, 2023-2024**

| Service(s) in addition to WS VMs | 2023 | | 2024 | |
|---|---|---|---|---|
| | Share of WS customers (%) | Share of WS customer Azure revenue (%) | Share of WS customers (%) | Share of WS customer Azure revenue (%) |
| **Storage, Virtual Network** | 99% | 99.58% | 99.89% | 99.97% |
| **Storage, Virtual Network, Bandwidth** | 97.88% | 98.72% | 99.23% | 99.40% |
| **Storage, Virtual Network, Bandwidth, Backup** | 81.13% | 86.54% | 82.05% | 87.05% |

*Notes: Based on Microsoft UK Azure data, this analysis is restricted to customers with positive WS AHB vCore hours and zero Linux hours as well as above $10,000 spend. The services are defined based on Service Level 2 (SL2) data; Service Level 2 is the level of granularity with which Microsoft records data in its financial systems*

## Annex 2 – Microsoft's updates to SQL Server Data

75. In Microsoft's follow-up response to the CMA's Licensing working paper,[57] we submitted an analysis of SQL Server SPLA licensing costs by assuming three flat ratios of Enterprise to Standard Edition for AWS' usage data: 50:50, 40:60, and 25:75. As Microsoft was unable to provide to the CMA the usage by edition for all SQL Server vCore hours, the CMA has adopted these same flat ratios for Azure customers and assumes that these ratios are broadly accurate.[58]

76. For the purposes of this response, Microsoft has provided a method to disaggregate SQL PAYG usage data by edition. We note that the Web edition (which has a comparatively lower price than the other editions) should be included in the analysis. The methodology used is as follows:

    a. The SQL IaaS IP usage data by edition is maintained in the form of VM hours in the ordinary course of business and not normalised for the number of cores.[59]

    b. Microsoft converted the SQL IaaS VM hours to vCore hours by multiplying the hours by the number of cores recorded for the VM. The number of cores is sourced from Microsoft's Service Level 5 (SL5) data, a level of granularity in Microsoft's internal financial systems.

    c. When the SL5 description contained a range for the number of cores (e.g. 1-4), Microsoft used the higher end of the range (i.e., 4 cores).

77. This results in an edition split for Azure as seen in Table 13. The normalised VM hours are not used directly in the analysis and instead the estimated SQL server edition splits are applied to the total SQL PAYG vCore hours. The table (which also contains AWS' and GCP's relative usage from tables T.1 and T.2 of Appendix T) shows that compared to the CMA's assumptions, customers spend relatively more on Standard and Web editions, which are lower in price than the Enterprise edition.

[57] Keystone presentation to the CMA on licensing, 15 August 2024, slide 19.

[58] CMA PDR, Appendix T, para T.31.

[59] A VM hour is the number of running hours for a VM, regardless of number of vCores on the VM.

78. The mapping from SL5s to number of cores is provided in the table below. The VM hours associated with each TPID and SL5 are multiplied by this normalisation factor to arrive at the number of vCore hours per SQL edition and aggregated over the revenue buckets to create the relative usage splits shown above.

**Table 23: Normalisation factor used to convert VM hours into vCore hours by SQL Server Edition**

| Service Level 5 | Number of cores assumed |
|---|---|
| 12 vCPU VM | 12 |
| 1-4 vCPU VM | 4 |
| 16 vCPU VM | 16 |
| 20 vCPU VM | 20 |
| 24 vCPU VM | 24 |
| 32 vCPU VM | 32 |
| 4 vCPU VM | 4 |
| 44 vCPU VM | 44 |
| 48 vCPU VM | 48 |
| 6 vCPU VM | 6 |
| 64 vCPU VM | 64 |
| 8 vCPU VM | 8 |
| 80 vCPU VM | 80 |
| 96 vCPU VM | 96 |
| 104 vCPU VM | 104 |
| 72 vCPU VM | 72 |
| 2 vCPU VM | 2 |
| 1 vCPU VM | 1 |
| 112 vCPU VM | 112 |

79. Therefore, in our AEC analysis, we do not use the 25:75, 60:40 or 50:50 edition splits. Instead, we use Microsoft's edition splits as estimated above. We also use the relative usage data of AWS and GCP in their respective analyses, as a sense check.

80. We also note that the AWS SPLA prices to be used in the analysis of the CMA should be corrected from $72.27 and $283.24 for standard and enterprise editions respectively, to $72.6 and $256.96.[60]

81. In addition, we also include Web edition, using the standard SPLA price, at $4.54 per core.

[60] The correct monthly prices are provided in AWS' HSPLA contract and the standard SPLA prices for the Enterprise and Standard editions respectively. Microsoft notes that the difference in the CMA's assumed price for Standard is due to a rounding error, taking the vCore hour price provided in the Keystone datapack submission and multiplying it by 730.

# Annex 3 – AEC test full results

## METHODOLOGY

82. The as-efficient competitor test analyses the resulting margins an as-efficient competitor could make if they had Azure WS and SQL Server usage but had to pay the HSPLA cost.

83. We use the following approach.

    a. **Revenues**. We gather WS VM revenues from all WS-using customers. For the test including SQL Server, we include the SQL Server IaaS PAYG IP revenues. For storage, virtual network and bandwidth, we approximate a customer's WS-share by pro-rating these SL2 revenues based on a customer's WS share of total compute (WS + Linux) vCore hours.

    b. **Hypothetical IP cost**. We multiply total Azure WS vCore hours (and SQL Server vCore hours) by the appropriate licensing cost. For SQL Server, this varies based on the ratios of edition split we assume and if we use AWS or GCP prices.

    c. **Hardware COGS**. To approximate the VM hardware COGS, we use the COGS % from an internal Windows VM profitability estimate, shown below in Table 24.[61] We note that this estimate is an internal one which is not vetted to the same extent as the metrics Microsoft maintains in the ordinary course of business for externally reported financials. It is also in terms of Microsoft's financial year (July – June) rather than calendar year, but we consider this to be a reasonable proxy.

    d. **Other services COGS**. To approximate the general COGS for the additional services, we use Azure standalone services COGS, sourced from the FY24 product breakdown report.[62]

**Table 24: Global Windows VMs revenues, COGS and gross margins**

| Windows Compute+IP | FY23 | FY24 |
|---|---|---|
| Compute revenue | $ 4,191,115,572.59 | $ 4,688,019,791.23 |
| IP revenue | $ 1,344,200,917.47 | $ 1,415,842,126.30 |
| **Net revenue [1]** | **$ 5,535,316,490.06** | **$ 6,103,861,917.53** |
| Depreciation | $ 471,443,069.60 | $ 557,390,109.40 |
| Hosting Cost | $ 891,471,734.79 | $ 901,724,727.01 |
| **Capacity COGS [2]** | **$ 1,362,914,804.39** | **$ 1,459,114,836.41** |
| CSS | $ 125,872,253.84 | $ 109,862,488.28 |
| Marketing | $ 276,349,844.78 | $ 246,636,387.76 |
| People Cost | $ 322,953,248.44 | $ 246,399,593.96 |
| Royalties/Others | $ 80,255,512.31 | $ 88,799,057.80 |
| **Non-Capacity COGS [3]** | **$ 805,430,859.36** | **$ 691,697,527.80** |
| **Total COGS [4] = [2] + [3]** | **$ 2,168,345,663.75** | **$ 2,150,812,364.21** |
| **Gross Margin [5] = [1]-[4]** | **$ 3,366,970,826.31** | **$ 3,953,049,553.32** |
| **Gross Margin % [6] = [5]/[1]** | **61%** | **65%** |

[61] Note that we apply FY24 COGS % to 2024 (and FY23 COGS % to 2023) in the AEC tables.

[62] Note that we apply FY24 Azure COGS % to 2024 (and FY23 Azure COGS % to 2023) in the AEC tables.

| | | |
|---|---|---|
| **COGS % [7] = 100% - [6]** | **39%** | **35%** |

*Source: Microsoft data. Notes: Depreciation costs include server depreciation given that datacenter depreciation is included in Hosting costs. Hosting costs include costs of datacenters consumed by compute server and prorate the share of available datacenter space. "CSS" stands for Customer Service and Support costs, which are allocated based on call volume. People costs include costs for people directly working on workloads plus a share of people working on Azure as a platform (outage, monitoring, etc.) "Royalties/Others" represents other costs which includes items such as acquisition amortization, royalties, billing costs like credit card fees, and commerce costs (system that interfaces with customers/collects payments from customers).*

84. The resulting margin is calculated as the sum of revenues minus licensing cost minus hardware VM cost minus the other services COGS.

85. As this is an as-efficient competitor analysis, we use Azure's own COGS. We believe this is a good approximation for AWS, who has been reported to have ~60% gross margins.[63] We also believe this to be a good approximation for GCP. Although GCP is smaller, evidence submitted by Google themselves suggest that GCP [redacted].[64]

86. Therefore, the results in these tests are reliable in predicting the (lack of) foreclosure effect of hyperscale SPLA licensing on AWS/GCP who are as-efficient competitors.

**WS USERS**

**Table 25: As-efficient competitor results for WS UK customers, 2024**

| | | [A] WS VMs, storage and virtual network | [B] WS VMs, storage, virtual network and bandwidth | Comments |
|---|---|---|---|---|
| *Revenues* | WS VMs | $500M | $500M | |
| *Revenues* | Storage | $463M | $463M | Pro-rated storage spend for all WS-using UK customers. |
| *Revenues* | Virtual Network | $33M | $33M | Pro-rated virtual network spend for all WS-using UK customers. |
| *Revenues* | Bandwidth | - | $20M | Pro-rated bandwidth spend for all WS-using UK customers. |
| *Hypothetical Costs – IP* | *Azure WS vCore hours* | *9,036M* | *9,036M* | |
| *Hypothetical Costs – IP* | Azure WS vCore hour licensing cost | $353M | $353M | Using the hyperscaler cost of $0.0391 per vCore hours. |
| *Costs – Compute* | *Azure HW COGS % of WS revenue* | 35% | 35% | 1 minus FY24[65] Azure Windows VMs gross margin. |
| *Costs – Compute* | Estimated compute WS VM costs | $175M | $175M | |

[63] *See* How Amazon Web Services makes money: Estimated margins by service.

[64] *See* [redacted].

[65] Microsoft's fiscal year runs from July to June i.e., FY24 runs from July 2023 to June 2024.

| | | | | |
|---|---|---|---|---|
| *Costs - Other* | *Azure COGS % (global)* | 40% | 40% | Sourced from Product Breakdown Report. It is 1 minus the gross margin of C+E Azure Standalone Services in FY24. |
| *Costs - Other* | Estimated other COGS | $201M | $209M | |
| ***Margin*** | **Margin over licensing costs and other COGS** | **$268M** | **$280M** | |
| ***Margin*** | **Margin over licensing costs and other COGS (%)** | **27%** | **28%** | |

*Source: Keystone analysis of Azure data. Notes: Based on Microsoft UK Azure data, this table is restricted to customers with positive WS vCore hours in a given year and at least $10,000 annual total cloud spend in a given year. Additional Service Level 2 (SL2) service revenues are prorated based on a customer's WS vCore share of total compute vCore hours. Service Level 2 is the level of granularity with which Microsoft records data in its financial systems.*

**Table 26: As-efficient competitor results for WS UK customers with 100% WS Compute, 2024**

| | | [A] WS, storage and virtual network | [B] WS, storage, virtual network and bandwidth | Comments |
|---|---|---|---|---|
| *Revenues* | WS | $33M | $33M | |
| *Revenues* | Storage | $32M | $32M | Pro-rated storage spend for all WS-using UK customers. |
| *Revenues* | Virtual Network | $1M | $1M | Pro-rated virtual network spend for all WS-using UK customers. |
| *Revenues* | Bandwidth | - | $2M | Pro-rated bandwidth spend for all WS-using UK customers. |
| *Hypothetical Costs – IP* | *Azure WS vCore hours* | *477M* | *477M* | |
| *Hypothetical Costs – IP* | Azure WS vCore hour licensing cost | $19M | $19M | Using the hyperscaler cost of $0.0391 per vCore hours. |
| *Costs – Compute* | *Azure HW COGS % of WS revenue* | 35% | 35% | 1 minus FY24 Azure Windows VMs gross margin. |

| | | | | |
|---|---|---|---|---|
| *Costs – Compute* | Estimated compute WS VM costs | $12M | $12M | |
| *Costs - Other* | *Azure COGS % (global)* | 40% | 40% | Sourced from Product Breakdown Report. It is 1 minus the gross margin of C+E Azure Standalone Services in FY24. |
| *Costs - Other* | Estimated other COGS | $14M | $14M | |
| ***Margin*** | **Margin over licensing costs and other COGS** | **$23M** | **$24M** | |
| ***Margin*** | **Margin over licensing costs and other COGS (%)** | **34%** | **35%** | |

*Source: Keystone analysis of Azure data. Notes: Based on Microsoft UK Azure data, this table is restricted to customers with positive WS vCore hours and zero Linux vCore hours in a given year and at least $10,000 annual total cloud spend in a given year. Additional Service Level 2 (SL2) service revenues are prorated based on a customer's WS vCore share of total compute vCore hours. Service Level 2 is the level of granularity with which Microsoft records data in its financial systems.*

**Table 27: As-efficient competitor results for WS AHB UK customers, 2024**

| | | [A] WS AHB, storage and virtual network | [B] WS AHB, storage, virtual network and bandwidth | Comments |
|---|---|---|---|---|
| *Revenues* | WS AHB | $174M | $174M | |
| *Revenues* | Storage | $215M | $215M | Pro-rated storage spend for all WS-using UK customers. |
| *Revenues* | Virtual Network | $17M | $17M | Pro-rated virtual network spend for all WS-using UK customers. |
| *Revenues* | Bandwidth | - | $8M | Pro-rated bandwidth spend for all WS-using UK customers. |
| *Hypothetical Costs – IP* | *Azure WS AHB vCore hours* | *4,343M* | *4,343M* | |
| *Hypothetical Costs – IP* | Azure WS AHB vCore hour licensing cost | $170M | $170M | Using the hyperscaler cost of $0.0391 per vCore hours. |
| *Costs – Compute* | *Azure HW COGS % of WS revenue* | 35% | 35% | 1 minus FY24 Azure Windows VMs gross margin. |
| *Costs – Compute* | Estimated compute WS VM costs | $61M | $61M | |

| | | | | |
|---|---|---|---|---|
| *Costs - Other* | *Azure COGS % (global)* | 40% | 40% | Sourced from Product Breakdown Report. It is 1 minus the gross margin of C+E Azure Standalone Services in FY24. |
| *Costs - Other* | Estimated other COGS | $94M | $97M | |
| ***Margin*** | **Margin over licensing costs and other COGS** | **$81M** | **$86M** | |
| ***Margin*** | **Margin over licensing costs and other COGS (%)** | **20%** | **21%** | |

*Source: Keystone analysis of Azure data. Notes: Based on Microsoft UK Azure data, this table is restricted to customers with positive WS AHB vCore hours in a given year and at least $10,000 annual total cloud spend in a given year. Additional Service Level 2 (SL2) service revenues are prorated based on a customer's WS vCore share of total compute vCore hours. Service Level 2 is the level of granularity with which Microsoft records data in its financial systems.*

**Table 28: As-efficient competitor results for WS AHB UK customers with 100% WS Compute, 2024**

| | | [A] WS AHB, storage and virtual network | [B] WS AHB, storage, virtual network and bandwidth | Comments |
|---|---|---|---|---|
| *Revenues* | WS AHB | $4M | $4M | |
| *Revenues* | Storage | $5M | $5M | Pro-rated storage spend for all WS-using UK customers. |
| *Revenues* | Virtual Network | $0.2M | $0.2M | Pro-rated virtual network spend for all WS-using UK customers. |
| *Revenues* | Bandwidth | - | $0.3M | Pro-rated bandwidth spend for all WS-using UK customers. |
| *Hypothetical Costs – IP* | *Azure WS AHB vCore hours* | *80M* | *80M* | |
| *Hypothetical Costs – IP* | Azure WS AHB vCore hour licensing cost | $3M | $3M | Using the hyperscaler cost of $0.0391 per vCore hours. |
| *Costs – Compute* | *Azure HW COGS % of WS revenue* | 35% | 35% | 1 minus FY24 Azure Windows VMs gross margin. |
| *Costs – Compute* | Estimated compute WS VM costs | $1M | $1M | |
| *Costs - Other* | *Azure COGS % (global)* | 40% | 40% | Sourced from Product Breakdown Report. It is 1 minus the gross margin of |

| | | | | |
|---|---|---|---|---|
| | | | | C+E Azure Standalone Services in FY24. |
| *Costs - Other* | Estimated other COGS | $2M | $2M | |
| ***Margin*** | **Margin over licensing costs and other COGS** | **$2M** | **$3M** | |
| ***Margin*** | **Margin over licensing costs and other COGS (%)** | **27%** | **28%** | |

*Source: Keystone analysis of Azure data. Notes: Based on Microsoft UK Azure data, this table is restricted to customers with positive WS AHB vCore hours and zero Linux vCore hours in a given year and at least $10,000 annual total cloud spend in a given year. Additional Service Level 2 (SL2) service revenues are prorated based on a customer's WS vCore share of total compute vCore hours. Service Level 2 is the level of granularity with which Microsoft records data in its financial systems.*

**Table 29: As-efficient competitor results for WS non-AHB UK customers, 2024**

| | | [A] WS non-AHB, storage and virtual network | [B] WS non-AHB, storage, virtual network and bandwidth | Comments |
|---|---|---|---|---|
| *Revenues* | WS non-AHB | $327M | $327M | |
| *Revenues* | Storage | $248M | $248M | Pro-rated storage spend for all WS-using UK customers. |
| *Revenues* | Virtual Network | $17M | $17M | Pro-rated virtual network spend for all WS-using UK customers. |
| *Revenues* | Bandwidth | - | $12M | Pro-rated bandwidth spend for all WS-using UK customers. |
| *Hypothetical Costs – IP* | *Azure WS non-AHB vCore hours* | *4,693M* | *4,693M* | |
| *Hypothetical Costs – IP* | Azure WS non-AHB vCore hour licensing cost | $183M | $183M | Using the hyperscaler cost of $0.0391 per vCore hours. |
| *Costs – Compute* | *Azure HW COGS % of WS revenue* | 35% | 35% | 1 minus FY24 Azure Windows VMs gross margin. |
| *Costs – Compute* | Estimated compute WS VM costs | $114M | $114M | |

| | | | | |
|---|---|---|---|---|
| *Costs - Other* | *Azure COGS % (global)* | 40% | 40% | Sourced from Product Breakdown Report. It is 1 minus the gross margin of C+E Azure Standalone Services in FY24. |
| *Costs - Other* | Estimated other COGS | $107M | $112M | |
| ***Margin*** | **Margin over licensing costs and other COGS** | **$187M** | **$194M** | |
| ***Margin*** | **Margin over licensing costs and other COGS (%)** | **32%** | **32%** | |

*Source: Keystone analysis of Azure data. Notes: Based on Microsoft UK Azure data, this table is restricted to customers with positive non-AHB WS vCore hours in a given year and at least $10,000 annual total cloud spend in a given year. Additional Service Level 2 (SL2) service revenues are prorated based on a customer's WS vCore share of total compute vCore hours. Service Level 2 is the level of granularity with which Microsoft records data in its financial systems.*

**Table 30: As-efficient competitor results for WS non-AHB UK customers with 100% WS Compute, 2024**

| | | [A] WS non-AHB, storage and virtual network | [B] WS non-AHB, storage, virtual network and bandwidth | Comments |
|---|---|---|---|---|
| *Revenues* | WS non-AHB | $29M | $29M | |
| *Revenues* | Storage | $27M | $27M | Pro-rated storage spend for all WS-using UK customers. |
| *Revenues* | Virtual Network | $1M | $1M | Pro-rated virtual network spend for all WS-using UK customers. |
| *Revenues* | Bandwidth | - | $2M | Pro-rated bandwidth spend for all WS-using UK customers. |
| *Hypothetical Costs – IP* | *Azure WS non-AHB vCore hours* | *397M* | *397M* | |
| *Hypothetical Costs – IP* | Azure WS non-AHB vCore hour licensing cost | $16M | $16M | Using the hyperscaler cost of $0.0391 per vCore hours. |
| *Costs – Compute* | *Azure HW COGS % of WS revenue* | 35% | 35% | 1 minus FY24 Azure Windows VMs gross margin. |
| *Costs – Compute* | Estimated compute WS VM costs | $10M | $10M | |

| | | | | |
|---|---|---|---|---|
| *Costs - Other* | *Azure COGS % (global)* | 40% | 40% | Sourced from Product Breakdown Report. It is 1 minus the gross margin of C+E Azure Standalone Services in FY24. |
| *Costs - Other* | Estimated other COGS | $11M | $12M | |
| ***Margin*** | **Margin over licensing costs and other COGS** | **$20M** | **$21M** | |
| ***Margin*** | **Margin over licensing costs and other COGS (%)** | **35%** | **36%** | |

*Source: Keystone analysis of Azure data. Notes: Based on Microsoft UK Azure data, this table is restricted to customers with positive non-AHB WS vCore hours and zero Linux vCore hours in a given year and at least $10,000 annual total cloud spend in a given year. Additional Service Level 2 (SL2) service revenues are prorated based on a customer's WS vCore share of total compute vCore hours. Service Level 2 is the level of granularity with which Microsoft records data in its financial systems.*

### WS AND SQL PAYG USERS

87. The tables below contain the AEC results presented in Section 6.1, but with additional details and steps underlying the calculations.

**Table 31: As-efficient competitor results for WS and SQL UK customers, 2024**

| | | [A] WS, SQL, storage and virtual network | [B] WS, SQL, storage, virtual network and bandwidth | Comments |
|---|---|---|---|---|
| *Revenues* | WS | $381M | $381M | Pro-rated storage spend for all WS-using UK customers. |
| *Revenues* | SQL PAYG | $69M | $69M | Pro-rated virtual network spend for all WS-using UK customers. |
| *Revenues* | Storage | $346M | $346M | Pro-rated bandwidth spend for all WS-using UK customers. |
| *Revenues* | Virtual Network | $25M | $25M | |
| *Revenues* | Bandwidth | - | $13M | |
| *Hypothetical Costs – IP* | *Azure WS vCore hours* | *6,958M* | *6,958M* | |
| *Hypothetical Costs – IP* | Azure WS vCore hour licensing cost | *272M* | *272M* | Using the hyperscaler cost of $0.0391 per vCore hours. |
| *Hypothetical Costs – IP* | *Azure SQL PAYG hours* | *645M* | *645M* | |

| | | | | |
|---|---|---|---|---|
| *Hypothetical Costs – IP* | Azure SQL PAYG hour licensing cost | *$78M - $106M* | *$78M - $106M* | Using the AWS and GCP SPLA prices for each SQL edition (Standard, Enterprise and Web) and Azure/AWS/GCP relative usage data by edition |
| *Costs – Compute* | *Azure HW COGS % of WS revenue* | *35%* | *35%* | 1 minus FY24 Azure Windows VMs gross margin. |
| *Costs - Compute* | Estimated compute WS VM costs | *133M* | *133M* | |
| *Costs - Other* | *Azure COGS % (global)* | *40%* | *40%* | Sourced from Product Breakdown Report. It is 1 minus the gross margin of C+E Azure Standalone Services in FY24. |
| *Costs - Other* | Estimated other COGS | *150M* | *155M* | |
| ***Margin*** | **Margin over licensing costs and other COGS** | **$159M - $187M** | **$167M - $195M** | |
| ***Margin*** | **Margin over licensing costs and other COGS (%)** | **19% - 23%** | **20% - 23%** | |

*Source: Keystone analysis of Azure data. Notes: Based on Microsoft UK Azure data, this table is restricted to customers with positive WS vCore hours, positive SQL Server IaaS PAYG vCore hours in a given year and at least $10,000 annual total cloud spend in a given year. Additional Service Level 2 (SL2) service revenues are prorated based on a customer's WS vCore share of total compute vCore hours. Service Level 2 is the level of granularity with which Microsoft records data in its financial systems.*

### Table 32: As-efficient competitor results for WS (100% compute) and SQL UK customers, 2024

| | | [A] WS, SQL, storage and virtual network | [B] WS, SQL, storage, virtual network and bandwidth | Comments |
|---|---|---|---|---|
| *Revenues* | WS | $14M | $14M | Pro-rated storage spend for all WS-using UK customers. |
| *Revenues* | SQL PAYG | $6M | $6M | Pro-rated virtual network spend for all WS-using UK customers. |
| *Revenues* | Storage | $11M | $11M | Pro-rated bandwidth spend for all WS-using UK customers. |
| *Revenues* | Virtual Network | $0.4M | $0.4M | |
| *Revenues* | Bandwidth | - | $1M | |

| | | | | |
|---|---|---|---|---|
| *Hypothetical Costs – IP* | *Azure WS vCore hours* | *190M* | *190M* | |
| *Hypothetical Costs – IP* | Azure WS vCore hour licensing cost | *7M* | *7M* | Using the hyperscaler cost of $0.0391 per vCore hours. |
| *Hypothetical Costs – IP* | *Azure SQL PAYG hours* | *65M* | *65M* | |
| *Hypothetical Costs – IP* | Azure SQL PAYG hour licensing cost | *$6M - $7M* | *$6M - $7M* | Using the AWS and GCP SPLA prices for each SQL edition (Standard, Enterprise and Web) and Azure/AWS/GCP relative usage data by edition |
| *Costs – Compute* | *Azure HW COGS % of WS revenue* | *35%* | *35%* | 1 minus FY24 Azure Windows VMs gross margin. |
| *Costs – Compute* | Estimated compute WS VM costs | *5M* | *5M* | |
| *Costs - Other* | *Azure COGS % (global)* | *40%* | *40%* | Sourced from Product Breakdown Report. It is 1 minus the gross margin of C+E Azure Standalone Services in FY24. |
| *Costs - Other* | Estimated other COGS | *5M* | *5M* | |
| ***Margin*** | **Margin over licensing costs and other COGS** | **$7M - $8M** | **$7M - $9M** | |
| ***Margin*** | **Margin over licensing costs and other COGS (%)** | **23% - 26%** | **23% - 27%** | |

*Source: Keystone analysis of Azure data. Notes: Based on Microsoft UK Azure data, this table is restricted to customers with positive WS vCore hours and zero Linux vCore hours in a given year, positive SQL Server IaaS PAYG vCore hours in a given year and at least $10,000 annual total cloud spend in a given year. Additional Service Level 2 (SL2) service revenues are prorated based on a customer's WS vCore share of total compute vCore hours. Service Level 2 is the level of granularity with which Microsoft records data in its financial systems.*

**Table 33: As-efficient competitor results for WS AHB and SQL UK customers, 2024**

| | | [A] WS AHB, SQL, storage and virtual network | [B] WS AHB, SQL, storage, virtual network and bandwidth | Comments |
|---|---|---|---|---|
| *Revenues* | WS AHB | $139M | $139M | Pro-rated storage spend for all WS-using UK customers. |
| *Revenues* | SQL PAYG | $51M | $51M | Pro-rated virtual network spend for all |

| | | | | |
|---|---|---|---|---|
| | | | | WS-using UK customers. |
| *Revenues* | Storage | $169M | $169M | Pro-rated bandwidth spend for all WS-using UK customers. |
| *Revenues* | Virtual Network | $13M | $13M | |
| *Revenues* | Bandwidth | - | $6M | |
| *Hypothetical Costs – IP* | *Azure WS AHB vCore hours* | *3,492M* | *3,492M* | |
| *Hypothetical Costs – IP* | Azure WS AHB vCore hour licensing cost | *137M* | *137M* | Using the hyperscaler cost of $0.0391 per vCore hours. |
| *Hypothetical Costs – IP* | *Azure SQL PAYG hours* | *475M* | *475M* | |
| *Hypothetical Costs – IP* | Azure SQL PAYG hour licensing cost | *$61M - $86M* | *$61M - $86M* | Using the AWS and GCP SPLA prices for each SQL edition (Standard, Enterprise and Web) and Azure/AWS/GCP relative usage data by edition |
| *Costs – Compute* | *Azure HW COGS % of WS revenue* | *35%* | *35%* | 1 minus FY24 Azure Windows VMs gross margin. |
| *Costs – Compute* | Estimated compute WS VM costs | *49M* | *49M* | |
| *Costs - Other* | *Azure COGS % (global)* | *40%* | *40%* | Sourced from Product Breakdown Report. It is 1 minus the gross margin of C+E Azure Standalone Services in FY24. |
| *Costs - Other* | Estimated other COGS | *74M* | *76M* | |
| ***Margin*** | **Margin over licensing costs and other COGS** | **$27M - $52M** | **$31M - $56M** | |
| ***Margin*** | **Margin over licensing costs and other COGS (%)** | **7% - 14%** | **8% - 15%** | |

*Source: Keystone analysis of Azure data. Notes: Based on Microsoft UK Azure data, this table is restricted to customers with positive WS AHB vCore hours in a given year, positive SQL Server IaaS PAYG vCore hours in a given year and at least $10,000 annual total cloud spend in a given year. Additional Service Level 2 (SL2) service revenues are prorated based on a customer's WS vCore share of total compute vCore hours. Service Level 2 is the level of granularity with which Microsoft records data in its financial systems.*

**Table 34: As-efficient competitor results for WS AHB and SQL UK customers (WS 100% Compute), 2024**

| | | [A] WS AHB, SQL, storage and virtual network | [B] WS AHB, SQL, storage, virtual network and bandwidth | Comments |
|---|---|---|---|---|
| *Revenues* | WS AHB | $1M | $1M | Pro-rated storage spend for all WS-using UK customers. |
| *Revenues* | SQL PAYG | $1M | $1M | Pro-rated virtual network spend for all WS-using UK customers. |
| *Revenues* | Storage | $1M | $1M | Pro-rated bandwidth spend for all WS-using UK customers. |
| *Revenues* | Virtual Network | $0.04 M | $0.04 M | |
| *Revenues* | Bandwidth | - | $0.1 M | |
| *Hypothetical Costs – IP* | *Azure WS AHB vCore hours* | *21M* | *21M* | |
| *Hypothetical Costs – IP* | Azure WS AHB vCore hour licensing cost | *1M* | *1M* | Using the hyperscaler cost of $0.0391 per vCore hours. |
| *Hypothetical Costs – IP* | *Azure SQL PAYG hours* | *13M* | *13M* | |
| *Hypothetical Costs – IP* | Azure SQL PAYG hour licensing cost | *$1M - $1M* | *$1M - $1M* | Using the AWS and GCP SPLA prices for each SQL edition (Standard, Enterprise and Web) and Azure/AWS/GCP relative usage data by edition |
| *Costs – Compute* | *Azure HW COGS % of WS revenue* | *35%* | *35%* | 1 minus FY24 Azure Windows VMs gross margin. |
| *Costs – Compute* | Estimated compute WS VM costs | *$0.3 M* | *$0.3 M* | |
| *Costs - Other* | *Azure COGS % (global)* | *40%* | *40%* | Sourced from Product Breakdown Report. It is 1 minus the gross |

| | | | | |
|---|---|---|---|---|
| | | | | margin of C+E Azure Standalone Services in FY24. |
| *Costs - Other* | Estimated other COGS | *1M* | *1M* | |
| ***Margin*** | **Margin over licensing costs and other COGS** | **$0M - $1M** | **$0M - $1M** | |
| ***Margin*** | **Margin over licensing costs and other COGS (%)** | **8% - 15%** | **10% - 16%** | |

*Source: Keystone analysis of Azure data. Notes: Based on Microsoft UK Azure data, this table is restricted to customers with positive WS AHB vCore hours and zero Linux vCore hours in a given year, positive SQL Server IaaS PAYG vCore hours in a given year and at least $10,000 annual total cloud spend in a given year. Additional Service Level 2 (SL2) service revenues are prorated based on a customer's WS vCore share of total compute vCore hours. Service Level 2 is the level of granularity with which Microsoft records data in its financial systems.*

**Table 35: As-efficient competitor results for WS non-AHB and SQL UK customers, 2024**

| | | [A] WS non-AHB, SQL, storage and virtual network | [B] WS non-AHB, SQL, storage, virtual network and bandwidth | Comments |
|---|---|---|---|---|
| *Revenues* | WS non-AHB | $242M | $242M | Pro-rated storage spend for all WS-using UK customers. |
| *Revenues* | SQL PAYG | $68M | $68M | Pro-rated virtual network spend for all WS-using UK customers. |
| *Revenues* | Storage | $177M | $177M | Pro-rated bandwidth spend for all WS-using UK customers. |
| *Revenues* | Virtual Network | $11M | $11M | |
| *Revenues* | Bandwidth | - | $7M | |
| *Hypothetical Costs – IP* | *Azure WS non-AHB vCore hours* | *3,466M* | *3,466M* | |
| *Hypothetical Costs – IP* | Azure WS non-AHB vCore hour licensing cost | *136M* | *136M* | Using the hyperscaler cost of $0.0391 per vCore hours. |
| *Hypothetical Costs – IP* | *Azure SQL PAYG hours* | *641M* | *641M* | |
| *Hypothetical Costs – IP* | Azure SQL PAYG hour licensing cost | *$78M - $106M* | *$78M - $106M* | Using the AWS and GCP SPLA prices for each SQL edition (Standard, Enterprise and Web) |

| | | | | |
|---|---|---|---|---|
| | | | | and Azure/AWS/GCP relative usage data by edition |
| *Costs – Compute* | *Azure HW COGS % of WS revenue* | *35%* | *35%* | 1 minus FY24 Azure Windows VMs gross margin. |
| *Costs – Compute* | Estimated compute WS VM costs | *85M* | *85M* | |
| *Costs - Other* | *Azure COGS % (global)* | *40%* | *40%* | Sourced from Product Breakdown Report. It is 1 minus the gross margin of C+E Azure Standalone Services in FY24. |
| *Costs - Other* | Estimated other COGS | *76M* | *79M* | |
| ***Margin*** | **Margin over licensing costs and other COGS** | **$97M - $125M** | **$101M - $129M** | |
| ***Margin*** | **Margin over licensing costs and other COGS (%)** | **19% - 25%** | **20% - 25%** | |

*Source: Keystone analysis of Azure data. Notes: Based on Microsoft UK Azure data, this table is restricted to customers with positive non-AHB WS vCore hours in a given year, positive SQL Server IaaS PAYG vCore hours in a given year and at least $10,000 annual total cloud spend in a given year. Additional Service Level 2 (SL2) service revenues are prorated based on a customer's WS vCore share of total compute vCore hours. Service Level 2 is the level of granularity with which Microsoft records data in its financial systems.*

**Table 36: As-efficient competitor results for WS non-AHB and SQL UK customers (WS 100% Compute), 2024**

| | | [A] WS non-AHB, SQL, storage and virtual network | [B] WS non-AHB, SQL, storage, virtual network and bandwidth | Comments |
|---|---|---|---|---|
| *Revenues* | WS non-AHB | $13M | $13M | Pro-rated storage spend for all WS-using UK customers. |
| *Revenues* | SQL PAYG | $6M | $6M | Pro-rated virtual network spend for all WS-using UK customers. |
| *Revenues* | Storage | $10M | $10M | Pro-rated bandwidth spend for all WS-using UK customers. |
| *Revenues* | Virtual Network | $0.4M | $0.4M | |
| *Revenues* | Bandwidth | - | $0.5M | |
| *Hypothetical Costs – IP* | *Azure WS non-AHB vCore hours* | *169M* | *169M* | |

| | | | | |
|---|---|---|---|---|
| *Hypothetical Costs – IP* | Azure WS non-AHB vCore hour licensing cost | *7M* | *7M* | Using the hyperscaler cost of $0.0391 per vCore hours. |
| *Hypothetical Costs – IP* | *Azure SQL PAYG hours* | *63M* | *63M* | |
| *Hypothetical Costs – IP* | Azure SQL PAYG hour licensing cost | *$6M - $7M* | *$6M - $7M* | Using the AWS and GCP SPLA prices for each SQL edition (Standard, Enterprise and Web) and Azure/AWS/GCP relative usage data by edition |
| *Costs – Compute* | *Azure HW COGS % of WS revenue* | *35%* | *35%* | 1 minus FY24 Azure Windows VMs gross margin. |
| *Costs – Compute* | Estimated compute WS VM costs | *4M* | *4M* | |
| *Costs - Other* | *Azure COGS % (global)* | *40%* | *40%* | Sourced from Product Breakdown Report. It is 1 minus the gross margin of C+E Azure Standalone Services in FY24. |
| *Costs - Other* | Estimated other COGS | *4M* | *4M* | |
| ***Margin*** | **Margin over licensing costs and other COGS** | **$7M - $8M** | **$7M - $8M** | |
| ***Margin*** | **Margin over licensing costs and other COGS (%)** | **23% - 27%** | **24% - 27%** | |

*Source: Keystone analysis of Azure data. Notes: Based on Microsoft UK Azure data, this table is restricted to customers with positive non-AHB WS vCore hours and zero Linux vCore hours in a given year, positive SQL Server IaaS PAYG vCore hours in a given year and at least $10,000 annual total cloud spend in a given year. Additional Service Level 2 (SL2) service revenues are prorated based on a customer's WS vCore share of total compute vCore hours. Service Level 2 is the level of granularity with which Microsoft records data in its financial systems.*