

Non-Confidential Version

CMA CLOUD SERVICES MARKET INVESTIGATION

MICROSOFT RESPONSE TO THE COMPETITIVE LANDSCAPE, COMMITTED SPEND AGREEMENTS AND EGRESS FEES WORKING PAPERS



Linklaters

1 July 2024

LATHAM & WATKINS

Cloud Services Market Investigation

Microsoft Response to the Competitive Landscape, Committed Spend Agreements and Egress Fees Working Papers

1 Introduction

- (1) Microsoft welcomes the opportunity to respond to the CMA's emerging thinking in its Updated Issues Statement ("UIS") and Working Papers.
- (2) This submission specifically addresses the CMA's overarching emerging thinking captured in the UIS and the Competitive Landscape Working Paper ("CLWP"); see **Section 2**. In addition, **Section 3** responds to the CMA's analysis of committed spend agreements ("CSAs") while **Section 4** responds to the Egress Fees Working Paper. Given their later publication date, Microsoft's responses to the Licensing and Technical Barriers Working Papers will follow separately.

1.1 The Working Papers and the CMA's emerging thinking as yet fail to engage with the direct and consistent evidence of positive customer outcomes. Nor do they articulate an "even more competitive" counterfactual, against which well-functioning cloud market performance would realistically be "not good enough"

- (3) The evidence collected so far during the CMA inquiry (and the Ofcom market study) clearly shows that the cloud market is known for rapid and dynamic innovation, record levels of investment, with at least three firms (Amazon, Google and Microsoft) investing tens of billions each and every year, falling prices in real terms, significant discounts, customer advantages with overall reduced costs compared to on-premises deployments, rapid growth, with at least six firms all growing at double-digit rates and likely far more than that, high customer satisfaction, and customers making use of multi-cloud deployments when it works for them. These baseline facts are essentially not disputed and are in fact supported throughout the CMA's Working Papers; so far however these outcomes are not given due weight and are glossed over in the conclusions. The emerging thinking in the Working Papers fails to articulate (let alone quantify or evidence) a "counterfactual" against which today's market would allegedly fall short, and harm customers; rather, it acknowledges explicitly that for innovation levels, quality and switching levels "there is no clear counterfactual to compare outcomes with what they might be in a well-functioning market"¹ and none is posited as to "even more competitive" pricing outcomes.²

1.2 In lieu of an overall assessment of customer outcomes, and intense dynamic rivalry, the Working Papers zero in on a few retrospective metrics to assert there are "indicators" of market power

- (4) Over the 2019-22 data period analysed, Amazon and Microsoft held relatively high share of supply in the UK and they currently enjoy profits that exceed the weighted cost of capital. Given these metrics, the Working Papers then speculate about theoretical harms and potential changes in incentives that *may* come to pass in the future and *may* lead to negative customer outcomes. The evidence, however, clearly shows that the cloud market rivalry is not static – it has not, in its almost 20 years of existence, reached a steady-state of equilibrium and it would be unwise to assume or predict a future date when it will. The CMA

¹ CLWP, paragraphs 6.4(b), 6.19 and 9.18; UIS at paragraph 27.

² Cf. CLWP, paragraphs 6.10-11 and 9.18; UIS at paragraph 27.

should therefore carefully consider any potential unintended consequences of intervening today against any possible hypothetical benefits of doing it.

1.3 The Working Papers take an erroneous step backward from Ofcom’s “Phase 1” review by marginalising dynamic rivalry from Google

- (5) Ofcom relied on the approach taken in all the main third-party industry studies to label Google as a “hyperscaler” in the same class as Amazon and Microsoft (differences being ones of degree, not competitor class). Google has the capability to operate at hyperscale because it enjoys single-firm economics and Google Cloud Platform (“GCP”) is one of the firm’s priority areas for investment: GCP is not an independent company or regulated subsidiary and yet the CMA primarily treats GCP as if it were; indeed, the Working Papers appear to make a category mistake in characterising Google as akin to Oracle, IBM and Aruba, rather than the Microsoft and Amazon for competitive assessment purposes. This is clearly incorrect on the evidence in front of the CMA. Whatever barriers to expansion others may face, these do not apply to – of all competitors – Google. At over \$2 trillion market capitalisation, it is worth more than Amazon. At over \$300 billion in annual revenue, and over \$85 billion in EBIT, Google has balance sheet firepower, evidenced by cloud capex spending on a par with Microsoft and with a capital intensity above Amazon’s.

1.4 The Working Papers miss the fact that CSAs are central to effective price competition between competing cloud providers, resulting in lower prices and supporting the significant levels of long-term investment in UK cloud infrastructure

- (6) As evidenced by Microsoft’s customer case studies, CSAs are a key mechanism through which all cloud providers compete for customers (in particular, larger customers) as they allow providers to offer significant discounts in exchange for a committed consumption of cloud services (and to assist with workload migration). Discounting between providers is part and parcel of effective competition and encourages customer switching and multi-clouding (given that there is no exclusivity obligation on the customer).
- (7) The Working Papers largely ignore these empirical customer benefits in the analysis and instead focus on a highly-stylised analysis of “sticky” v. “contested” demand (which does not align with the evidence in front of the CMA given the differentiated nature of cloud services / demand, and that [X]). Any proposed intervention by the CMA would likely lead to worse outcomes for UK customers – i.e. higher cloud consumption prices, less competition for switching / migrating / multi-cloud customers and less certainty required for long-term infrastructure investment.

1.5 The Working Paper on egress fees is based purely on theory and ignores the evidence collected clearly showing that such fees do not prevent UK customers from switching or multi-cloud deployment

- (8) The entirety of customer evidence consistently shows that egress fees are not a material concern for cloud customers considering switching or multi-cloud, particularly given that cloud providers compete on egress fees. Egress fees are an irrelevance compared to the multitude of other factors customers take into account when considering complex switching / multi-cloud decisions. The CMA’s theory of harm, on the other hand, is rooted in the hypothetical customer that requires significant data to be moved across clouds on a constant basis.
- (9) If the CMA were to force cloud providers to supply egress services for free or at below cost, this would reduce cloud providers’ incentives to innovate on bandwidth infrastructure and

could lead to excessive and inefficient use of the cloud. As a result, if the CMA feels compelled to intervene (notwithstanding the submission that egress fees do not result in an adverse effect in competition or “**AEC**”), Microsoft submits that any potential remedies should be limited to: (i) internet egress via ISP, retaining the flexibility for cloud providers to provide premium egress services; and (ii) allowing cloud providers to charge fees that recoup their costs and a reasonable return on their investments.

1.6 Inadequate foundations for AEC findings on a provisional basis

- (10) The CMA has collected a great deal of information, including the CMA’s Cloud Services Market Investigation Qualitative Customer Research Final Report (“**Jigsaw Research**”), and Microsoft has submitted a wealth of evidence to the CMA on innovation and pricing outcomes. While it is understood that the CMA carried out economic analysis in parallel across different areas of concern and, in some cases, did not have the time to adequately absorb the evidence received from the parties before publishing its Working Papers, Microsoft is concerned that too often the Working Papers seem to favour hypothetical concerns over the actual market evidence currently available to the CMA. In places, they also seem to rely upon anecdotal statements from self-interested parties over carefully produced industry research, such as that provided by Gartner and Flexera, as well as direct customer feedback.
- (11) Caution regarding intervention is warranted in such circumstances. The CMA’s emerging thinking focuses primarily on indicators (which do not justify intervention), not bad outcomes or customer detriment (which may justify intervention). At the same time, the CMA’s new DMCC powers are ideally suited to intervene at some point in the future if and when outcomes merit intervention. Unlike previous UK-centric MIRs on domestic markets, the public policy issue with any well-intentioned but ultimately distortive regulation of the UK portion of the cloud market is that UK customers might end up being worse off than their counterparts (and cross-border competitors) in the EEA, U.S., and the rest of the Western economy.

2 Response to the Competitive Landscape

2.1 The missed starting point: there is no well-defined counterfactual or definition of what a well-functioning cloud market is if not the market that exists today

2.1.1 The emerging thinking is still surprisingly preliminary on today’s clearly well-functioning market outcomes

- (12) The emerging thinking does not yet grasp the nettle on market *outcomes* – and yet it should be agreed by now given the significant available evidence that this market is currently generating unprecedented degrees of innovation, vast R&D and capex investment for the future by all three hyperscalers (Amazon, Google, and Microsoft) and more, falling prices in real terms, heavy discounting and generally supportive customer opinion evidence (Jigsaw). Pricing outcomes are inconsistent with customer lock-in (due to alleged barriers to multi-cloud and switching) and consistent with strong customer choice:
- (a) **Price declines.** Quality-adjusted real effective prices for the top five Azure products ([X]) that constitute [X]% of Azure’s total revenue, have fallen from 2018-22. Real effective prices for the top five products fell by at least [X]% as of 2022.

- (b) **No erosion in discounts to support theory of “lock in”.** The value of discounts Azure customers received has remained [redacted]% (2022-23) while committed spend discounts (“CSDs”) for UK customers [redacted]; the CMA finds this result also for Amazon.
- (c) **Innovation via feature upgrades.** Microsoft’s cumulative number of feature updates for the top 20 UK cloud services was 927 in 2023. Amazon, Google and others similarly are innovating rapidly.
- (d) **Capex that converts to customer benefits (scale, quality, innovation).** Since Azure launched, the combined capex of Amazon, Microsoft and Google has been *circa* \$700 billion. Microsoft’s R&D and capex, in particular, has accelerated since Azure first launched in 2010. Microsoft has more than tripled absolute R&D spend and seen a 12x rise in absolute capex.
- (e) **UK capacity expansion.** Microsoft has expanded the UK capacity of its dedicated Azure data centres by [redacted].
- (f) **Amazon’s early dominance now challenged.** The CMA finds that barriers to expansion for smaller cloud providers could be insurmountable, but Microsoft has surmounted those barriers to challenge Amazon and Google is doing the same. While Google’s profitability is [redacted],³ their momentum, investment, and year-on-year growth all establish that they too have managed to overcome any claimed barriers to expansion in the cloud market.
- (13) Microsoft fully appreciates that analysing [redacted], together with that from other sources, requires the CMA to undertake careful diligence, ideally market-wide pricing data. Microsoft will constructively engage on questions of interpretation as the MIR moves forward. It submits, however, that the outcomes only point one way.
- (14) For example, a fair question might be: do large incumbents (with larger shares of customer spending) discount less while smaller players, who struggle, have to discount more? Microsoft has submitted evidence showing that: (i) its prices have been falling in inflation-adjusted real terms (i.e. they are not going up); and (ii) existing customers are not “exploited” with lower discounts on renewal than when they signed up to Azure in the first place. Google and AWS have similar types of CSDs to Azure despite the fact that AWS has a much “bigger” installed base revenue share than Azure and GCP has a much “smaller” installed base revenue share. There is no evidence adduced that discounting declines with increasing market share (or rather, UK share of supply).
- (15) In short, it is difficult to know what more to expect from a “*well-functioning*” IT market – the benchmark against which market performance in an MIR is judged – in circumstances where the UIS recognises that any IT market inherently has switching costs.

2.1.2 An “unclear” counterfactual of a well-functioning market is in the “too-hard basket”

- (16) The MIR Guidelines provide that the CMA “*has to find a benchmark*” known as the counterfactual.⁴ That critical counterfactual benchmark is of a well-functioning (but not idealised textbook) market (*id*). On this essential question, the UIS and CLWP conclude this is “*unclear*” as regards innovation and quality (and makes no suggestion on prices).⁵

³ [redacted]

⁴ CMA MIR Guidelines CC3, paragraph 320.

⁵ UIS, paragraph. 27; and the CLWP at 6.19-22.

Microsoft submits that the reason is that it is difficult to construct a realistic and yet substantially *more* “well-functioning” market than the one before the CMA – against which today’s market would theoretically be malfunctioning or performing poorly or exhibiting weak competitive pressure. While the UIS makes one helpful but generic pass in this direction,⁶ this foundational flaw infects much of the specific Working Paper analysis.

- (17) In fact, the Working Papers argue that a better functioning market would likely see more switching. However, the customer evidence suggests that customers are generally satisfied with their cloud provider but would switch if there were a reason to. Of course, deploying a full enterprise’s IT infrastructure, whether in the cloud or on premises, is a complex task and not one customers do lightly. Similarly, the Working Papers suggest that more multi-cloud deployments would suggest a better performing market, largely ignoring industry studies confirming that multi-cloud is the new norm when appropriate. Rather, the Working Papers suggest only one kind of integrated multi-cloud shows a well-functioning market even though customers disagree. Microsoft understands that simple economic models suggest that, *all else equal*, lower switching costs and an ability to “mix and match” applications across providers are good market characteristics. Unfortunately, the Working Papers do not assess whether the current level is *right* or not, given the expected benefits and the very real (and well-documented) costs of switching and/or building integrated multi-cloud architectures.

2.2 In the face of abundant evidence of intense dynamic rivalry, the Working Papers rely far too much on simple measures of concentration (i.e. UK shares of supply from 2019-22)

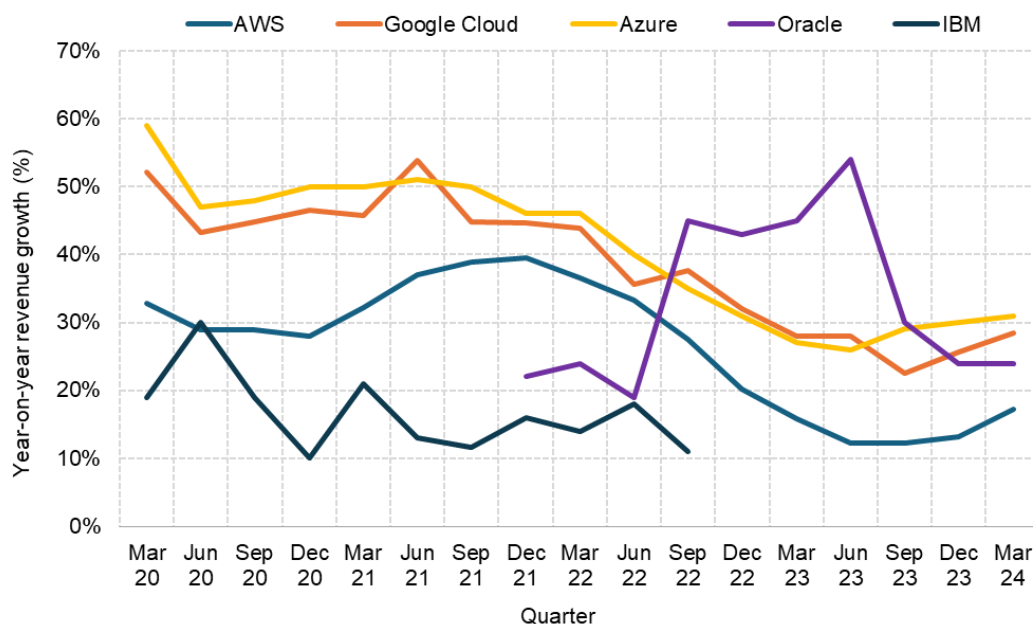
2.2.1 There is abundant evidence of intense rivalry and none of cosy or static behaviour between the top several competitors

- (18) In the face of competition, AWS has held on to a declining but still large share of UK cloud revenue (percentages in the 40’s in IaaS, in the 20’s in PaaS for 2022). Microsoft has – through trial and error, effort and expense – successfully made gradual inroads into eroding the long-term market leader’s share, with 2022 UK revenue share percentages in the 30’s in IaaS and in the 20’s in PaaS. Similarly, Google, although somewhat smaller in share, grew rapidly from 2020 to 2022 (and continues to grow). These results are the fruit of competition, not a lack of it. Competition is not soft, weak or cosy: Amazon, Microsoft and Google are “punching each other hard”.
- (19) This applies not only in the short term of winning customer workloads but looking forward on investments in data centres, chips and other competitive assets so that none of these players cedes strategic ground in the future to the other two. They also leverage differentiated competitive advantages derived variously from first mover status (Amazon) and from adjacent markets (Microsoft in on-premises enterprise IT, but also Google in data and the “ad stack”, and its vertical integration throughout the “AI stack”). Any advertiser who wants to understand how their ad campaigns perform on Google, can only do so by using GCP, giving Google an immediate advantage to leverage that “must have” workload into an additional GCP workload. A good sign of the lack of cosiness is that Amazon and Google are competing not only on the merits in the marketplace but are seeking to utilise the regulatory process for commercial gain, as well as undermine Microsoft’s intellectual property, under the guise of a consumer welfare concern (developed further in the Licensing Working Paper response to follow).

⁶ See paragraphs 86-90.

- (20) Nor is it accurate to describe cloud simply as a “three hyperscaler” race. They are joined by other players such as Oracle and IBM with smaller shares but that are themselves large tech players with distinct on-premises IT pedigrees and aggressive expansion plans. For instance, along with the three hyperscalers, Oracle’s cloud business has sustained double-digit revenue growth consistently since Q4 2021 (the period from when Oracle IaaS and PaaS cloud revenue data is available). Moreover, IBM’s cloud vertical also sustained double-digit revenue growth consistently between Q1 2020 and Q3 2022 (after which the revenue data becomes unavailable).⁷

Figure 1: Cloud provider revenue growth 2020-2023



Source: S&P Capital IQ and 10-K financials.⁸

- (21) So, while the two-firm share ratio (CR2) or HHI may show a concentrated market, and while there is academic literature supporting a general link between market concentration and market outcomes, such analysis of the competitive health of a market should always involve a case-by-case assessment, most of all true of in an in-depth MIR.

2.2.2 The surprising error in marginalising Google, a hyperscale rival worth \$2 trillion as a fringe player

- (22) Microsoft competes head-to-head with Google in cloud; what is important is not whether a rival is the closest but whether it is a close competitor.⁹ Google is a close competitor today; its capabilities and unique assets make its future prowess as a competitor even more impressive.
- (23) Most alarming is that the Working Papers actually take a step backward from Ofcom’s review when it comes to Google. Ofcom (consistent with the approach taken in third-party industry

⁷ Oracle’s financial quarters do not align with calendar quarters and the closest quarter is shown in the figure. For IBM, between Q4 2021 and Q3 2022, only hybrid cloud revenues are reported.

⁸ Note: Revenue growth has been calculated by computing quarterly year-on-year growth in the revenues of the listed firms’ cloud divisions. Oracle data is at the firm’s fiscal quarter level, which does not align with the calendar year; nearest fiscal quarter has been used in this case. Oracle does not report IaaS and PaaS cloud revenue data prior to Q4 2021 and IBM does not report the revenues of its cloud arm post Q3 2022.

⁹ Cf. CMA, *Merger Assessment Guidelines* (2021), paragraph. 4.8

studies to label Google as a “hyperscaler”) put GCP in the same class as Amazon and Microsoft (differences being ones of degree, not competitor class). Google has the capability to operate at hyperscale because it enjoys single-firm economics: GCP is not an independent company or regulated subsidiary and yet the CMA primarily treats GCP as if it were; indeed, the Working Papers appear to make a category mistake in characterising Google as more like Oracle, IBM and Aruba than the other hyperscalers for competitive assessment purposes.

- (24) This is surprising. Whatever barriers to expansion others may face, these do not apply, of all competitors, to Google. At over \$2 trillion market cap, it is worth more than Amazon. At over \$300 billion in annual revenue, and over \$85 billion in EBIT, Google has balance sheet firepower, evidenced by capex spending on a par with Microsoft and with a capital intensity above Amazon’s (see below). In addition, GCP has had a positive operating margin since Q1 2023, is #2 in terms of the number of cloud customers reflecting their strong position with cloud native firms, offers (by its own marketing) over 170 products comparable to AWS and Azure, and can already leverage its advantages in data, ads, and the AI stack (in which it is the uniquely vertically integrated player from chips and cloud through search to consumer mobile). As the CLWP notes, analyst reports “*recognise that [GCP] is expanding its capabilities across IaaS and PaaS and that in some areas it has been influencing the rest of the industry (e.g., Kubernetes)*”¹⁰ and key strengths include: “*its offerings in sovereign cloud, multi-/hybrid-cloud through Anthos and data insights and analytics and AI; its sales execution; its ability to attract cloud native applications; its ability to offer its services globally; its industry-focused solutions; and its storage and network capabilities.*”¹¹
- (25) The Working Papers also fail to recognise the unique benefits Google derives from its dominant position in digital advertising. The CMA is well aware from its work in digital advertising that Google has overwhelming strength in search advertising throughout the ad tech ecosystem. When campaigns are run using Google’s advertising services, the resulting data flows directly into Google Ads Data Hub (ADH).¹² ADH is built on infrastructure from Google Cloud and Google BigQuery. It gives advertisers access to the data needed to shape their campaign. Once there, it is stuck. Advertisers can input their other data into Google’s system for comparison, but they cannot export their Google data out to gain draw specific insights about their campaigns. For example, if an ecommerce retailer wants to understand the path their audience takes to conversion (purchase), they can bring their data from their CRM systems and marketing databases into BigQuery and combine it with the known Google data.¹³ They can then use BigQuery’s data analytics capabilities to draw insights.¹⁴ But they cannot do the same in another cloud. Similarly, Google limits which clouds can interoperate with its new privacy sandbox in Google Chrome and does not support Azure.¹⁵
- (26) The CMA documented these facts in the context of its Online Platforms and Digital Advertising Market Study.¹⁶ At the time, it was not focused on the competitive advantage that

¹⁰ CLWP, paragraph 2.117.

¹¹ CLWP, paragraph 2.118(a).

¹² See: <https://developers.google.com/ads-data-hub/marketers> which explains “with Ads Data Hub, you can upload your first-party data into BigQuery and join it with Google event-level ad campaign data”.

¹³ See WebFX, “Google Ads Data Hub: Everything a Marketer Needs to Know,” available [here](#).

¹⁴ See Google Cloud, “Ad Agencies Choose BigQuery to Drive Campaign Performance,” Sept. 9, 2021, available [here](#).

¹⁵ <https://developers.google.com/privacy-sandbox/relevance/aggregation-service/>.

¹⁶ See Final Report, Appendix O: Measurement Issues in Digital Advertising at paras. 122-124. There, the CMA was concerned about the impact of such restrictions on the measurement of digital advertising campaigns noting: “Using Google ADH is the only way to access all of Google’s post-campaign data and measurement analytics services in a single

such practices might impart on Google in the context of the market for cloud computing services. Not surprisingly however, they do confer a significant advantage. GCP's, largest industry segment is media companies who depend heavily on digital advertising for monetization.¹⁷

- (27) In addition, Google's strength in digital advertising places it in a position to offer ad credits on those exceedingly high-margin products to win GCP usage. That includes its YouTube, Google Search and advertising services. An example of this was documented by the Wall Street Journal in 2021, in relation to GCP winning a cloud contract with Univision Communications *"Google has secured one of its largest-ever cloud-computing contracts, Spanish-language broadcaster Univision Communications Inc., in a deal that shows how the tech giant is leveraging other parts of its operation to drive business to its cloud division...The Alphabet Inc. unit, which announced the deal Monday, beat out rival services by packing its offering with benefits across its YouTube video platform."*¹⁸
- (28) The Working Papers also fail to take account of Google's obvious advantages in AI and the future importance of those workloads. The promise of AI has led Nvidia to become the most valuable company in the world and Google, through its 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 that neither Amazon nor Microsoft enjoy, to secure a strong leg up in future competition. Judging by Google's continued multibillion dollar per quarter investment, it recognises these advantages. Google's own statements confirm this advantage: *"Google Cloud is bringing decades of AI research, innovation, and investment to the world with the launch of Generative AI support in Vertex AI and Generative AI App Builder"* said Ritu Jyoti, Group Vice President, Worldwide Artificial Intelligence (AI) and Automation Research, IDC. *"With this, Google Cloud is poised to enable a whole new generation of builders, innovators, developers and doers to harness the power of AI in novel ways."*¹⁹ The CEO of GCP describes Google cloud as *"one of the top enterprise companies in the world."*²⁰ The decades long Google search monopoly enables GCP services. Google explains that as Google *"build[s] generative AI into Search, we are drawing on years of insights and technology development to harness and deploy AI at scale. Google Cloud is how we provide these same insights, tools and platforms to help you innovate too, sharing the best of our AI directly."*²¹
- (29) The Jigsaw Research customer survey reflects Google' strength, noting that:
- (a) *"Google is seen as superior in particular to AWS in terms of its AI offering including for tasks such as machine learning and translation. It also offers AI powered Google Marketing for Google Cloud users."*²²
- (b) *"The main providers are seen as AWS, Microsoft and Google among participants. For some, this is the main or only consideration set in terms of who might even be on a*

environment, and ADH does not permit customers to export the information to any other measurement or ad tech partner they may prefer."

¹⁷ See HG Insights, "Google Cloud Market Share and Buyer Landscape in 2024," available [here](#).

¹⁸ See "Google Bundles Products to Land Univision Deal," Wall Street Journal, April 26, 2021, available [here](#). This is not a one-off example – see also: "Activision Blizzard And Google Enter Into Multi-Year Strategic Relationship To Power New Player Experiences", prnewswire.com, available [here](#), a deal with included YouTube as exclusive streaming partner.

¹⁹ <https://cloud.google.com/blog/products/ai-machine-learning/generative-ai-for-businesses-and-governments>

²⁰ [Cloud Next 2023: Sharing the best of our AI with Google Cloud \(blog.google\)](#)

²¹ Id.

²² Jigsaw Research– page 30.

shortlist of providers in the event of a review of the market or a switch, though most were not aiming to make any changes. They each have an excellent reputation, are seen to deliver a reliable service and offer a wide range of solutions that cover many needs."²³ [emphasis added]

(c) *"Many felt that services provided by these main players were '80% the same', but there were distinctions perceived e.g., some see AWS as the first mover and innovator, Microsoft as having superior IAM (identity and access management), and Google as having a strength in analytics (specifically, Google's analytics product BigQuery as having better performance)."*²⁴ [emphasis added]

(30) This issue is more than semantics: if an AEC were found and intervention were applied to regulate the conduct of some hyperscalers (Amazon and Microsoft) but not others (Google), this would distort inter-hyperscaler competition and damage arguably critical and aggressive industry-competitive dynamics.

(31) Google does not need the support of the CMA to compete and to artificially give it a leg up would simply allow it to build market share whilst competing less effectively than they do today. UK customers would suffer as competition softens. Uneven regulation would soften Microsoft's ability to win customers (back) from GCP and deprive them of the benefits of competition. Selective intervention would fail the most basic of remedial tests: first, do no harm.

2.3 Working Papers draw inaccurate inferences from "textbook" indicators of competitive outcomes

2.3.1 Profitability

The utility-inspired ROCE analysis fails to reflect dynamic competition – especially on AI investment

(32) The CMA's Guidelines explain that "[*high price-cost margins ... [or] profitability*]" are only indicators and do "*not provide conclusive evidence that the market could be more competitive*".²⁵ However, aside from the problematic comparative assessment of Google on the one hand and Microsoft on the other (see above), and problems with how the analysis has been done (see Section 2.3.2), Microsoft's fundamental concern is that the CMA's textbook ROCE vs. WACC profitability analysis, well-suited to mature markets such as energy, water, and telecoms, is off the mark for a market where: (a) operational excellence and scale requires significant investments multiple years before reaching positive returns (Amazon: 2015; Microsoft: 2016, Google: 2023), and (b) more importantly has manifestly not reached an equilibrium steady-state but is in a capex spending race of unprecedented proportions. In short, it is a "good thing" for competition (and in turn customers), not a "bad thing", that each of Google (\$32.3 billion capex at capital intensity of 10.5%), Microsoft (\$35.2 billion capex at capital intensity of 15.5%) and Amazon (\$52.7 billion capex at capital intensity of 9.2%) are generating positive returns which can justify the kinds of capital expenditure they are undertaking for the future.

(33) The CLWP addresses falling or flat ROCE by arguing that these recent trends are the result of "*increased levels of investment in cloud infrastructure ... [to] [support] the development of*

²³ Jigsaw Research - paragraph 1.3.6. and 3.4.4. See also paragraph 3.4.

²⁴ Jigsaw Research - paragraph 1.3.8. and 3.4.6.

²⁵ CMA CC3 Guidelines, paragraph 126

AI services”²⁶ which is “not ... a result of competitive forces”.²⁷ It is not clear to Microsoft on what basis this can be correct. Indeed, Gartner and other third-party industry reports have long acknowledged that the cloud market will continue to grow, with AI being one workload. Microsoft entirely missed the mobile revolution which it ceded to and is now dominated by Google (Android) and Apple (iOS) ecosystems. Google is also vertically integrated into AI chips (Tensor Processing Units), cloud for AI, Deep Mind foundation model development, AI-powered Google search, Google Gemini AI assistants, distribution to consumers on Android and via Google Play Store, and AI-powered services on Android and across its applications and services. The notion that Microsoft’s capex race to serve AI – and not get left in the dust by rivals and in particular Google – has nothing to do with “competitive forces” is clearly wrong.

2.3.2 Microsoft’s UK share of supply

Phase 1 indicators of Microsoft market power are relied on without any further critical evaluation or in context

- (34) The above gaps make the assessment of Microsoft’s market power particularly disappointing because the CMA’s MIR Guidelines define “significant” market power with respect to the durability of bad market **outcomes** against a “competitive level” counterfactual:

*“endur[ing] over time and ... the ability to maintain prices above the competitive level, or restrict output or quality [or innovation] below competitive levels”.*²⁸

In this case, as noted, that counterfactual is a “well-functioning” cloud market. The market power question, by contrast, is not answered by high **shares of supply**:

*“a large market share does not always indicate that competition within the market is weak. It may simply indicate that the firm(s) possessing it is capable and relatively efficient, having low costs, an attractive product, or both.”*²⁹

- (35) On this point, Ofcom devoted 12 months to the conclusion that it “suspected” that Amazon and Microsoft hold “a degree of market power” relying on indicators: retrospective metrics of high shares (2020-22), barriers to expansion, and supplier profitability. The emerging thinking to date does not fundamentally advance the ball. It uses essentially the same mechanical checklist to conclude that there are “indicators” that Amazon and Microsoft hold “significant market power”.
- (36) Instead of focusing on a single share of supply metric, a better indicator of future relative market position would be UK cloud providers’ share of global capex. On this basis, the 2022 shares are: Amazon at 41% (\$52.7 billion), Microsoft at 28% (\$35.2 billion), Google at 25% (\$32.3 billion), Oracle at 5% (\$6.6 billion) and IBM with 1% (\$1.2 billion):

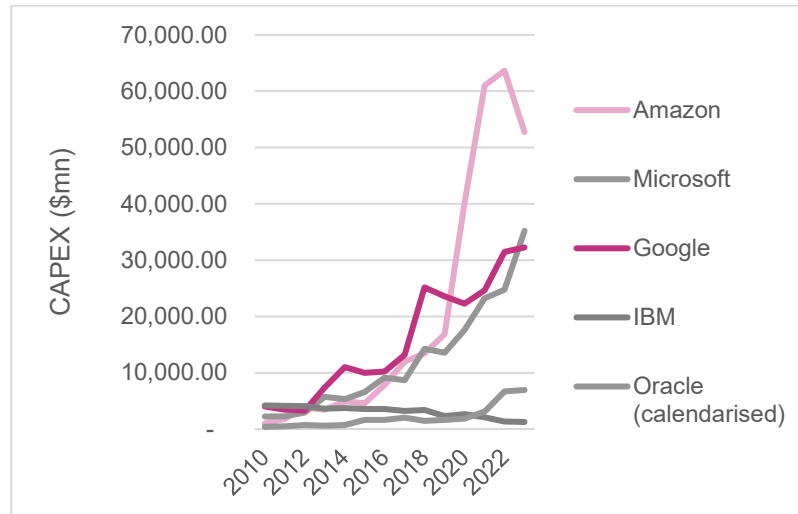
²⁶ CLWP, paragraph 6.41.

²⁷ CLWP, paragraph 6.43.

²⁸ CMA MIR Guidelines CC3, paragraph 9.

²⁹ CMA MIR Guidelines CC3, paragraph 190.

Figure 2 – 2010-2022 Capex Spend by Cloud Provider



Source: S&P Capital IQ³⁰

- (37) It remains the case, however, that any indicator is just an indicator – and not evidence of poor customer outcomes that could underpin an AEC finding and intervention.

2.3.3 Absolute switching rates and multi-cloud adoption

- (38) No IT market ever is or has been a textbook widget market with perfect competition, high switching with the only variable being price (competed down to marginal cost). Though much more fluid than on-premises IT, cloud is no exception. As Gartner has explained, switching costs or “*vendor lock-in*” is baked into the essence of differentiated IT markets, and the only question is whether customers address portability up-front (prior to vendor selection and customer investment) or down the track (once they commit to an environment) and where in the stack the lock in occurs.³¹ For example, the CMA cites abstraction layers as being a positive development, but of course every customer that develops to that abstraction layer is, to a certain extent, locked into that abstraction layer. It may make switching datacentres on balance easier, but it creates its own complexity and level of lock-in should the customer want to switch technologies. It is good for customer welfare that IT services markets are differentiated not commodity goods markets and that customers be given the choice to decide what services they want to build and deploy and what “multi-cloud” solutions make the most sense to them. The fact that customers do not use services like Kubernetes or Terraform as much as the CMA believes they “should” does not reflect limitations in the services themselves, but rather the fact that different customers will make different choices about how to optimise their cloud architecture.
- (39) Given the benefits of differentiation, customers devote time, effort and expense to develop familiarity (and customisation) of one differentiated IT environment over another. Switching involves inevitable hassle. In interpreting the customer evidence, this response will point out truths that are self-evident in IT. Customers will likely:

³⁰ Notes: Calendar year CAPEX calculated by summing over the firms’ quarterly figures. Annual CAPEX figures have been used for years (2013-2015) where quarterly CAPEX was unavailable for Google and IBM. Oracle’s annual CAPEX is calculated by weighting the quarterly CAPEX by the number of months in the reported quarter from the calendar year.

³¹ Relative to switching from on premises to different on premises or from on premises to the cloud. See Microsoft’s response to Ofcom’s Interim Report – paragraph. 25.

- (a) Not “switch” cloud vendor unless the benefits outweigh the costs: “if it ain’t broke” they won’t try to fix it. A me-too offering will not be worth it, even for relatively small price differentials. An innovative, differentiated offering with serious new value might be.
 - (b) Not “multi-cloud” if their cloud spending is small and their cloud workloads are few: high-earners may have multiple investment pots and bank accounts; low-net worth individuals will probably have accounts with one bank – a paternalistic sense that they should or could be able to “multi-bank” very much as it makes no sense for them. Of course, here the MIR involves not vulnerable consumers (cf. *Energy, Banking, Funerals*) but instead corporate enterprises, large and medium-to-small.
- (40) With that background, the CMA’s analysis shows that multi-cloud prevalence weighted by revenue (capturing the behaviour of bigger spending customers) is approximately 5x that of unweighted prevalence (capturing the behaviour of a larger number of smaller-spending customers, some of whom will only have one workload (multi-cloud is largely meaningless) or a few (not as yet relevant): 34.4% weighted vs. 7.1% unweighted in 2022. That weighted figure has increased from 31.3% to 34.4% (3.1 percentage points or around 10% growth from 2020)³² which the CMA recognises is a likely underestimate.³³ If weighted spend were not higher than unweighted spend there might be cause for concern that even as customers increase cloud spending, with more workloads, there is no concomitant increase in multi-cloud. This is contrary to the data which shows a 5-fold increase.
- (41) **Absent a well-founded counterfactual that in a well-functioning market this “should” be a 7-, 8- or 10-fold increase, there is no basis to conclude that multi-cloud adoption is “too low” and signals weak competition or bad outcomes.**
- (42) This is reinforced by Microsoft’s evidence that individually-negotiated discounts [X], showing the presence of consistent competitive pressure even absent switching of the workloads concerned.³⁴

2.3.4 Cogent customer concerns and dissatisfaction

No weight of cogent customer complaints that the market is responding weakly to open-source and portability demand signals

- (43) The weight of the Jigsaw customer survey evidence does not, as a whole, “*attribute an unmet desire to supplier conduct*” (which most lends itself to the case for government intervention) versus customer inertia because inherent costs of IT switching often outbalances the gains. In other words, there is a lack of evidence that customers would *like to do more switching or multi-cloud* but a supplier-driven *inhibitor prevents them*. As noted by the survey:
- (a) “*Switching cloud providers is seen as the equivalent of moving other kinds of infrastructure, such as ‘moving house’ or moving a business from one country to another. It is not something to undertake lightly or consider at all unless it leads to significant business benefits long term that override the perceived cost and risk of changing.*”³⁵ “*Where people had switched, either all or some of their infrastructure, this was driven by an anticipated reduction of cost or clear increase in performance. In some cases switching was a requirement following a merger or acquisition to rationalise*

³² CLWP, paragraph 3.73 and Table 3.1.

³³ CLWP, paragraph 3.76.

³⁴ [X]

³⁵ Jigsaw Research – paragraph. 3.6.3.

providers.”³⁶ “Overall, for those who had switched, the process not only brought cost and operational risk, but took IT staff away from core work and typically ended up being more challenging and time consuming than anticipated”.³⁷ “Another aspect of concerns about potential disruptions to the IT service relates to the resources a switch or multi-cloud approach would take away from the IT department. This is not a direct technical challenge but rather an exacerbating factor, making a switch or multi-cloud attempt more difficult.”³⁸ “Especially in the case of switching cloud provider, the human knowledge that a company’s engineering team has built up over many years working within a specific cloud environment may be seen as ‘stranded assets’ by decision makers – i.e. the company may have invested a lot of time and money into building a body of knowledge with one particular cloud provider which makes it even harder to justify a switch into another environment where a considerable amount of that knowledge is no longer applicable.”³⁹ “Many research participants expressed concern that switching provider or adopting a multi-cloud architecture would absorb too much of their software development resources.”⁴⁰

- (44) Alternatively, it could be the case that where there is customer demand signalling for “open source” and “portability”, the suppliers in the market respond to those signals with solutions, such as supporting Kubernetes. This is indeed what can be observed, particularly in relation to Azure. This is explored further in our response to the technical barriers Working Paper.

2.3.5 Failure to capture the implications of the AI wave for its ‘engine room’ – cloud

- (45) Constant innovation in technology creates competitive pressure. No one hyperscaler can capture that innovation as much of that innovation happens outside of the hyperscalers and often in the “commons” (notably, open source).
- (46) As noted above, with the AI wave, powered by the cloud, the hyperscalers cannot sit still or they will become the obsolescent if not obsolete cloud infrastructure hosts for the AI of tomorrow. That is why hyperscalers are investing massive amounts in growing their capacity, investing in Deep Mind foundation models, supporting tooling and applications.
- (47) Meanwhile, Nvidia, now the world’s most valuable company due its pre-eminence in GPU chips for AI applications, is playing a significant role in promoting competition through its strategic allocation of GPUs and sponsoring expansion of players such as CoreWeave. CoreWeave – a relative unknown when the MIR commenced – announced a \$1.1 billion investment, placing its valuation at \$19 billion. Equally, it recently announced expansion into Europe with the UK as its European headquarters and is preparing for an IPO.
- (48) Moreover, CoreWeave is also a demonstration of the fact that because “range of services” and scale is not as relevant for all uses of AI infrastructure, generative AI has created opportunity for such new entrants.
- (49) The CLWP focuses on the potential impact of AI on how competition works in cloud services in terms of: (i) the role of AI developers as customers, (ii) the importance to cloud suppliers of providing accelerated compute for AI, and (iii) competition to supply accelerated compute.

³⁶ Jigsaw Research – paragraph. 3.6.7.

³⁷ Jigsaw Research – paragraph. 3.6.9.

³⁸ Jigsaw Research – paragraph. 4.2.12.

³⁹ Jigsaw Research – paragraph. 4.3.5.

⁴⁰ Jigsaw Research – paragraph. 4.3.9.

The analysis is however relatively narrow as if AI were just another customer workload class within cloud, alongside retail, financial services, etc.

- (50) The fact that AI is cloud-powered in a backdrop of a competitive arms race for AI development suggests the cloud market will look even more competitive than the well-functioning market it is today. Not less so.

2.4 Conclusion

- (51) In summary, the AEC test is fundamentally a present-tense test: assessing features that prevent, restrict or distort competition and not features that may, at some point, if certain hypothesised pre-conditions are met, cause such effects. As set out above, there is no compelling evidence to demonstrate that the AEC test is met or likely to be met in the coming years. If concerns about the future evolution of the market do emerge at some point down the line, beyond the appropriate timeframe for the current MIR assessment, the CMA will have the capacity to monitor the market through the newly established Digital Markets Unit and intervene, as needed..

3 The CMA's concerns on CSAs are merely theoretical as CSAs account for only a part of a customer's cloud demand, they result in lower prices and certainty for UK customers, encourage customer switching / multi-cloud and support investment in UK cloud infrastructure

3.1 Introduction

- (52) Microsoft fundamentally disagrees with the CMA's analysis of CSAs and their effect on competition in the cloud services market. Based on all available evidence gathered by the CMA (and consistent with Microsoft's experience of market dynamics), the CMA's emerging view that CSAs harm competition is purely theoretical.
- (53) CSAs are integral to the competitive process in the cloud market – it is how cloud providers compete head-to-head. CSAs result in lower prices for customers and the commitments on which these discounts are based are what underwrite this price investment. This price competition in turn encourages customer switching and multi-homing and, at a broader level, supports investment in public cloud infrastructure.
- (54) The theoretical issues the CMA is currently exploring are “textbook” concerns that are not borne out in light of the specific features of the cloud computing industry or the structure of CSAs. In reality, committed spend agreements cover only part of a customer's overall cloud demand – [X] – and customers have substantial excess demand and flexibility to contract with other cloud providers. Given the lack of evidence of actual customer harm, Microsoft submits that the threshold for an AEC has not been reached. Any intervention restricting the ability for some, or all, cloud providers to offer CSAs would distort the competitive process and lead to negative impacts on users of cloud services based in the UK. In particular, the suggestion to exempt CGP from CSA remedies would simply insulate them from competition, artificially increasing their market share, leading to worse outcomes for UK consumers.

3.2 CSAs are pro-competitive – they facilitate customer switching and investment in UK cloud infrastructure

- (55) CSAs benefit customers, providing lower prices (through discounts), facilitating migration support (through cloud credits and migration investments) and encouraging cloud providers'

long-term investment in cloud infrastructure services (through the use of long-term commitments).

- (56) Microsoft and its customers primarily rely on CSAs to deliver large scale, transformational IT projects. Deployment of these workloads can take several years (or more) and the commitment volumes and contract duration are designed to accommodate a “hockey stick” acceleration of use (i.e. a slow ramp phase followed by a sharp increase in consumption). CSA and committed consumption volumes provide customers (and Microsoft) with a greater certainty of revenue / costs for the full deployment period, underpinning significant upfront investment it makes in its customers through discounts, credits and other fund incentives.
- (57) CSAs are typically utilised by larger, more sophisticated businesses.⁴¹ These customers generally hold CSAs with two (or more) cloud providers and their high levels of cloud spend give them an opportunity to negotiate sizeable discounts for each cloud. Customers therefore make a case-by-case decision on where to deploy workloads and do so on a range of factors (such as product and service quality, geographic coverage, integration with other products etc.). The existence of CSAs is, thus, not a determining factor and does not prevent customers from switching or developing multi-cloud where a supplier offers a superior product.
- (58) Against this backdrop, Microsoft is concerned that the CMA’s emerging thinking fails to fully reflect the customer benefits and pro-competitive features of CSAs:
- (a) **CSAs are a key element of competition in the market, are offered by cloud providers of different sizes and allow customers to receive better prices than they otherwise would have.** CSAs are not exclusively offered by AWS or Microsoft. They are a feature of competition among all cloud providers, including hyperscalers and smaller providers, who compete aggressively for customers on a workload-by-workload basis. CSAs are a key aspect of price competition between cloud providers and Microsoft typically offers moderate discounts of between [X]%⁴² (and an average discount of [X]%),⁴³ based on [X]. Discounts are particularly valuable to customers in the early stages of deployment, where set up (switching or migration) costs are high and cloud consumption is low (and the benefit they derive from the services is also minimal).
- (b) **CSAs give customers and cloud providers greater certainty of pricing and demand.** Fixed duration consumption agreements provide customers with an assurance that the agreed price / discount will persist over the life of their cloud contract from the start to the end, providing greater visibility and control of their cloud spend.⁴⁴ Microsoft typically enters into [X] (which underpin the necessary return on investment to offer upfront investment); however, ultimately, the length is driven by its customers and their specific needs. For example, Microsoft entered into [X]. Other long-term CSAs include [X]. Additionally, fixed-duration contracts allow cloud providers to better forecast revenues that will be available for reinvestment as well as demand for capacity over the medium / long term, which increases investment certainty and innovation and lowers costs.

⁴¹ As the CMA notes in paragraph 2.11 and paragraph 2.14, [X].

⁴² [X]

⁴³ [X]

⁴⁴ Jigsaw Research, paragraph 6.2.3.

- (c) **CSAs enable cloud providers to invest in better support and services for customers than they otherwise would have.** Commitment agreements allow cloud providers to make significant upfront investment during the early stages of deployment. Cloud providers use this to fund the costs of technical support and solution development, which customers would have otherwise incurred themselves during the deployment process.⁴⁵ Microsoft’s End-Customer Investment Funds (“**ECIF**”) are provided to assist customers with project-specific needs [§<]⁴⁶ [§<].⁴⁷ Microsoft can make these investments confidently because of the guaranteed revenue of the commitment. [§<].
- (d) **CSAs enhance switching between cloud providers.** Commitment agreements allow cloud providers to subsidise and overcome inherent costs involved with switching from rivals upfront. Microsoft offers the Azure Credit Offer (“**ACO**”)⁴⁸ to reduce “dual run costs” for customers switching to Azure. As noted above, ACOs can be a significant part of the value provided to customers under a CSA:⁴⁹ [§<]. As with ECIFs, [§<], this allows it to make these investments based on guaranteed revenue. All things being equal, Microsoft would therefore expect less customer switching if CSAs are prohibited or limited as it would be difficult for customers to persuade cloud providers to pay upfront for migration costs.
- (e) **CSAs provide customers with more flexibility to develop terms that are tailored to their needs.** CSAs are heavily negotiated and designed to reflect the individualised circumstances of each cloud migration based on projected demand, length of deployment and the nature of the workload / industry. As mentioned in the CMA’s CSA Working Paper,⁵⁰ Microsoft works with customers to set and meet their commitment levels and customers not utilising their commitment is the exception rather than the norm. For example, [§<].⁵¹ [§<].⁵²

(59) The *commitment* in the CSA is not an ancillary part of this offering. On the contrary, the price investments (i.e. discounts) made following these intense negotiations are *only possible* because of the commitments that come with them. It is akin to bulk purchasing and the procompetitive discounting that commonly comes with larger purchases. The kinds of customers and workloads with CSAs are those with complex needs and typically extended timescales to both shift and scale workloads in the cloud. Commitments enable the cloud provider to invest in meeting those needs, even where profitable revenues are not expected to flow back to the cloud provider months, or years, later. If the commitment was to be removed, cloud providers’ ability to make these investments would suffer and customers would bear more of the costs of shifting, configuring and optimising workloads (both from on-prem and between clouds). In summary, CSAs are part and parcel of the way cloud

⁴⁵ For example, Microsoft provides funds through the ECIF programme, which can cover pre-migration activities (e.g. training, solution design, ISV tool development) and post-migration support (e.g. remediation of technical issues, application development, iterative testing and architecture design). These funds provide an incentive to customers to increase their investment in cloud computing.

⁴⁶ [§<]

⁴⁷ [§<]

⁴⁸ ACOs are free credits customers can use against Azure consumption. They are often provided within a set timeframe and can also be provided contingent on certain thresholds of spend being met.

⁴⁹ [§<]

⁵⁰ CSA Working Paper, paragraph 1.30

⁵¹ [§<]

⁵² [§<]

providers compete aggressively on price and non-price terms in the cloud services market. Customers use the negotiation process to their advantage to secure lower prices and greater investments in their infrastructure. To the extent that CSA remedies are limited to Azure and AWS (and exempt GCP), there will be less competitive pressure on cloud providers to offer better terms – artificially increasing their market share at the expense of UK cloud customers. Rather than harming competition, CSAs are therefore key to healthy competition in this market and enable switching between cloud providers that are offering new and innovative services.

3.3 The CMA’s analysis does not reflect either the empirical evidence gathered or the structure and competitive dynamics of the cloud services market.

3.3.1 Concepts of “sticky” and “contestable” demand don’t make sense given the dynamic and differentiated nature of the cloud market

- (60) Microsoft submits that the CMA’s framework for assessing sticky demand is inappropriate in the context of a dynamic cloud market, where customers’ preferred products are constantly changing, and cloud providers cannot determine with any precision for any customer what proportion of cloud spend is “sticky” and what portion might be “contestable”. This is not like (by way of example) the computer chip (CPUs) market, where a chip supplier has visibility over how many computing devices a computer OEM ships and thus can determine how much of its customers’ demand is contested. A cloud provider is not able to know how much more a customer may be looking to spend in the cloud, nor can the cloud provider know how willing the customer is to switch workloads to other cloud providers.
- (61) Assessing the cloud market through a lens of bright-line sticky and contestable demand provides a static view of the market. This framework can be useful when considering stable and mature industries that are not subject to rapid change, where providers can clearly identify (and lock in) sticky demand, but it does not make sense given the dynamic nature of the cloud market.
- (62) Intense competition between cloud providers to develop differentiated services has led to an increase in the range / quality of products available to customers and a decrease in prices over time.⁵³ A service or product that seems sticky today is very likely to face fierce competition from competing services when a contract is up for renewal. As the Jigsaw Research makes clear, users are willing to switch to competing services where there is value in doing so – including where this will reduce costs or increase the performance of their cloud environment.⁵⁴ In assessing the impact of CSAs on the competitive process, the CMA needs to consider the following:
- (a) **Constant and significant growth in contestable demand.** As the CMA notes, worldwide IaaS and PaaS spending is forecast to increase c.50% and c.45% respectively between 2022 and 2024.⁵⁵ The vast majority of this demand growth is for new cloud services, either existing workloads migrating to the cloud or new workloads generated by cloud native applications, is contestable and cloud providers compete aggressively to win this new demand.⁵⁶ All things being equal, this growth means that

⁵³ [X]

⁵⁴ Jigsaw Research, paragraph 1.3.14.

⁵⁵ Competitive Dynamics Working Paper, paragraph 2.12.

⁵⁶ For example, in the Jigsaw Research, paragraph 6.1.18 “*participants from start-up businesses described the cloud providers as quite active in trying to win their business.*”.

the alleged “contestable” demand should be significantly larger than any alleged “sticky” demand at any given time.

- (b) **Competition for new CSAs.** The CMA ignores the intense competition between cloud providers that takes place for first-time CSAs. As the CMA’s research makes clear, currently only a small number of customers have CSAs ([redacted] % of total Azure customers in 2022) and cloud providers use CSAs to aggressively compete for customers who do not already benefit from these arrangements.⁵⁷ CSAs are typically used to support large and complex workloads, so in these circumstances there is likely to be relatively little existing spend (which may or may not be “sticky”) compared with the future spend (which would be new “contestable” demand) under the new agreement.
- (c) **Technological change.** In addition to the migrated and new workloads, technological innovation is a key driver for switching and multi-cloud and customers report using APIs and cloud-agnostic applications to enable multi-cloud and switching.⁵⁸ Any assessment on the existence, or proportion, of sticky / contestable demand and customer incentives to switch requires an understanding of how technology changes over time. Again, this requires a dynamic view of the market that is not captured by a framework more suited to mature commoditised markets.
- (d) **Real-world evidence of pricing.** Assuming that Microsoft did have a high share of sticky demand and little incremental commitments to fight for, it follows that discounts would be higher for new CSAs (where all the demand is contestable) relative to renewals (where, according to the CMA’s theory, customers are locked in and all demand is thus sticky) and prices would increase over time. However, the empirical evidence submitted by Microsoft shows that the opposite is true: [redacted].⁵⁹ As shown by Figure 3 below, competition is as intense when customers renew a CSA as when they entered into their first CSA.

Figure 3: [redacted]

[redacted]

Source: [redacted].⁶⁰

3.3.2 CMA’s analysis is theoretical, not supported by empirical evidence and ignores how customer contracts are structured in reality.

- (63) As set out above, the CMA’s analysis of sticky / contestable demand is largely theoretical. It is also not supported by robust empirical evidence from customer contract data or actual customer usage.
- (64) First, no robust evidence is offered to support a finding that sticky demand would be high in practice. Instead, the CMA has made a high-level assumption that Microsoft’s (and AWS’) share of sticky demand must be substantial due to: (a) a lack of alternative suppliers, which does not reflect the plurality of cloud providers in the UK market (including well-resourced competitors AWS, GCP, OCI and IBM); (b) high barriers to switching / multi-cloud, which are lower in the cloud industry compared to any other IT infrastructure and decreasing over

⁵⁷ UIS, paragraph 38

⁵⁸ UIS, paragraph 38

⁵⁹ [redacted]

⁶⁰ [redacted]

time due to technological change;⁶¹ and (c) customer's feedback on how much demand they would be *willing* (not *able*) to switch to another provider, which varied considerably between respondents (as noted below).

- (65) Second, if Microsoft were able to accurately forecast a customer's proportion of sticky demand, it would be expected to develop a consumption commitment that captures all of a customer's sticky demand. In fact, the CMA's analysis shows that [redacted]. Importantly, this analysis does not factor multi-cloud arrangements, the presence of other CSAs, or changes in behaviour based on the size of commitment (i.e. that customers may spend only so much as to qualify for their discount).⁶² Taken together, committed consumption would represent (at most) a modest proportion of customers' overall cloud spend across providers and thus would have a much more limited influence on decisions to allocate new and existing workloads between cloud providers.
- (66) Third, the CMA's analysis does not reflect the highly bespoke nature of CSAs. As the Jigsaw Research recognises,⁶³ CSAs are not a "one size fits all" agreement and the terms can vary significantly between customers on volume and length and can include multiple types of discounts (such as broad-based commitment discounts, specific SKU discounts, Committed Use Discounts, Reserved Instances and Savings Plans) and customer investments (such as ACO and ECIF). The CMA's analysis is centred around theoretical discrete units of demand for a "representative customer" without controlling for these variations.⁶⁴ This is not applicable to how demand in the cloud market works in reality.
- (67) Finally, the CMA's empirical evidence on incremental discounts and commitments is flawed in that the CMA considers them separately when they are related. As indicated from [redacted],⁶⁵ [redacted] the two cannot be analysed separately as the CMA does.⁶⁶ [redacted].

Figure 4: [redacted]

[redacted]

Source: [redacted].⁶⁷

3.3.3 CMA's research shows that CSAs are not a material constraint on customer behaviour

- (68) Feedback regarding the impact of CSAs on the allocation of new workloads, switching and multi-cloud contradicts the CMA's theory that CSAs are a meaningful constraint on switching and multi-cloud.
- (69) In relation to the allocation of new workloads, contrary to the CMA's theory that customers are likely to place these with incumbent cloud providers in order to benefit from existing discounts, only a handful of customers "*explicitly gave the reasons for choosing either Microsoft or AWS as being the presence of a CSD or commitment*". Instead, customers indicate they choose Microsoft for a diverse range of reasons, including quality of service, product range and geographic reach and not just purely price. In fact, Jigsaw Research had

⁶¹ CSA Working Paper, paragraph 2.102

⁶² CSA Working Paper, see paragraph 2.75

⁶³ Jigsaw Research, paragraph 6.1.6.

⁶⁴ CSA Working Paper, paragraphs 1.19 and 120(c).

⁶⁵ For example, [redacted].

⁶⁶ CSA Working paper, paragraph 2.92. The CMA finds that [redacted].

⁶⁷ [redacted]

to proactively ask many customers as regards discounts in CSAs as it was not brought up by all participants themselves.⁶⁸

- (70) In relation to switching existing workloads, some customers indicated that CSAs influence their propensity to stay with their current provider *to some extent*.⁶⁹ However, the CMA's approach to gathering customer evidence on the share of sticky demand is flawed. The CMA asked customers what share of demand they are *willing* to switch, not what they are *able* to switch. This approach is incorrect, as what matters in the CMA's concept of sticky demand is the extent to which the customer "*can exercise effective choice over that demand [which] is limited by factors such as lack of suitable alternatives or barriers to switching*."⁷⁰
- (71) This distinction is evident in the CMA's survey results. It is clear from the survey that customers do not consider themselves to be locked in by CSAs and they have the ability to switch cloud providers who offer a better service. As the CMA notes in paragraph 2.61: "**Many customers... said it was possible to switch away at least some services from their main provider to alternative providers**".
- (72) Customers may be unwilling to switch a certain share of their demand for reasons other than the demand being sticky; for example, they may be happy with the level of discounts (and other non-financial benefits) from cloud providers. However, customer responses to Jigsaw Research acknowledge a willingness to switch if there are incentives. This can be due to price factors (such as the provision of discounts) or non-price factors (access to innovative products, greater performance or better integration with other services). For instance, the CMA notes that "*customers suggested **decisions on allocation of workloads are on a per service basis and consider a range of factors including the quality of the services, the business needs and current capabilities***".⁷¹ Additionally, smaller, lower spend cloud users also report switching cloud providers purely based on the financial incentives offered to them (e.g. credits and other monetary incentives).⁷²
- (73) The CMA's assessment overstates the proportion of sticky demand contracted by Microsoft (and AWS), which is based on the unevidenced assumptions that higher barriers to switching and multi-cloud mean a higher share of sticky demand. However, this conclusion is directly contradicted by feedback from some Microsoft customers who indicated that they were "*unwilling to switch **25% of [their] demand***" or "*were willing to switch all or **virtually all of their Microsoft demand***".⁷³ Similarly, customers are also unlikely to be able to accurately quantify the share of their demand which other cloud providers are *unable* to service. As some customers may not be presently considering switching (or may not have considered it in the past), they may be overestimating the share of their demand that only Microsoft can provide. This is supported by the fact that most customers could not in fact specify a proportion of spend for services they would not be willing to switch.⁷⁴
- (74) Microsoft also wishes to highlight its concerns about the methodology of the survey, which may affect the results of the survey. Notably: (1) public sector customers are also heavy users and beneficiaries of CSAs (see the UK Government's committed spend discount schemes) yet have been excluded from the Jigsaw Research; and (2) in some cases,

⁶⁸ Jigsaw Research, paragraph 6.1.1

⁶⁹ CSA Working Paper, paragraphs 2.4.1 and 6.3.5.

⁷⁰ CSA Working paper, paragraph 1.10

⁷¹ CSA Working Paper, paragraph 2.61

⁷² Jigsaw Research, paragraph 1.4.25 and CMA paragraph 2.40

⁷³ CSA Working Paper, paragraph 2.60

⁷⁴ CSA Working paper, paragraph 2.62

moderators in the Jigsaw Research had to proactively ask customers as regards discounts in CSAs.⁷⁵ These aspects of the survey may lead to a bias towards expressing concerns about / negative effects of CSAs.

3.4 The CMA's analysis does not support the AEC finding and all of the CMA's proposed remedies come with significant downsides

- (75) Given the significant benefits provided to customers through CSAs and the harm that would follow from limiting or removing these benefits, Microsoft considers that the starting threshold for CMA intervention should be high.
- (76) Intervention into a well-functioning market is likely to have significant consequences, including increasing prices for cloud users (vis-à-vis their global competitors) and reducing investment in cloud infrastructure services from both Microsoft (including in UK data centres and cloud applications) and customers (in productivity-enhancing cloud migrations).
- (77) The CMA has invited submissions on whether the benefits of CSAs could be achieved using alternative discount structures. The removal of commitments from the arrangements (i.e. permitting traditional volume tier discounts only) will simply remove the certainty of revenue streams and demand that cloud providers rely upon to fund substantial upfront investments in cloud deployments and workload switching. It would also remove the ability for cloud providers to offer tailored agreements to meet their customers' needs.⁷⁶ Ultimately Microsoft does not see a meaningful way of amending the discount structures while achieving these benefits. Microsoft considers that the proposed remedies are likely to lead to increased prices for UK customers and make it harder for cloud providers to compete for the next workload.
- (78) The CMA's current proposal is also unworkable as it disadvantages those customers who are not able to move their deployments to another geography. Given the global nature of cloud services, customers who might otherwise deploy in the UK can and do contract in the U.S. or other jurisdictions (at their request). Intervention into global markets that reduces flexibility, or makes discounting structures less beneficial to UK customers, is likely to lead to a change in behaviour from global customers who might otherwise contract in the UK. We anticipate these customers will take advantage of the regulatory arbitrage by contracting in a foreign jurisdiction, providing them with a cost advantage over domestic rivals. In both cases, this would create uneven competition and lead to worse outcomes for UK customers.
- (79) More generally, as set out in the introduction, we would strongly oppose the suggestion that Microsoft and AWS would be subject to these restrictions while GCP is exempt, despite its significant resources and market power across the digital economy. GCP is, in fact, a well-resourced competitor on a par with AWS and Azure and treating it differently would ignore their competitive advantages in adjacent markets and, critically, their vertically integrated AI capabilities. Currently, all cloud providers compete aggressively head to head, including through CSAs which offer customers the most competitive prices and terms for new workloads. Exempting GCP from remedies on discounts would simply insulate it from competition and soften their incentives to compete on price with Microsoft and AWS. While this might allow them to win new workloads (including some workloads they would have likely won anyway) and artificially build some market share, they would be able to do so by

⁷⁵ Jigsaw Research, paragraph 6.1.1

⁷⁶ CSA Working Paper, paragraph 2.113(d)

offering worse terms to UK customers. This would clearly be distortive and simply involve the CMA picking winners to the detriment of UK customers.

4 The CMA's theory of harm on egress fees is hypothetical and does not meet the AEC standard as Microsoft's egress fees do not prevent customers from switching / multi-cloud

4.1 Introduction

- (80) The CMA is considering whether to intervene in the cloud market to potentially set prices or force cloud providers to offer a service – which has real and significant supply costs – for free. Such an intervention can only be based on real-world evidence that egress fees are preventing, restricting or distorting competition in the UK cloud market and that intervention would address this distortion of competition (or any resultant detrimental effect on customers).
- (81) Microsoft's experience is that egress fees are not an area of concern for customers either when joining Azure or when moving workloads off Azure. Consistent with this experience, while there may be inherent challenges in switching between cloud providers, the CMA identifies no evidence that egress fees have a meaningful impact on consumer behaviour. To the extent they have any impact on competitive behaviour, cloud providers compete aggressively to minimise these (through specific discounts, offers and credits to overcome these frictions).
- (82) Instead, the CMA's UIS and Egress Fees Working Paper set out a hypothetical theory of harm that egress fees may create barriers to switching / multi-cloud on the basis that hypothetical "one-off" switching costs could make it more difficult for customers to switch cloud providers or use multiple clouds.⁷⁷ Microsoft has made egress free for customers switching away from Azure and, insofar as multi-cloud is concerned, the CMA's theory is in turn based on a hypothetical customer who – instead of optimising the cloud architecture – designs it in a way that requires significant data to be moving between clouds on a constant basis. This "use case" for multi-cloud does not exist in practice, and the CMA's stylised examples of potential frictions created by multi-cloud (as set out in its recent heat map) are not realistic and mainstream. The CMA's response that these scenarios could hypothetically exist but for egress fees completely ignores that there are many other more significant factors and costs – such as security, reliability and technical / performance costs – associated with developing a workload that involves the mass ongoing movement of data between cloud infrastructures.⁷⁸
- (83) While egress fees do not currently drive customer behaviour, if the CMA were to intervene by forcing egress fees to a level below cloud providers' costs (and a return on investment) or to prohibit them completely, this would likely impact customer behaviour. While it would be unlikely to have a meaningful impact on switching or multi-cloud solutions, it would likely lead to excessive and inefficient usage of the cloud and disincentivise cloud providers' investment / innovation in data transfer services and network infrastructure. It would also

⁷⁷ Egress Fees Working Paper, paragraph 2.8 and 2.70; and Jigsaw Research, paragraph 1.4.8.

⁷⁸ Moreover, the CMA's analysis in the Egress Fees Working Paper compares egress fees associated with switching to the customers' annual 2022 spend on cloud computing (paragraphs 2.29 and 2.30). However, the reality – as the CMA also notes in the Egress Fees Working Paper – is that when customers switch, egress fees are amortised or "*defrayed over a number of years*" (paragraph 2.12(c)) and therefore looking at egress fees as a one-off proportion of total annual spend will severely overstate the overall cost and impact of egress fees on customer decision-making. Egress fees are simply not a material factor to customers considering switching or multi-cloud.

undermine other goals set out by the CMA such as clarity and predictability of cloud spend for customers as it will eliminate the connection between customers' consumption of data transfer services and their payment for that service. These concerns particularly apply if not all data transfer fees are allowed following any remedial action and if complex mechanisms were created to determine accepted charges. Predictability and simplicity are critical considerations for all cloud stakeholders.

4.2 Microsoft's strategy on egress is not to lock-in customers to Azure but rather to encourage migration to the cloud and competition continues to drive egress fees down, where they are relevant

- (84) Bandwidth has a real and significant cost. In today's cloud market, the way those (and other costs) are recovered by cloud providers – including Microsoft – is the direct result of the competitive process through which the market has developed.⁷⁹
- (85) Historically, cloud providers charged for both egress and ingress. As competition in the cloud market intensified, Microsoft (as well as other cloud providers) moved to a model of free ingress (i.e. no charge to put your data in the cloud), but more significant charges for egress. These are rational economic reasons why in a competitive cloud market it would be expected that ingress fees become free and egress remains positively priced – specifically this facilitates migration to the cloud, while mitigating overall costs for integrated multi-cloud solutions and for customer switching.⁸⁰
- (86) Competition has also driven egress fees down. Whilst quality-adjusted egress fees have fallen for customers in the UK, [redacted].⁸¹ [redacted].⁸² [redacted].⁸³
- (87) Where there are material benefits to switching or multi-cloud, Microsoft and other cloud providers have designed egress fee strategies to minimise their effects on customers. As recognised by the Jigsaw Research, customers with larger data transfers that incur egress fees beyond the free tiers have multiple options for minimising their costs,⁸⁴ e.g. general cloud discounts and credits, egress-specific discounts and dedicated capacity offers. Microsoft has also entered into the Bandwidth Alliance (a group of cloud providers committed to discounting or waiving data transfer fees for shared customers)⁸⁵ and developed an integrated OCI-Azure service with Oracle (which lowers egress charges and increases latency in a single operating environment).⁸⁶

⁷⁹ See, for example, [redacted].

⁸⁰ See, for example, [redacted].

⁸¹ [redacted]

⁸² For example, improvements to ExpressRoute's data path performance, the expanded availability of ExpressRoute direct locations (see [redacted] and [General availability: Azure ExpressRoute Global Reach: 2 new locations | Azure updates | Microsoft Azure](#)) and Microsoft's roll-out of connectivity with IPv6 internal protocols since August 2022 (see [General availability: ExpressRoute IPv6 Support for Global Reach | Azure updates | Microsoft Azure](#)). As acknowledged by the CMA Guidelines, paragraph 108: "the pattern of prices over time can...indicate the nature of competition" and "static or continually rising prices may indicate a lack of competition".

⁸³ [redacted]

⁸⁴ Jigsaw Research, paragraphs 6.1.26-6.1.27

⁸⁵ [Bandwidth Alliance | Reduce Data Transfer Fees | Cloudflare](#)

⁸⁶ This integrated service provides customers with direct access to Oracle database services running on Oracle Cloud Infrastructure (OCI) and deployed in Azure data centres – see further here: [Microsoft and Oracle Expand Partnership to Deliver Oracle Database Services on Oracle Cloud Infrastructure in Microsoft Azure](#)

(88) In sum, Microsoft’s strategy on egress fees has been driven by intense competition for data transfer services, as well as broader strategies to facilitate switching and multi-cloud solutions.

4.3 The CMA’s hypothetical theory of harm on egress fees is unsupported by the CMA’s evidence which finds that almost no customers consider egress fees as a meaningful barrier to switching or to multi-cloud

(89) Consistent with all available evidence – including the analysis conducted on Microsoft’s behalf [redacted],⁸⁷ Microsoft’s previous submissions and the results of Ofcom’s customer survey⁸⁸ – the CMA’s own customer evidence in the Jigsaw Research does not demonstrate weak competition and, in fact, illustrates that egress fees are not a meaningful barrier to switching nor to multi-cloud:

- (a) The Jigsaw Research “*highlights that in almost no cases were egress fees considered [by customers] as the main or even one of the main barriers to switching or to a multi-cloud approach.*”⁸⁹ Whereas, for example, in the *Mobile Operating Systems Market Study*, the CMA’s survey evidence that 64% of users considering switching referred to at least one of the barriers to switching under investigation was sufficient for the CMA to consider that they pose “*material perceived barriers to switching*”.⁹⁰
- (b) Customers who did switch or multi-cloud consider the egress fees incurred “*a price worth paying*”.⁹¹ In other words, customers consider egress to be a valued service given the technology, infrastructure and complexity involved in the data transfer.
- (c) As supported by [redacted],⁹² egress fees are negligible compared to customers’ overall cloud spend and are, thus, in the CMA’s analysis “*barely a factor*” in choosing a cloud provider (both in terms of the initial cloud provider or ongoing use of cloud services) even among customers that accumulate considerable egress fees.⁹³ This is partly because all major cloud providers charge some egress fees and cloud users tend to view them as a “*cost of doing business*”.
- (d) Based on interviews with medium to large public cloud customers in the UK, the Jigsaw Research did not find any major concerns as regards egress fees impacting their ability to multi-cloud. This is consistent with Microsoft’s experience of its significant customers in the UK (including public sector customers) generally pursuing multi-cloud strategies. Larger customers were also more likely to be paying more egress fees (due to e.g. free tiers applying up to certain data thresholds and significant amounts of data transfers), further demonstrating that egress fees do not act as a barrier or deterrence to multi-cloud.⁹⁴ Moreover, such larger customers are able to exert bargaining power and are

⁸⁷ [redacted]

⁸⁸ [redacted] and the results of Ofcom’s customer survey in which only 6% of respondents identified egress fees as the largest challenge to switching (Ofcom Cloud Services Market Study Final Report, paragraph 5.153).

⁸⁹ Jigsaw Research, paragraph 5.2.2; Egress Fees Working Paper, paragraphs 2.46, 2.48, 2.64 and 2.66, respectively.

⁹⁰ CMA Final Report, paragraphs 3.90-3.91, available here: [Mobile_Ecosystems_Final_Report_amended_2.pdf](https://publishing.service.gov.uk) (publishing.service.gov.uk)

⁹¹ Ibid.

⁹² [redacted]

⁹³ Egress Fees Working Paper, paragraph 2.18; and Jigsaw Research, paragraphs 5.1.2-5.1.3.

⁹⁴ Jigsaw Research, paragraph 5.3.4.

more likely to offer a stronger customer response, for example, by seeking discounts on egress fees (as acknowledged by the Jigsaw Research).⁹⁵

(e) Reluctance to switch or to multi-cloud is limited to a minority of customers where: (a) it is involved in exceptionally high data volumes; (b) mature companies have accumulated a large amount of data over the years;⁹⁶ or (c) in “*limited multi-cloud use cases*” such as using a different cloud provider for backup solutions.⁹⁷ This is consistent with [redacted].⁹⁸ These larger customers are more likely to have a stronger bargaining power and represent a significant competitive restraint [redacted].⁹⁹

(f) The CMA admits that only a few customers “*spontaneously*” identified egress fees as a challenge – including where they “*could be significant even if they could be mitigated to some extent*”.¹⁰⁰ Indeed, “*moderators needed to probe on this as only a few participants brought up egress fees themselves*”,¹⁰¹ and “[I]n very few cases, participants brought up egress fees as a potential barrier themselves.”¹⁰²

(90) In addition, Microsoft (as well as some of the other large cloud providers) have announced that free egress is available for all customers taking their data completely out of the Azure infrastructure via the internet to switch to another cloud provider or an on-premises data centre.¹⁰³ This offer is available to all Azure customers around the world and from any Azure region – including in the UK. Microsoft appreciates the CMA has very recently issued further questions regarding Microsoft’s egress strategy and the impact of this. Microsoft welcomes the opportunity to explore this with you.

(91) Contrary to its previous practice in other Market Investigations and Market Studies and its guidance,¹⁰⁴ the CMA’s concerns in the cloud market are not founded on the actual evidence gleaned from customers in the UK cloud market but on a “hypothetical” scenario – e.g. a hypothetical customer concerned about the potential costs of keeping databases in sync in a multi-cloud architecture – or on the “potential” implications on how egress fees “might” impact customer switching or multi-cloud behaviour.¹⁰⁵ Put simply, there is no evidence of

⁹⁵ The Jigsaw Research acknowledges that: (i) as regards lower spend cloud users in particular, it is common for the new provider to offer significant cloud credits to offset any egress fees incurred with the incumbent provider; and (ii) significant egress fees discounts are offered as part of the overall cloud package of services provided to customers (paragraphs 1.4.25, 3.3.3, 3.6.10 (IT Customer), 6.1.18, 6.1.26-6.1.27 and 6.3.6).

⁹⁶ Jigsaw Research, Qualitative Customer Research Final Report at paragraphs 5.2.2-5.2.4; and Egress Fees Working Paper, paragraphs 2.47 and 2.66, respectively.

⁹⁷ Jigsaw Research, paragraph 1.4.16.

⁹⁸ See, for example, [redacted].

⁹⁹ [redacted]

¹⁰⁰ Egress Fees Working Paper, paragraphs 2.39-2.40.

¹⁰¹ Jigsaw Research, paragraph 5.1.1.

¹⁰² Jigsaw Research, paragraph 5.2.1.

¹⁰³ [Now available: Free data transfer out to internet when leaving Azure | Azure updates | Microsoft Azure](#)

¹⁰⁴ Guidelines for MIRs: Their role, procedures, assessment and remedies, paragraph 127. See also the *Mobile Ecosystems Market Study* as referred to in paragraph (89) and fn. 90 above. The CMA’s hypothetical and abstract analysis is further illustrated in its “Hypothetical Scenarios – Microsoft Heatmap methodology note” where it recognises that certain aspects of egress transfer volumes included in its analysis “*may not be technically feasible for certain customers, or likely would not be implemented in practice*” (paragraph 22(b)) and that the results of its analysis “*are purely intended to be illustrative scenarios*” (paragraph 40) rather than based on real-life evidence of customer behaviour in the UK cloud market.

¹⁰⁵ Jigsaw Research, paragraphs 1.4.15, 1.4.19 and 5.26.

In addition, the CMA dismisses evidence from quantitative surveys (listed below) in the public cloud market due to lack of validity, lacking quality and coverage of respondents. Whilst the methodology and method of data collection vary by survey, it is unreasonable for the CMA to find methodological faults and to dismiss each such survey / report given that cloud providers and other market participants rely on these industry reports in the ordinary course of business to ascertain

harm to consumers as a result of egress fees as customers remain able to switch or to multi-cloud when it makes commercial sense to do so.¹⁰⁶

4.4 Any proposed remedy banning or restricting egress fees below costs risks doing more harm than good

- (92) The remedies proposed by the CMA in the Egress Fees Working Paper – including a prohibition on charging egress fees and capping egress fees by reference to other data transfer fees or to the costs incurred by the cloud provider – would not be effective or proportionate interventions (as is required per the CMA’s MIR Guidelines)¹⁰⁷ given that the evidence shows that egress fees do not have an adverse effect on competition, but rather allow for recoupment of the costs of data transfer which require significant investment (e.g. in subsea and optical cables). In addition, the CMA’s proposed remedies would have unintended and significant adverse consequences (as detailed further below).

4.4.1 No evidence that the CMA’s proposed remedies would facilitate more switching or multi-cloud by customers

- (93) As outlined in Section 4.3 above, the CMA has not adduced any evidence to suggest that egress fees restrict customer switching or multi-cloud; there is, therefore, no evidence to suggest that the CMA lowering egress fees would lead to higher levels of customer switching and multi-cloud. In particular, such a remedy would not help customers at all with the significant non-egress fees-related costs, efforts and complexities of switching cloud infrastructure providers or pursuing multi-cloud strategies (as outlined in Section 4.1 above). The Jigsaw Research describes switching cloud providers as akin to “*moving house*”¹⁰⁸ – and just as removing or lowering the costs of removal companies will not incentivise more people to move homes, it is by no means clear that any remedies targeted at egress fees would necessarily increase switching in the cloud.
- (94) The CMA therefore has no statutory basis on which to take remedial action as the proposed remedies would not be remedying, mitigating or preventing an AEC on competition in the cloud market or any detrimental effects on customers resulting from any AEC.¹⁰⁹

4.4.2 Banning or restricting egress fees below costs (and return on investment) could lead to significant adverse consequences for all cloud users

- (95) The pricing structure under a ban or reduction of egress would mean that cloud customers fees charged are divorced from their usage, leading to negative economic externalities. Banning or setting egress fees artificially low through regulation would distort prices such that they are not reflective of the true underlying costs and value provided and would break the existing cloud model of customers paying for actual services consumed. The risks are

customer cloud spending trends and inform their senior-level decision-making as well as their product and business strategy. These industry reports are widely accepted as being informative of the cloud sector and should be considered, at least cumulatively, by the CMA as important evidence of the prevalence of multi-cloud. None of the reports (listed below) refer to egress fees as a barrier to multi-cloud and they all show that multi-cloud is prevalent among cloud providers.

Sources: Flexera, ‘Flexera 2024 State of the Cloud Report’, 2024; Posey M, ‘Multicloud in the mainstream’, S&P Global Market Intelligence (commissioned by Oracle), February 2023; Public First, ‘Public First Poll for CCIA (Cloud Users)’, ‘Table 56’, June 2023; Bhagavath V and Mehra R, ‘Multi-cloud Networking Will Inflect in 2024; Public and Private AI, Application Resiliency and Cybersecurity Are Top Demand Drivers’, IDC Market Perspective, March 2024; and DeMattia A and Grady J, ‘Hybrid, Multi-cloud Management Maturity’, Enterprise Strategy Group (commissioned by Infoblox), April 2024.

¹⁰⁶ [3<]

¹⁰⁷ CMA MIR Guidelines CC3, paragraph 329.

¹⁰⁸ Jigsaw Research, paragraph 1.3.13.

¹⁰⁹ Section 134(4) of the Enterprise Act 2002.

also highlighted in [redacted]. Significant disadvantages to removing egress fees that would impact all cloud customers and stakeholders would include the following (*inter alia*):

- (a) potential abuser traffic — i.e. the inefficient and excessive usage of the cloud infrastructure network for egress and other data transfers. This is a material risk given that customers would not internalise some or all of the wider costs of their consumption under zero / very low egress fees, leading to wasteful overconsumption of egress / data transfer services and the underlying network infrastructure; and
- (b) data resilience security risks arising from the already significantly high and increasing volume of data traffic via cloud infrastructure.¹¹⁰

(96) As a result of these significant disadvantages (as well as the fact that [redacted], which implies that customer switching is a relevant commercial and competitive constraint), Microsoft concludes [redacted].¹¹¹ Microsoft urges the CMA to consider these factors when considering its potential remedies.

4.4.3 Any proposed remedy must allow cloud providers to recoup their costs and return on investment

(97) Egress fees are required by cloud providers in order to recoup their actual costs incurred for providing data transfer services and so that they can invest in high-quality and innovative services and network infrastructure. [redacted].¹¹² This is also consistent with customers' views in the Jigsaw Research, which found that amongst UK cloud customers there is a "*degree of acceptance of egress fees, i.e. the view that data transfers do naturally incur costs and cloud providers are expected to charge customers for this.*"¹¹³

(98) A ban or reduction in egress fees (below costs and return on investment) would reduce the incentives of cloud providers to invest in high-quality network infrastructure and data transfer services, to the detriment of all cloud market participants. Given that ingress fees have already been driven to zero by competitive pressures, if the CMA were to ban or to lower egress fees, this would also lead cloud providers to seek to recover their costs and return on investment in a less-targeted manner.

(99) For multi-product firms such as cloud providers with high fixed costs that cannot be attributed to a particular product / service and low marginal costs, setting prices equal to marginal costs would fail to recover their fixed costs and would reduce incentives for future investment. In this case, the best approach based on economic principles is to allow a multi-product firm to raise prices above marginal costs for different products / services just enough to cover the common costs and a normal return for the business overall.¹¹⁴

¹¹⁰ [redacted]

¹¹¹ *Ibid.*

¹¹² [redacted]

¹¹³ Jigsaw Research, paragraph 5.1.9.

¹¹⁴ These permitted mark-ups should be inversely related to their elasticity of demand. Under Ramsey pricing, services that are highly elastic will attract lower prices than services that are inelastic, even if they have very similar costs. See Baumol, William J., and David F. Bradford. "Optimal Departures From Marginal Cost Pricing." *The American Economic Review*, vol. 60, no. 3, 1970, pp. 265–83. JSTOR, P267. "*Each price be set so that its percentage deviation from marginal cost is inversely proportionate to the item's price elasticity of demand. According to this result, the social welfare will be served most effectively not by setting prices equal or even proportionate to marginal costs, but by causing unequal deviations in which items with elastic demands are priced at levels close to their marginal costs. The prices of items whose demands are inelastic diverge from their marginal costs by relatively wider margins.*"

- (100) For these reasons, should the CMA maintain that remedies are required on egress despite the overwhelming evidence that they have no meaningful impact on the cloud market, Microsoft considers that any remedies considered by the CMA must:
- (a) be limited to internet egress via ISP, which are typically cheaper egress services – thereby impacting the most price-sensitive customers and price as a competitive parameter as well as retaining incentives for cloud providers to invest in and to build out low-latency and premium offers (e.g. Microsoft’s Premium Global Network); and
 - (b) not prohibit egress fees completely but allow cloud providers to recoup their costs and return on investment in order to retain incentives for cloud providers to invest and innovate in these data transfer services.¹¹⁵

¹¹⁵ Microsoft also notes that the CMA’s analysis of cloud providers’ costs is fundamentally defective as it relies on data from AWS, whilst this egress-related costs analysis is both provider and customer / workload-specific (as recognised by the CMA at paragraph 3.40 of the Egress Fees Working Paper) due to differences in cloud providers’ costs and their egress services. For example, Microsoft has made significant investments in its Premium Global Network, which will allow it to build low-latency infrastructure (e.g. for future AI applications) at higher costs and technical performance. Premium Global Network also has higher costs as this allows Microsoft to provide services with higher committed levels of services and technical performance under the terms of Service Level Agreements. AWS, on the other hand, routes data through the public internet which has resulted in cheaper egress services.