OFCOM CLOUD SERVICES MARKET STUDY
INTERIM REPORT AND CONSULTATION

MICROSOFT RESPONSE

31 May 2023

Linklaters

Table of Contents

1 Introduction and Executive Summary

(1) Microsoft Corporation ("**Microsoft**") appreciates the opportunity to respond constructively in this submission to Ofcom's cloud services market study interim report (the "**Interim Report**") and Ofcom's related proposal to make a market investigation reference ("**MIR**") to the CMA of the "*public cloud infrastructure services market*"[1] in the UK (the "**Consultation**"), both published on 5 April 2023 (the Microsoft "**Response**").

(2) In summary, it is Microsoft's view that the case for a MIR is not made out. The reasons for this view are set out in this Response, which may be introduced and summarised as follows.

1.1 The market is fundamentally competitive – and the Final Report should say so

(3) Today's global cloud services market is fundamentally a product of intense dynamic competition. As the Interim Report recognises, cloud computing "*offers significant benefits*" to customers that have resulted from billions of pounds of investment and ongoing rapid innovation.[2] This is not likely to change in the foreseeable future.

(4) Microsoft agrees that cloud services supplied in the UK are undoubtedly worthy of first assessment: they are "*a fundamental building block for a diverse range of software applications that ultimately benefit consumers and businesses across the economy*" with UK public cloud spending of £4.5 billion to £5 billion in 2021[3] and growing. The companies benefiting from cloud innovation themselves contribute significantly to economic growth and productivity gains across all major economies, including the UK.

(5) It is particularly important for the competition analysis of cause and effect in cloud services to be highly robust for at least two reasons. First, the size and strategic importance of cloud services. Second, the fact that this fast-growing market is fundamentally competitive and market rivalry is fundamentally dynamic, not static or mature. It will be common ground that intervention should not unintentionally compromise or ossify these dynamics at the expense of the UK public sector, small business or large enterprise customers.[4]

(6) Competitive intensity is directly evidenced by actual market outcomes and dynamics on innovation, quality of services and price – as detailed in this Response (see especially Section 3.2). These market outcomes are driven by:

   (i) multi-billion pound capex and R&D spending in both infrastructure (such as building data centres, laying undersea cables) and innovation (development of new features and applications) by multiple competitors including no less than three so-called "*hyperscalers*" (Amazon, Google and Microsoft) and providers with strong existing enterprise customer bases such as Oracle and IBM;

   (ii) a large number of on premises workloads anticipated to move to the cloud and the expected emergence of cloud-native workloads in the future; and

   (iii) growth among a large array of platform services vendors.

[1] References in this document to "cloud", "cloud services" or "cloud market" below should be read as limited to Ofcom's definition of "public cloud infrastructure services" unless otherwise noted.

[2] Interim Report, para 3.11.

[3] Ofcom's Consultation Document, paras. 3.13-3.15.

[4] References to "customer(s)" in this Response refer to public sector and corporate customers, rather than individual consumers.

(7) The billions in supplier expenditure are direct evidence of intense rivalry to win workloads, whether those of first-time cloud users or – in an environment in which multi-cloud procurement is the emerging norm – those of existing cloud users.

(8) Specifically, in Infrastructure as a Service ("**IaaS**"), this competitive investment has flowed into developing custom-built hardware and software to increase performance and facilitate lower infrastructure costs; and in Platform as a Service ("**PaaS**"), into designing and deploying cutting-edge innovations such as AI, machine learning and enhanced security applications. These innovations are accessible to all cloud customers, regardless of size or cloud maturity and are relatively simple to integrate into an existing cloud environment, facilitating multi-cloud adoption. These investments and innovations have unlocked access to cloud computing resources that, as Ofcom notes, are "*faster to deploy, more flexible and potentially cheaper*", in turn supporting "*innovation and growth*" by cloud customers.[5]

(9) While many of the benefits of public cloud services over on premises IT or privately run server farms may be inherent to cloud technology, the pace of the roll-out and scale of these benefits is due to intense competition. As set out in Section 6, Microsoft Azure's published prices have generally decreased over time and both new customers and existing customers renegotiating their contracts alike are eligible for substantial discounts when purchasing via plans that give Microsoft greater insight into capacity required in the future, generating efficiencies that in a competitive market are also passed on to customers.

(10) These outcomes should be central to Ofcom's analysis and must directly inform how Ofcom weighs the probability that ultimate intervention following an 18-month MIR could make a competitive market incrementally more competitive still, against the risk that any ultimate such intervention would do more harm than good (by inhibiting or distorting competition on price, cost recovery, or innovation and services differentiation).

(11) The Interim Report marginalises this critical context. It acknowledges strong competition between three hyperscalers to win customers – that "*is leading to some positive outcomes*"[6] – but claims that once customers are won, they are "locked in" and either actually exploited or at risk of being exploited in the future (e.g. through higher prices on renewal).

(12) While the recognition of "*some positive outcomes*" is an understatement, the central contention is simply not correct. Azure does not exploit "locked in" customers on price while it competes for new ones, not least because this dichotomy is false. Nor is there a realistic possibility that Microsoft or any other cloud vendor can profitably slow their rapid pace of innovation as a result of IT lock-in effects. As IT vendor/customer relationships endure over time and are not "one shot games", Azure is perpetually competing for the next workload (see Section 1.2 below).

(13) In developing an overarching risk-of-exploitation thesis, the Interim Report places emphasis on aspects of market structure or performance used as indicators of (or factors that appear consistent with) a "big picture" of weak competition and poor customer outcomes, notably (i) "*high*" concentration including "*stable*" UK revenue shares; (ii) *"low"* switching rates and (iii) *"high"* supplier profitability. This narrative is fundamentally inconsistent with a detailed assessment of current competitive dynamics. It focuses on certain theoretical indicia of expected long-run market outcomes and not direct evidence of market outcomes. It also fails to situate today's competition in the evolution of market dynamics (i.e. how Azure challenged Amazon Web Services ("**AWS**") when it entered the market in 2010; and Google, Oracle and IBM's more recent challenger strategy to do the same). Ofcom's limited profitability analysis in

[5] Interim Report, para 1.1.

[6] Interim Report, para 1.15.

the Interim Report is compromised by its inability properly to consider the extensive fixed-cost investments made over time and the relative immaturity of the investment cycles at issue.

(14) In fact, competitive dynamics tell a different story than a static account of market structure and critically, direct evidence – including that set out in the Interim Report[7] itself as well as supplemental evidence contained or referred to in this Response – shows strong customer benefits and no customer harm. A cost/benefit conclusion favouring a proposed MIR cannot be squared with this direct evidence.

(15) The proposal to make an MIR, therefore, would need to rest, at its highest, on theories of future potential harm rather than the observable and often measurable competitive conditions and customer outcomes of *today.*

(16) Microsoft respectfully submits that a Final Report should include a clear diagnosis on the current state of competition and customer outcomes. Microsoft's case, as evidenced and explained in this Response, is that the market can only fairly be described as highly competitive and that the customer benefits this rivalry has yielded are very significant.

(17) Microsoft therefore urges Ofcom's Final Report to say so, not simply for the record – in order properly to reflect a 12-month competition analysis – but to build any cost/benefit analysis of possible MIR intervention on a sound foundation.

1.2 Competition for the "next workload" is enduring and any dichotomy between new and "locked-in equals exploited" customers is false

(18) The Interim Report suggests that, in Ofcom's view, the risk of exploitation of "lock in" is high because switching costs are "*high*", and inter-cloud switching is "*low*". There are several problems with this view. The first, is that the Interim Report has no comparison or benchmark as to what switching level should be ideal or expected: low relative to what? Second, even if switching levels can robustly be characterised as (comparatively) "*low*", this finding is wholly inconclusive: low switching is perfectly consistent with a competitive market in which existing customers are <u>not</u> exploited relative to new customers but instead derive essentially the <u>same</u> competitive benefits. If customers already have a competitive deal, the relatively modest if any benefits of switching would not very often outweigh the unavoidable costs (given all IT environments have a degree of lock-in; see paragraph (24) below).

(19) Of course, a dichotomy between (1) inert and possibly vulnerable existing customers, subject to a "loyalty penalty", and (2) new customers, who reap the benefits of competition may very well be a sensible explanatory tool for UK consumers and their mobile phone and broadband contracts,[8] or their gas and electricity bills.[9] But applying this narrative extrapolated from various retail markets (a mature B2C setting) to global cloud services to enterprises (a dynamic B2B setting) would be misplaced, for two main reasons.

(20) First, it rests on a simplistic dichotomy between "new" and "locked in" customers and the incorrect notion that the market will in the foreseeable future reach some steady "mature" state in which there are few or no "new" customers left to compete for. Second, it erroneously equates the existence of switching costs with the absence of choice and with supplier market power (or

[7] Ofcom recognises in para 6.12 of the Interim Report that there are "*positive outcomes for customers that are likely to be driven by the competitive dynamics in cloud infrastructure services*" which include product innovation, customer choice of ISVs, customer choice of open-source products / technologies and decreases in prices in IaaS.

[8] These being the two "Ofcom markets" out of the five that were the focus of the 2018 *Loyalty Penalty* super-complaint to the CMA, the three others being cash savings, mortgages and household insurance.

[9] See the CMA's *Energy* MIR (2016).

"exploitation" of customers with IT environment "lock-in"). We address each of these issues in turn.

Competition is driven by the scale of opportunity for "next workloads" in multi-cloud regardless of the "cloud maturity" of the customer or workload

(21) Cloud is ultimately a volume (capacity-utilisation) driven business with high upfront investment costs. A new workload is incremental revenue, and the core directive of Microsoft's sales associates is to increase server consumption to fill capacity and keep utilisation high.

(22) First, the market opportunity for suppliers to convert on premises IT workloads to inaugural or "first time in the cloud" workloads is huge and dwarfs the size of the current cloud market. According to various industry reports, the cloud market is estimated to be at least 7 - 12 years away from reaching any stable equilibrium between on premises and cloud IT spend.[10] Chasing conversion of on premises IT spend into cloud spend drives fierce competition as the cloud market continues to grow at a healthy pace.

(23) Second, even for "existing" cloud customers, a dichotomy between stable and "locked-in" or "mature" cloud workloads and contestable "new" or "immature" IT workloads is incorrect.

(i) **The new normal is multi-cloud procurement:** Contrary to the implication in the Interim Report, multi-cloud procurement is the prevalent model going forward. Examples of Azure multi-cloud customers in the UK include [✂], among many others. None of Azure's significant customers are pursuing a single-cloud strategy over the medium-to-long term. Multi-cloud drives competition because it means a customer's cloud IT "share of wallet" is up for grabs; it is not the case that once an enterprise has cloud spend with one provider, it will not award new workloads to another supplier. Providers therefore compete for each workload as the customer considers whether to deploy to the cloud and which one.

(ii) **New and emerging cloud-native workloads create new opportunities:** In addition to workloads originally migrated from on premises, Azure continues to compete for new and emerging cloud-native workloads that have evolved to take full advantage of the flexibility that the cloud provides. This trend is set to continue and represents another important driver of competition regardless of the "maturity" of migrated workloads.

(iii) **Like on premises, cloud IT is a process of perpetual renewal (replacements and upgrades) which generate "next workload" opportunities:** Finally, the rule of perpetual IT renewal drives competition for customer relationships and the next workload. Not least in a multi-cloud paradigm, there is scope for competition for an existing workload at the point of renewal, replacement, and upgrade (an endemic feature of IT procurement). Given that business-critical services are run through the cloud, such renewal, replacement and upgrade is not an optional "nice to have" but critical to customers' ability to compete in their own markets. This means that a sizeable share of notionally "old" or "existing" cloud workloads are, in principle, contestable and

[10] During a recent AWS earnings call, Andy Jassy made the statement that "*it's also useful to remember that 90% to 95% of the global IT spend remains on-premises. And if you believe that, that equation is going to shift and flip, I don't think on-premises will ever go away, but I really do believe in the next 10 to 15 years that most of it will be in the cloud.*" Given this statement was made by AWS and other sources (see The case for a $2 trillion addressable public cloud market) have indicated that approximately 20% of on premises workloads have already moved to cloud, we have assumed the midpoint of 12 years as a reasonable upper bound.

Cloud as a percentage of IT spend is expected to reach 25% (see The Cloud Is Still a Multibillion-Dollar Opportunity. Here's Why (forbes.com)) where it is currently at 17% of total IT spend. Given this ratio would be extrapolated to reach 25% in 2030, we have estimated 7 years as a reasonable lower bound.

subject to constant review if the current technology becomes outdated or the switching benefits otherwise outweigh the switching costs.

(24) Microsoft Azure's cloud business does not distinguish or adapt its competitive terms based on the "cloud maturity" of either the customer or the workload in question. Given the enduring nature of IT spend in a multi-cloud paradigm, and competition for the next workload, Azure has not adopted a strategy of compromising customer relationships by unfairly ratcheting up prices once they are "locked in" and nor does it plan to: its real-world marketplace actions are consistent with the belief that volume gains from winning future workloads (and not short-term margin extraction) are the right strategic focus. In fact, analysis of pricing trends for Azure customers described in more detail in Section 6 below, shows the opposite trend of prices generally either decreasing (i.e. discounts increasing) or remaining stable (i.e. discounts remaining stable) upon customer contract renewal.

IT lock-in is neither caused nor "exploited" by Azure

(25) There is no question that customers who want to change providers will incur switching costs of varying degrees based on the choices of the customer when designing its cloud deployment and also based variously on personnel familiarity with an IT environment, IT solution customisation by and for customers, plus logistical factors. It is these switching costs which Ofcom refers to as "lock in". But as Gartner points out, switching costs are inherent to all IT solutions although Microsoft believes that all else equal, switching between cloud environments is probably the easiest of all settings.[11] Customers have the ability today to rely upon technologies that are effectively standardised across clouds and would allow for easier switching but doing so also reduces the benefits that customers receive from differentiated or custom features in IT solutions. This trade off explains why the Interim Report finds that there are technologies that "*may in principle facilitate switching . . . (but) take-up of some of these technologies may be limited*."[12] Thus, increased switching costs are generally understood and, on balance, willingly borne by sophisticated customers in order to build the most effective solution for their particular needs. Nor do these features (again, endemic to all IT procurement) imply that customers do not have a choice; on the contrary, cloud customers make informed decisions about the level of putative "lock in" they will accept and when to incur costs (i.e. upfront – to invest in a provider-agnostic architecture; or deferred – to use a differentiated, provider-specific architecture upfront and to incur switching costs if and when a decision to switch is taken).

(26) Customers must overcome substantial "on premises IT lock-in" to switch to the cloud in the first place. But notwithstanding the costs in time, money and resources, customers are doing this at scale because cloud benefits greatly outweigh the on premises to cloud-switching costs. Ofcom should not expect customers to move existing workloads between clouds if the customer is already receiving competitive prices and high levels of innovation and quality (which they are; see Section 3) because the gains from switching may be limited. This may also introduce significant complexity along a variety of dimensions without any meaningful benefits, such as requiring customers to manage different commitments and capabilities from different providers with respect to security, privacy, regulatory compliance, resiliency, sustainability, procurement and management (and more) for a single cloud workload. For these reasons, while multi-cloud procurement makes sense for many good reasons, the notion of dividing or replicating one cloud workload across multiple providers (what the Interim Report refers to as "integrated multi-cloud") would rarely make economic sense.

[11] Relative to switching from on premises to different on premises or from on premises to the cloud.

[12] Interim Report, para 5.14.

(27) Instead, competition is perpetually for the next workload, whether it be a new cloud native workload or a new switch from on premises, or the threat of customers moving existing workloads (if they see sufficient benefit in doing so) as part of their multi-cloud deployment.

1.3 The Interim Report's specific concerns do not reflect remediable problems of a genuine "lack of competition"

(28) The Interim Report identifies three market features that it considers may limit competition. While Microsoft agrees there are inherent IT barriers to switching also in the cloud (see paras. 25-26 above), it disputes that committed spend discounts, provider-imposed interoperability factors, and/or egress fees either aggravate "lock-in" or should be construed as exploiting putatively "captive" customers (or being capable of meaningfully facilitating such action in future).

(i) **Committed spend volume discounts are the essence of how Azure competes on price for larger users and plans for needed capacity:** In common with many other high fixed cost markets, committed spend discounts are mutually beneficial: they allow customers to access lower-cost services and better predict cloud costs and usage; at the same time, they assist cloud providers with vitally important long-term capacity planning, given the need for high-fixed cost infrastructure investments that are location-critical. Such discounts are not designed or imposed by cloud providers to capture the entirety of a customer's demand. Rather, commitments are generally planned by customers on the basis of expected spend on specific workloads, similar to planning on premises infrastructure investments but with greater flexibility to change plans on the cloud. These volume discounts are not, therefore, disguised exclusivity arrangements and do not aim to foreclose competitors; customers themselves have flexibility to determine their committed spend targets based on their expected workflow and to determine the contract length. As pro-competitive volume discounts, any future intervention to limit their usage in the UK would directly (and perversely) *harm* the very UK customers – relevant Azure customers include the [✂] – that intervention should aim to *help*. At the same time, intervention would undermine efficient capex spending by limiting cloud providers' ability to forecast customer demand.

(ii) **As customers choose their IT solution trade-offs, Azure's strategy is to mitigate interoperability friction:** Microsoft's aim and incentive as the historical "challenger" in cloud has been two-fold. First, to design its Azure services to work as well as possible with third-party applications; a system cannot be made technically interoperable in such a way as simultaneously to make both joining easy but leaving difficult. Second, to provide new and innovative value, including that Azure's own services work seamlessly together, to make it worthwhile for customers to switch. Both elements increase the chances of winning customers away from AWS (and now others). The flip side of seamless integration *within* an environment is less-than-perfectly-seamless integration *across* differentiated environments, but contrary to some implications in the Interim Report, there is simply no nefarious motivation to make Azure's services "work less well" (or not at all) with other environments. The "presence" of lock-in does not signal an automatic "absence" of interoperability. Nor is interoperability friction a binary state. Rather, customers – who are typically large and sophisticated enterprises – make strategic choices to acquire a range of IaaS and PaaS products and have a preference either to enhance portability or focus a larger portion of their workloads on a single platform to take advantage of well-integrated services and better performance. While customers may have an understandable preference for an ideal state in which all innovative systems could work seamlessly together, a simplistic "desire" belies the feasibility of facilitating "frictionless" operation of disparate systems, which if

implemented would – as history has shown – inevitably lead to standardisation around the "lowest common denominator" and reduce the value of the services concerned. Given the dynamic nature of competition, Microsoft has serious concerns that any ill-designed intervention in this area will stifle innovation. This would sacrifice (degrees of) dynamic competition and innovation which has served cloud customers so well thus far.

(iii) **Egress fees are the organic result of competition, not evidence of "exploitation" or lock in:** The current market practice of allowing free data ingress and charging for data egress has evolved organically as part of intense competition. When Azure entered the market and cut ingress fees entirely in 2011, AWS then also dropped its ingress fees to zero,[13] reflecting the competitive pressure that Azure exerted on the incumbent. The egress-only pricing model lowers the migration costs for customers moving to the cloud for the first time and reflects the strong competition between providers to win workloads migrating to the cloud. Meanwhile, charging for egress allows cloud providers to recover the costs of transferring data in and out of the cloud. Data storage and transfer has a cost and therefore in Microsoft's view, charging for data transfer (including between Azure regions) is important. For customers, egress fees provide transparency of costs and the fees are payable only by customers that consume such egress services. Customers consider egress costs in designing the flow and usage of data in their solutions, and thus charging for data aligns the customer and provider incentives to seek the most efficient organisation of such data flow. An intervention that made data egress/transfer "totally free" would result in excess (suboptimal) usage, while the Interim Report is wise to flag the risks of distortion and cost inherent in price regulation.

1.4 Conclusion: the case for intervention is not made out

(29) Cloud services are important, but this alone does not suffice to justify extending a 12-month Ofcom market study inquiry into an additional 18-month MIR by the CMA. It is Microsoft's view that the case for an MIR is not made out. The potential concerns identified by Ofcom, when read though the correct lens of a highly competitive market today, would suggest allowing dynamic market evolution to play out before commissioning any detailed assessment with a view to intervention.

(30) Set against the unquantified, incremental benefits of addressing (at present, theoretical and, as this Response will show, misguided) future competitive harm are the real costs of intervention. On this point, Microsoft commends the effort shown towards balance in many of the Interim Report's conclusions, in particular the recognition of the risks of side effects and unintended consequences of candidate interventions intended to compel the market to function "better". While the Interim Report is the stronger for these acknowledgements, it still fails to make a sufficiently compelling case for an MIR to "fix" significant live or imminent competition problems in the cloud market via "smart regulation".

(31) As Ofcom correctly acknowledges, the possible remedies identified in the Interim Report are difficult to design, monitor and enforce in a market that is fast-evolving, comprising different and rapidly evolving services and with a complex pattern of national and international contracts. The historic national-market MIRs conducted in the UK do not resemble the cloud market in which providers innovate globally, offer standard global policies, where pricing is not at national level, and where customers often require services in multiple countries.

(32) Given the fast evolving and highly dynamic nature of cloud, any intervention could very well be net negative and leave UK enterprises and public sector customers worse off – if, for example,

[13] Amazon matches Azure free inbound data offer - Networking - Cloud - iTnews.

UK customers were deprived of discounts, or of innovation through "interoperability" requirements which favour portability considerations over innovation and differentiated value add services. It would be a particularly unfortunate outcome if UK businesses and public sector customers faced less vibrant and competitive cloud solutions on a global stage than those available to their rivals in the EU, the US and China.

(33) Microsoft continues to seek constructive engagement with Ofcom to help sharpen the focus of inquiry and provides its comments in the interest of the robustness of the assessment. The below response and critique should be read in this constructive spirit.

2 Structure of this Response

(34) This remainder of this Response is divided into three parts:

(i) **Part A** explains that current market outcomes are highly competitive, need to be understood in the context of market evolution, and have unarguably led to highly beneficial outcomes for customers (**Section 3**);

(ii) **Part B** refutes Ofcom's core contention that customer lock-in means that suppliers have market power over their existing customers and **exploit** them (or have the ability and incentive to do so in future) and that pernicious supplier practices **reinforce and augment lock-in** (**Sections 4** - **7**); and

(iii) **Part C** addresses the specific competitive concerns flagged in the Interim Report that form the basis for candidate "adverse effect on competition" ("**AEC**") findings and the proposed MIR, namely committed spend discounts (**Section 8**), interoperability (**Section 9**) and egress fees (**Section 10**).

Part A: The competitive context of today's positive customer outcomes

3 Intense cloud competition has yielded significant customer benefits

(35) Ofcom concludes in the Interim Report that it "*can see positive outcomes for customers that are likely to be driven by the competitive dynamics in cloud infrastructure services*"[14] citing (i) "*continual*" innovation; (ii) "*broad*" customer choice of ISVs; (iii) customer choice of open-source products; and (iv) "*generally decreas(ing)"* pricing trends for compute.[15]

(36) This assessment is correct. It is not possible clinically to separate and quantify the intrinsic benefits of cloud (over on premises IT) versus those incremental benefits that flow from the intensity of cloud rivalry. However, one can broadly compare the benefits that on premises customers received when they migrated to AWS as the first-mover to the subsequent growth of the cloud market spurred on by challenger competition from Azure, Google, IBM and Oracle, among others; Ofcom cites the 50% price drop in AWS general purposes virtual instances over time, for example.[16] Absent such intense competition, it is far from plausible that the market would have expanded so fast and offered such vibrancy of innovation.

(37) As part of a national economy that was an early adopter of the cloud, UK businesses have already benefited significantly from rivalry-fuelled innovation.

(38) Microsoft believes it is critical that Ofcom's analysis fully reflects how the market has evolved to date and continues to evolve today. In this section, we outline the evolution of the market and how competitive dynamics have driven customer benefits over time (including to major customers in the UK such as [✂]) – both in relation to non-price benefits (e.g. innovation and the proliferation of value-added services) and price benefits. We then outline how the synergy of competition at the IaaS and PaaS infrastructure levels has driven customer benefits.

3.1 Today's market structure is the current snapshot of a process of dynamic rivalry as Azure and others expanded to erode AWS' first-mover position

(39) In evolving technology markets, it is settled wisdom that a static or snapshot assessment is not appropriate. This applies in full to cloud services. Before launching any analysis heavily premised, for example, on current global market structure and market shares (or the subsegment of UK shares of supply by revenue), it is important not to start the narrative immediately in 2023 (using data from e.g. 2020-22) but to contextualise how the market has evolved to its current state.

(40) Cloud-based services were launched throughout the 2000s. AWS was the first to offer a modern public cloud service when it launched the Amazon Elastic Compute Cloud (EC2) and Simple Storage Service (S3) in 2006.[17] As the first-mover in the industry, AWS sold aggressively and built what has become a strong position as the market leader in public cloud.[18]

(41) Microsoft announced its intention to enter public cloud two years later in 2008, and its original entry was a PaaS offering catered to building Windows-based applications. Substantial R&D investment was necessary to offer higher performance and reliable cloud services. In FY2010

[14] Interim Report, para 6.12.

[15] Interim Report, para 6.12(a) to (d).

[16] Interim Report, para 6.12(d) and footnote 699.

[17] History of AWS - W3schoolsHistory of AWS - W3schools.

[18] Gartner Says Worldwide IaaS Public Cloud Services Market Grew 31.3% in 2018; Gartner Says Worldwide IaaS Public Cloud Services Market Grew 31.3% in 2018.

alone, Microsoft invested $8.7 billion in R&D, most of that devoted to cloud technologies, and had 70% of its 40,000 engineers working on cloud-related products and services[19] with substantial increases in R&D in later years.

(42) Azure was initially not well received: executives explained that it had prioritised performance and reliability over the delivery of other features.[20] Microsoft's initial PaaS was also less well-suited than IaaS to the economic and technical requirements of customers seeking to migrate to the cloud.[21] In response, Microsoft dedicated further R&D investments to improve Azure and in 2012, it launched its Azure IaaS offering.[22] Azure even embraced Linux and moved away from Windows-only on cloud (even where that meant lower revenue for its own products).

(43) Google first launched cloud services in 2008. Its first cloud products were tightly coupled to Google's own App Engine, a PaaS offering, instead of offering innovative infrastructure services. This approach reduced its attractiveness to enterprise clients. The final issue was Google's lack of experience catering to those enterprise clients. Its cloud services initially targeted developers already familiar with cloud computing, resulting in weaker documentation and a less flexible platform.[23] These strategic missteps may have contributed to its smaller market share than AWS and Azure so far, but Google has revised its strategy and a recent analyst report found that "*we see no structural reason why Google Cloud operating margins should not mimic that of its larger peers such as Azure and AWS*".[24] For example, to address its unfamiliarity with enterprise customers, Google hired Thomas Kurian (an ex-Oracle Vice President) as the head of its cloud division in January 2019. Under Kurian's direction, marked by an aggressive focus on enterprise customers[25][26] and industry-specific vertical solutions,[27] Google Cloud's annual revenue increased by approximately 230% in 2021.[28]

(44) This is consistent with Microsoft's own experience which is that Google Cloud Platform ("**GCP**") is and will remain an ever more credible competitive threat to win business from Azure and no doubt from AWS and others.

(45) As the move to cloud accelerated, some existing business with traditionally strong positions managing on premises infrastructure, like Oracle and IBM, initially questioned the value proposition of cloud for customers and instead considered it to be marketing hype.[29] Their cloud strategy has since pivoted, having realised their initial strategic mistake.

(i) **IBM** was early to launch public cloud services, in 2007. However, its focus remained on customer data centres, and it did not offer the ability to run applications on its cloud.[30]

---

19 MSFT Annual Report 2010 (microsoft.com).

20 Microsoft: Features still missing in Windows Azure | InfoWorld.

21 Time for a Second Look at Windows Azure? | WIRED.

22 Microsoft Offers New IaaS Services To Compete With Amazon -- Redmondmag.com.

23 Analyze Google's cloud computing strategy | TechTarget.

24 Benjamin Black et al, "One Alphabet Avenue – The Unorthodox Real Estate Play," *Deutsche Bank Research*, Mar 10, 2022, p 16.

25 Google Cloud boss Thomas Kurian's rocky path to profit (cnbc.com).

26 Thomas Kurian on his first year as Google Cloud CEO | TechCrunch.

27 How Thomas Kurian's 'Quite Simple' Strategy Is Transforming Google Cloud | CRN.

28 Google Cloud customers and partners grade CEO Thomas Kurian - Protocol.

29 In 2008, Oracle CEO Larry Ellison called the cloud trend "idiocy" and shared his view that it was a meaningless designation: "I can't think of anything that isn't cloud computing with all of these announcements. The computer industry is the only industry that is more fashion-driven than women's fashion. Maybe I'm an idiot, but I have no idea what anyone is talking about. What is it? It's complete gibberish" (Oracle's Ellison nails cloud computing - CNET). Similarly, in the early days of the cloud, IBM CEO Sam Palmisano outlined that IBM's focus was not on developing a cloud to support customer needs and that the company was secure in its existing businesses, stating in 2010, "Enterprise will have its own unique model. You can't do what we're doing in a cloud." (Palmisano outlines IBM's 2015 roadmap; Earnings to double; Consumerization mocked | ZDNET).

30 See Red Hat's KVM virtualization proves itself in IBM's cloud and What is each server for: Dev, Test, UAT, Staging, Demo and Production.

IBM failed to invest enough in cloud despite having the capital means to do so.[31] More recently, IBM has optimised its cloud platform for regulated industries (e.g. financial services). It has a publicly stated hybrid cloud strategy, based on leveraging the RedHat acquisition along with their incumbent expertise of their client's infrastructure and workloads to create what they referred to as a "*differentiated hybrid cloud value proposition*."[32]

(ii) **Oracle** launched its IaaS offering, Oracle Cloud Infrastructure ("**OCI**"), in 2016. Oracle refocused its cloud strategy in 2016[33] and invested in cloud offerings[34] after it delayed investment (reportedly due to little incentive to cannibalise existing revenues[35]). It has recently invested more in data centres and, despite being a late entrant, has a comparable number of data centre regions as AWS, Azure and GCP.[36] While OCI has been regarded as a destination for traditional Oracle workloads, it is also winning commitments in the broader market for non-Oracle workloads, serving over 18,000 customers including sizeable customers like Zoom and TikTok.

(iii) **OVHCloud** began offering public cloud IaaS in 2011 after over a decade as a hosted server provider. It has a global presence but operates primarily in Europe. OVHCloud provides valuable service adjacency, given that it also offers dedicated private servers, web hosting, and hosted private clouds, which makes it a helpful partner for companies still early in their transition to the public cloud.[37]

[31] "Instead of investing in cloud data centers, IBM has managed by press release for years, with current CEO Virginia Rometty continuing the policy of her predecessors. IBM hypes blockchain, it hypes Artificial Intelligence, and it hypes its Watson front-end. But it lacks the capital firepower to capitalize." (A Tech Giant No More: IBM Is Too Small to Compete in the Cloud Era; Cloud computing and AI: Can IBM finally catch the wave this time around?).

[32] IBM 2018 Q4 Earnings Call.

[33] Oracle's cloud strategy is simple – woo and win the latecomers.

[34] In 2018, Oracle unveiled a key technical innovation, **Autonomous Database,** to further strengthen their IaaS position. Autonomous Database is a fully automated database service that makes it easy for all organizations to develop and deploy application workloads regardless of complexity, scale, or criticality. In addition, Oracle focused on developing their cloud ERP products, Fusion, and NetSuite, to compete in the SaaS layer of the cloud. Between 2018 and 2020, Oracle also focused on technically differentiating OCI's infrastructure from AWS, with a particular focus on security.

Sources: Oracle 2018 Q3 Earnings Call; Oracle 2019 Q1 Earnings Call; Oracle 2019 Q4 Earnings Call.

[35] Analysts noted that "Moreover, we see the dismissal of Cloud Computing as a fad consistent with a strategy to 'delay, distract, and learn' while Oracle assesses a way to address this architectural shift. We continue to believe Oracle will face challenges navigating what we believe is the largest technological shift in the software industry in 20 years as there is little incentive for Oracle to embrace and cannibalize over $35b in revenues by investing in generationally more efficient multi-tenant SaaS." (Deutsche Bank, 'Oracle World brings few surprises', 26 September 2008).

[36] See AWS Global Infrastructure; Global Locations - Regions & Zones | Google Cloud; Azure global infrastructure; Oracle Regions and Availability Domains; IBM Cloud global data centres.

[37] "IDC MarketScape: Worldwide Public Cloud Infrastructure as a Service 2022 Vendor Assessment", November 2022, p 14.

(46) Annex 1 provides further information on the history of Microsoft and other cloud providers' strategies in the cloud market and how it affects competition in the market today.

3.2 Direct evidence of market outcomes benefitting customers are undeniably the product of intense competition

(47) Ofcom's Interim Report finds that "*our study has identified various indications that the market is not working well*".[38] Before addressing the indirect indicia[39] cited in support of this conclusion, Microsoft submits that the most direct and compelling evidence of whether a market is working well is the outcomes on the demand-side – i.e., those experienced by customers whose interests are served by strong competition (and the intended direct beneficiary of any pro-competitive regulatory interventions). As it is often said by the CMA, "dynamic" competition that drives innovation is more important than "static" price competition for long-term consumer outcomes. This section therefore starts with evidence on innovation.

3.2.1 Evidence: innovation and proliferation of value-added services

(48) Capgemini concludes that the hyperscalers maintain "*a very high rate of innovation*"[40] as cited by the Interim Report, which acknowledges that AWS, Azure, Google and Oracle "*are consistently adding new products and features across the entire cloud stack"* requiring *"continuous investment and innovation" [41]* and that the hyperscalers are "*continually developing new products*" as well as the features of existing products, and that they are at the "*forefront of cutting-edge technologies such as ML (machine learning) and AI (artificial intelligence)."[42]* Ofcom also cites evidence of the hyperscalers investing in making physical cloud infrastructure components (e.g. chips) more powerful and efficient.[43]

(49) These summary observations are evidenced by further data points not cited in the Interim Report. AWS, Azure and GCP alone have each introduced hundreds of product improvements and new products in the last five years alone:

[38] Interim Report, para 6.13. Emphasis added.

[39] These are (i) low switching; (ii) limited integrated multi-cloud, and AWS 'and Azure's (iii) high shares of supply and alleged "high" profitability; cf. Interim Report, para 6.13(a)-(d). These points are addressed in Part B.

[40] Interim Report, para 4.52(b).

[41] Interim Report, para 5.227.

[42] Interim Report, para 6.12(a).

[43] Interim Report, para 6.12. See also Annex 2 for a summary of some of the many key innovations by each hyperscaler and by product category in cloud infrastructure services.

Figure 1: Cumulative product releases, 2017-2022[44]

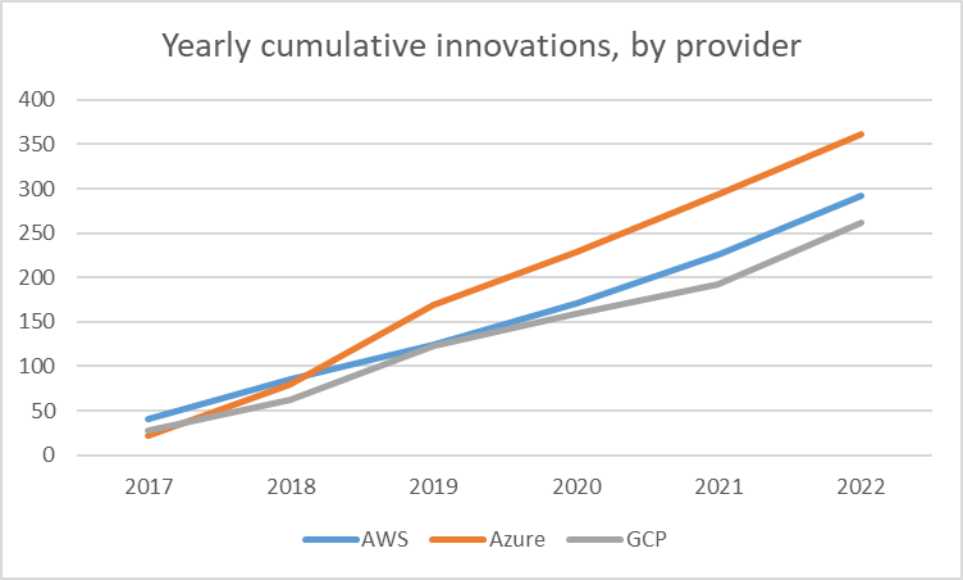

(50) These releases include enhanced capabilities of existing products which amount to quality improvements as well as entirely new products. In conjunction, cloud providers have invested heavily in the cloud infrastructure by building custom chips that improve performance and reduce costs, and they are investing in improved middleware that reduce infrastructure usage for customer as well as in storage, networking, security.[45]

(51) Such innovation is consistent with large investments. Economic consultants retained by Microsoft assessed investments in cloud from 2018 to 2022 through capital expenditure[46] and R&D expense (software development for example). As AWS, GCP and Azure do not report capex and R&D specifically for cloud, reasonable estimates have been calculated as follows.

(52) As a conservative approach, R&D expenses were allocated based on the percentage of cloud revenue. For capex, the range of capex was based on multiple allocation assumptions. In the conservative case, capital expenditure was allocated based on the percentage of cloud revenue. For the base case, some 80% of Microsoft and GCP's[47] and 40% of AWS' total capital expenditure[48] to cloud were allocated based on comments made by each vendor's respective CFOs during earnings calls. Based on this analysis, it is estimated that AWS, Microsoft, and GCP have invested in the following respective ranges for cloud between 2018 and 2022: $58-110 billion, $62-101 billion and $19-112 billion. Taking the lower bound estimate for each

[44] Based on 2022 annual cloud conferences of each of AWS, GCP, Azure.

[45] Middleware are just not additional features but significantly reduce infrastructure usage for customers as it enables much more efficient use of server resources. In the context of cloud computing, middleware abstracts the underlying communication between components, simplifying the process of developing, deploying, and running applications. Examples of cloud middleware include container management services, runtime environments, integration middleware services, and process optimization and decision management services (see What is middleware? (redhat.com)).

[46] Capital expenditure accounts for providers' investments in servers, including their chips, networking equipment, real estate, etc.

[47] Ruth Porat (CFO, Alphabet), Q4 2022 Earnings Call: "an important mix shift. We're increasing our investments in technical infrastructure. And that's not just for AI. That's to support investments across Alphabet, in particular in Cloud as well."; Q2 2022 Earnings Call: "Turning to CapEx. The largest investments in the second quarter were in servers followed by data centers...".

[48] Brian Olsavsky (CFO, Amazon), Q4 2021 Earnings Call "Let's talk a little bit about capital expenditures. And I'm going to do this with the inclusion of equipment finance leases, which is the residual that we sometimes lease on our infrastructure assets (…) So, when you look at those numbers and how they've grown over the last few years, I'll give you the proportions, which I'm not sure we've initially shown before is about 40% -- just under 40% of that CapEx is going into infrastructure, most of it's feeding AWS".

provider, the aggregate sum is $139 billion, at the higher bound it is $323 billion.

(53) With respect to Azure specifically:

(i) Microsoft has released more than 350 new and updated cloud products between 2017-2022, across a range of categories including analytics, databases, compute, storage and more.

(ii) Ofcom notes that Azure is chosen by customers because of, among other things, the number of features it offers (31%), the quality of service (39%), the level of security (34%) which as Ofcom notes are elements of innovation (which was not directly asked of customers in the survey).[49]

(iii) In 2022, Microsoft alone spent $24.5 billion on R&D and deployed more than 100,000 software engineers across the business.[50]

(iv) Some key recent examples of infrastructure innovations include:[51]

(a) New VMs containing ARM-based processor Ampere Altra, which provide high price-to-performance ratio for web services, databases, and microservices. Microsoft is also in the process of developing in-house chips for specific processes, such as the Athena chip to reduce costs for training generative AI models;

(b) Azure Premium SSD v2 Disk Storage provide storage enhancements for performance-critical workloads that consistently need sub-millisecond latency with high input-output operation throughput. This allows customers to increase storage performance without increasing storage capacity;

(c) Azure Cosmos DB now supports distributed PostgreSQL, providing faster and more scalable service for open-source relational databases; and

(d) GitHub Advanced Security for Azure DevOps enables developers to implement security earlier in the software development life cycle.

(54) Overall, there is no evidence that supports the view that because there are only three large IaaS players, there is a "lazy oligopoly" that innovates below competitive levels. On the contrary, the rate of innovation and spend summarised above makes cloud services one of the world's most innovation-heavy markets with vast and transformative customer benefits. The Interim Report is therefore currently unbalanced in the length it goes to paint these features as a source of concern – namely "*barriers to entry and expansion for smaller cloud players*"[52] – while failing to pause and give due recognition to the fact that this evidence cuts directly across the conclusion that the "*market is not working well*" for enterprise customers. On the contrary, the evidence seems almost incontrovertible that any "health check" [53] should be broadly positive.

[49] Q25 Customer Survey and see Ofcom comments at footnote 543 noting that *"(t)he top five reasons (in order) all arguably include elements of innovation: service quality; best value for money; supplier reputation; proposed level of security; and number of features*".

[50] Microsoft 2022 Annual Report; Welcome to the Engineering@Microsoft Blog.

[51] Keynote highlights from Microsoft Ignite 2022; Integrate security into your developer workflow with GitHub Advanced Security for Azure DevOps; Microsoft's AI Chip Strategy Reduces Costs and Nvidia Dependence.

[52] Interim Report, paras. 5.216ff.

[53] Cf. Consultation, para 3.15.

3.2.2 Evidence: efficiency gains and cost and price reductions

(55) Turning to price, as evidenced in the Interim Report,[54] access to cloud services have reduced customers' costs significantly. Cloud computing reduces total IT costs and allows companies to transform upfront capex into a smaller annual opex on cloud spend. As such, firms of all sizes can access computing resources affordably and focus efforts on business value-add efforts by outsourcing IT infrastructure. To give only a handful of examples from third-party research:

(i) Research from Gartner found that the TCO of migration to Cloud from on premises computing resulted in quarterly cost savings of 55% over a three-year period.[55]

(ii) OpsRamp, an IT operations management platform, conducted a survey in 2017 among IT decision makers: 94% of respondents said the cloud would reduce setup and maintenance costs, and 55% expected to reduce IT budgets by 30% with cloud infrastructure.[56]

(iii) A case study into the migration of an IT system of a Greek industry - from an in-house data centre to Google Cloud – was estimated to result in cost savings between 24% and 50%, depending on the solution.[57]

(iv) Research from the IDC sponsored by Microsoft found that Azure customers benefit from a 16% reduction in annualised infrastructure costs and a 37% reduction in three-year cost of operations, compared to on premises.[58] Similarly, an IDC study sponsored by Amazon estimated that AWS customers save 51% on operating cost than running on premises infrastructure.[59]

(56) While a degree of cost savings relative to on premises motivates cloud migration, competition between cloud providers has led to fierce price competition and price cuts over the past decade and a half – in other terms, vast rivalry-specific benefits.[60] There are many examples of competitive price reactions over the history of the industry.[61] For instance, in June 2011, not long after Microsoft Azure entered the cloud space, it announced the elimination of ingress fees.[62] As set out above, within a week of doing so, AWS made the same announcement,[63] matching Azure not only on the elimination but also on the start date of this new pricing.

(57) Aggressive price competition continues to the current day. In a recent report, analysts noted that "*On the pricing front, discounts are greater than 6-12 months ago.*"[64] Analysts also

---

[54] Context Consulting research report, Slide 28.

[55] Gartner, Is Public Cloud Cheaper Than Running Your Own Data Center?

[56] OpsRamp, Five Trends Reveal: The Emergence of Cloud-First Enterprises.

[57] Chatzithanasis, Georgios & Michalakelis, Christos. (2018). The Benefits of Cloud Computing: Evidence From Greece. International Journal of Technology Diffusion. 9. 61-73. 10.4018/IJTD.2018040104.

[58] IDC, The business value of migrating and modernising to Microsoft Azure.

[59] IDC: Fostering Business and Organizational Transformation to Generate Business Value with Amazon Web Services (awscloud.com); Fostering Business and Organizational Transformation to Generate Business Value with Amazon Web Services.

[60] AWS note, for example, that they have reduced prices 107 times between launching in 2006 and April 2021 (Amazon EC2 – 15 Years of Optimizing and Saving Your IT Costs).

[61] Price increases also occur in the industry across products. However, this is unlikely to be due to a lack of competition but more likely a result of quality improvements, lifting of experimental offers for innovative products or macroeconomic factors such as exchange rates. For example, Azure increased prices in Europe in 2023 to reflect currency fluctuations relative to USD and align itself to a "pricing model most common in (the) industry". This is referring to the fact that both AWS and GCP charge in USD for their products so when customers pay with foreign currency, the prevailing exchange rate is used (see Consistent global pricing for the Microsoft Cloud - Microsoft News Centre Europe).

[62] Microsoft offers free inbound data on Azure – Software – Cloud – iTnews.

[63] AWS Lowers its Pricing Again! – No Inbound Data Transfer Fees and Lower Outbound Data Transfer for All Services including Amazon CloudFront | AWS News Blog.

[64] UBS, 2Q21 Preview of AWS, Microsoft Azure and Google Cloud.

remarked that "*We are seeing more aggressive pricing from Microsoft to motivate large enterprises to commit to large Azure migration projects*."[65] This recent increase in price competition is not limited to the hyperscalers. In 2022, for instance, Oracle won an infrastructure deal against AWS in large part due to its aggressive discounting relative to AWS.[66]

(58) Last, but by no means least, reinforcing this general picture of intense price competition, empirical analysis of Azure's own data shows that prices have indeed come down, as discussed with Ofcom and detailed further in Section 6.

3.3 The story of rivalry behind these positive outcomes consists of the synergy of IaaS investment competition and PaaS innovation competition

(59) An important comment that Microsoft has on the Interim Report as a whole is that while it recognises early on the differences between IaaS and PaaS, these differences are frequently elided in the competitive assessment. It is, of course, useful to examine the cloud "stack" as a whole but the competitive landscape in PaaS is not the same as IaaS, and it is wrong to conflate the two if the focus is largely or only on "vertically-integrated hyperscalers" and that any player not active in both does not count.

(60) This section therefore examines IaaS and PaaS separately so that the competitive whole is understood as the sum of its parts, in which IaaS necessarily benefits PaaS and PaaS expansion drives greater demand for IaaS.

3.3.1 Competition in the IaaS "investment race" has been driven by intense rivalry, primarily between three major players

(i) Huge investments and scale economies

(61) Global cloud infrastructure is very expensive. Data centres cost billions of dollars to build and millions of dollars to operate and maintain. According to Dgtl Infra, it costs between $7-12 million per megawatt of commissioned IT to build a data centre.[67] In 2021 alone, around USD $54 billion was invested in data centre infrastructure – up from USD $24 billion in 2020.[68]

(62) Amazon, Microsoft and Google each have dozens of geographic regions and more than 350 data centres globally.[69] In Europe, the number of data centres continue to grow. In the last five years, around 240 data centres have been completed across Europe, and 32 new data centre projects are in the pipeline up to 2025.[70]

(63) When investing in cloud infrastructure, providers undertake substantial risk. They not only have to invest in and successfully execute their investments, but also need to select the appropriate efforts and optimise the level of investment amidst external uncertainty. The risk of making an incorrect R&D investment was borne out by the realised consequences of Microsoft's initial cloud offering as explained above.

(64) The upfront investment required to provide IaaS and the substantial economies of scale and scope tend to favour the large-scale players (i.e. hyperscalers) given their existing large-scale infrastructure. However, several players have achieved and sustained efficient scale. An

[65] UBS, 2Q21 Preview of AWS, Microsoft Azure and Google Cloud.

[66] UBS, 1Q23 Preview of AWS, Microsoft Azure and Google Cloud – Street Growth Expectations are Still too High.

[67] Dgtl Infra, How much does it cost to build a data center? (May 2022).

[68] Global investment in data centres more than doubled in 2021 with similar trajectory this year.

[69] Google Cloud's Data Center Locations: Regions and Availability Zones (Sep 2022), Amazon Web Services (AWS) Data Center Locations: Regions and Availability Zones (Jun 2022), Microsoft Azure's Data Center Locations: Regions and Availability Zones (Apr 2022).

[70] Savills Data Centre Advisory EMEA, European Data Centres: Deep dive in the data sphere.

economically efficient market structure for IaaS would not be expected to be fragmented across large numbers of competitors.

(ii) Scale benefits are passed on to customers and to PaaS providers

(65) Economies of scale emanate from, *inter alia*, decreasing cost of power, infrastructure labour costs, greater buying power, and multi-tenancy.[71] In particular, a single piece of hardware can host multiple tenants, and the cost of maintaining computing infrastructure is distributed across many servers.[72] According to Emerson, the difference in costs between a data centre of a maximum of 5,000 m² and a data centre exceeding 50,000 m² is more than 200% for both energy costs and operating costs.[73]

(66) Data centres aggregate demand and generate economies of scale by smoothing out demand variability across users and across time in order to maximise the utilisation of cloud resources. By pooling resources and demand into the public cloud, the variability of demand is diversified away, increasing the overall utilisation rate of the cloud resources, thus making it more efficient and cheaper to meet end-user demands.[74] The benefits to scale increase the upfront investment required to provide IaaS – as a consequence, the natural market structure for IaaS points away from fragmentation.

(67) The benefits of economies of scale have resulted in customers saving on their compute and storage costs, as discussed above.[75] It also highlights the scale associated with effective provision of cloud infrastructure services and the cost benefits to customers as providers grow. These requirements for scale reduce the likely number of providers in the IaaS market, compared to a product or service with a lower minimum efficient scale (however partnership and leasing provide a route for entry into cloud computing – e.g. Civo and OVHcloud have entered the data centre space through leasing and buying of colocation space).[76]

(68) Despite these high economies of scale, no single firm dominates: several players have achieved and sustained efficient scale, allowing them to remain competitive and pass on the benefits of scale efficiencies to consumers. Hyperscalers also pass on the benefits from scale economies in IaaS: as Ofcom acknowledges, ISVs have access to the IaaS layer of the hyperscalers' stack.

(69) As set out above, a market with several players that have achieved the required scale to achieve cost efficiency can result in intense competition among them when the products are relatively homogeneous, as is the case for IaaS compared to PaaS, albeit still not without significant innovation in IaaS (reflected in the development of innovative new hardware, for example AWS' investment in Graviton, Inferentia and Trainium chips to deliver better performance and enhanced capabilities[77] and Microsoft's investment high-performance computing to optimise AI workloads). This market structure has also ensured that most, if not all, cost efficiencies resulting from scale will be passed on to customers through the competitive process.

[71] The Economics of the Cloud (Nov 2010), Microsoft.

[72] The Economics of Cloud Computing, Microsoft, p.12.

[73] Emerson & Ponemon Institute, "Cost to Support Compute Capacity", 2016.

[74] Microsoft, The Economics of the Cloud.

[75] IaaS vs. PaaS—What is the Difference?

[76] Interim Report, para 5.192.

[77] Improve performance, lower costs, and enhance security with AWS silicon, AWS.

3.3.2 Dynamic competition in differentiated PaaS services

(i) Growth in verticalization and expansion of ISVs

(ii) In the early years of cloud computing, IaaS was the primary growth area. Cloud suppliers invested in IaaS capabilities such as hardware, security, privacy, compliance, data centre management, scalability, location, network options, storage tiers and other features to differentiate themselves. However, as workloads continued to migrate to the cloud, companies started thinking more strategically about their cloud applications and specifically how to take advantage of cloud architecture and new degrees of freedom and innovations enabled by this architecture. Customers evolved from leveraging basic compute and storage products to managing and deploying code rapidly in an agile manner – called a "lift and evolve" strategy.[78] Consequently, cloud providers "moved up the stack" to provide developer and management tools that leveraged their cloud infrastructure – i.e. "verticalization" along the stack.

Figure 2: Waves of public cloud adoption

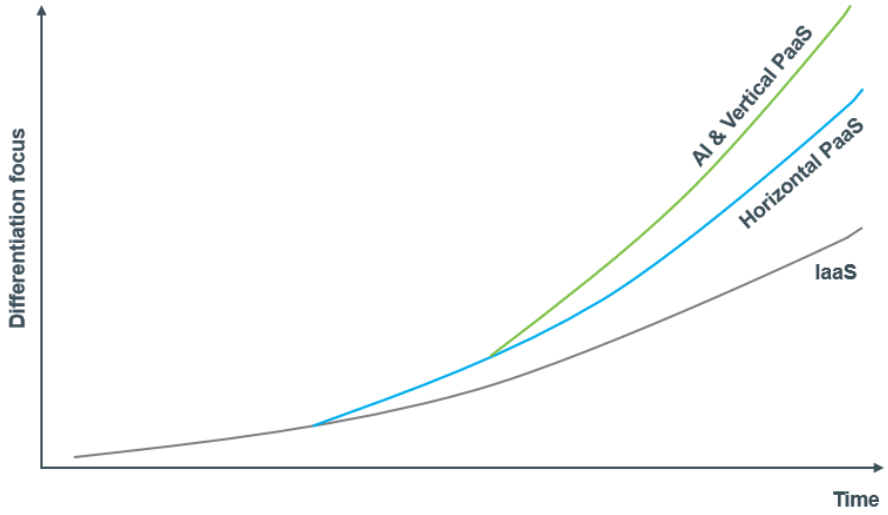

(70) Cloud providers worked to create new development PaaS capabilities to attract new customers that needed a more complete solution to drive adoption. Moving up the stack was considered a competitive advantage a few years after cloud was introduced. An analyst noted that "*the overall TAM is growing at such a rate that the IaaS market will support multiple 'winners' for the foreseeable future and longer-term we believe the bigger 'battle' will be which of the hyper-scale public cloud vendors can most effectively move 'up the stack' into the PaaS (platform-as-a-service) market.*"[79]

(71) In more recent years, cloud adoption has expanded to customers applying and customising cloud technologies to more specialised business requirements. Certain industries such as financial services also needed to address regulators' concerns about privacy and security in their IT infrastructures. As customer needs have become more heterogenous, it has opened opportunities for existing and new cloud providers to differentiate themselves in different industry and workload verticals, without the need for hyperscaler scale. Oracle, for instance,

[78] AWS re: Invent 2016: Lift and Evolve – Saving Money in the Cloud is Easy, Making Money Takes Help.

[79] Microsoft Corporation, Evercore ISI, 8 February 2018.

has focused on developing their cloud ERP products, Fusion, and NetSuite[80] and gene sequencing where it has won large customers from AWS.[81]

(72) This trend of verticalization has provided a manifestly viable entry and expansion path for smaller cloud providers, on the demand side, by enabling them to specialise and compete effectively in specific verticals. On the supply side, if one distinguishes between the entry costs of investment in hardware and software, hardware/bare metal costs for small data centres are unlikely to be prohibitive upfront costs if an entrant is only focusing on a certain vertical. Colocation is also an alternative to the high upfront costs of data centre investment. For instance, one reason Oracle's cloud spending is modest relative to some of the hyperscalers is because Oracle relies heavily on colocation partners for their cloud buildouts.[82] On the software side, set-up costs could be minimised by using open-source software such as Linux on servers. These two factors suggest that technological costs of entry may not be prohibitive for new entry. This is supported by evidence on recent entry of providers like CoreWeave, focusing on AI workloads (see Annex 1).

(ii) "Vertically-integrated" IaaS players have fostered PaaS competition from ISVs

(73) The use of the term "vertically integrated" in the context of the hyperscalers can raise the implication of a concern raised in some consumer-facing digital markets that the platform owner (here, by analogy the IaaS platform) may engage in "self-preferencing" to favour its own applications and foreclose (at least partially) rival applications from third parties (e.g. PaaS ISVs).

(74) The opposite of any such notion is true in PaaS. IaaS players have demonstrated in their actions that they benefit if they adopt a strategy in which an array of value-add and differentiated niche applications work well with their infrastructure, in order to drive volume to that infrastructure as they compete for volume via new workloads. Of course, owning IaaS does not mean ceasing to innovate in PaaS or surrendering that market to third parties either – on the contrary, the pace and scale of innovation also by AWS, Azure and GCP is centred on differentiated PaaS services, on top of innovations in IaaS. This competition similarly encourages PaaS-only providers to continue to innovate and to offer ever-improving value added services beyond what is provided by existing services.

(75) But the PaaS market is so fast-growing that there is ample scope for and evidence of innovation by hyperscalers, other IaaS players, and pure-play PaaS/SaaS vendor alike. For PaaS and similar cloud services, scale is less relevant because competition centres on tailored service offerings that meet the unique technical and business requirements of each client. Smaller players can compete on an equal footing with larger players. Going forward, it is increasingly likely that smaller players will enhance their competitive position, as multi-homing becomes more prevalent and cloud-interoperable solutions become more important and attractive.[83]

(76) The net effect of this dynamic interplay has resulted in a vibrant and competitive landscape of providers at the PaaS layer, enabled by the significant investments that have been made at the IaaS layer and fostered by the encouragement of the infrastructure players as they compete to

[80] Oracle 2019 Q1 Earnings Call.

[81] Oxford Nanopore, a large gene sequencing company is moving from AWS to OCI where they are going to store gene sequencing, but not only just storing sequences in OCI (see Oracle 2023 Q2 Earnings Call).

[82] Cloud Adoption Creating a Land Grab in the Data Centre Market - Ecosystm Insights (ecosystm360.com).

[83] In the following, and in accordance with the approach taken by Ofcom, we focus our assessment on IaaS and PaaS as the core of cloud business services. SaaS, on the other hand, which while also hosted in the cloud is much more stand-alone service category, is focused more on consumer operability than back-end functionality.

make their cloud environment the most attractive to the most number and diversity of enterprises and "industry verticals".

3.4 Conclusion on competitive market outcomes: the market is working well

(77) Microsoft therefore urges Ofcom to give central prominence to today's market outcomes, both non-price and price outcomes, in reaching a diagnosis on whether, in overall summary terms, the market is "*working well*" or "*not working well*". The evidence appears incontrovertible that the market is working well, and the explanation lies in the overall intensity of cloud competition. These facts are based on direct outcomes evidence, and should be decisive where, as is the case here, such evidence is demonstrably available – and not on whether only three players deserve the arguably negative connotations of the "*hyperscaler*" label or based on other indirect indicia.

(78) The Consultation also refers to the competitive "*health*" of the market.[84] Taking this as analogy, in the context of the resource intensity and cost of an 18-month+ MIR, the perceived competition "deficiency" (scale and gravity of alleged AEC) and, conversely, the candidate remedial intervention "upside" of an MIR can only robustly be undertaken by first declaring whether the vital signs of the patient on the operating table show that they are "*working well*" (the patient is fundamentally "fit, strong and healthy") or "*not working well*" (the patient is fundamentally "weak and ill"). In terms of side effects and risks, different choices flow in medicine from whether to makes sense to intervene heavily when a patient is fundamentally healthy.

(79) Having set the appropriate context of marketplace outcomes and rivalry, this Response now directly address Ofcom's over-arching concerns in Part B and specific concerns in Part C.

[84] Ofcom's Consultation Document, para 3.15.

Part B: Ofcom's central hypothesis of "lock in" exploitation is misplaced

(80) The central contention of the Interim Report is that the market is "*not working well*" and where "*effective competition*" is "*inhibited*" because of a core dichotomy: rivalry is "*predominantly... to attract [new] customers*" on the one hand, while there is "*limited scope for workload competition*" for the workloads of existing cloud customers, on the other.[85]

(81) This is said to arise because for any new or additional IT needs, customers are very likely to select the current provider: *"[m]any decision-makers acknowledged that a de facto lock-in exists*".[86]

(82) This section explains why Ofcom's central contention is fundamentally misplaced. Of course, customers are correct that there are always and inevitably degrees of de facto lock-in: it is a basic law of IT solutions that this is the case and is simply not controversial (see Section 4 below). However, Microsoft vigorously disputes that this lock in:

(a) must be caused, at least in substantial part, by supplier practices that aggravate and reinforce this; and

(b) must result in market power over existing customers and rent extraction/exploitation through paying higher prices than if they were "new" customers, a reduced ability to switch for better innovations and so forth.

(83) As to (a), the first is evident in part from the survey itself. When it comes to the specific assessment of interoperability and pricing practices, Ofcom should not lose sight of the fact, as the Interim Report sometimes does, that customers say that the root cause of the de facto lock-in *"… was often perceived to be a function of internal factors rather than provider-imposed restrictions*".[87]

(84) As to (b), there is no empirical evidence of Azure holding significant market power over existing customers such that this lock-in is exploited – indeed, the evidence points the other way. As will be explained (see Section 6), Azure's existing "locked-in" customers do not have it worse off than "new" customers. In a nutshell, the prime reason is that there is competition for the next workload, regardless of the maturity of the cloud customer or workload in question (see Section 5). Consequently, the Interim Report's narrative of weak competition and market power (relying on indicia evidence other than outcomes evidence) is wrong (Section 7).

4 All IT solutions markets exhibit degrees of "lock-in" and cloud is no exception

(85) "Lock-in" refers to the situation where a customer is unable (easily) to switch from one cloud service to another (either on the same platform or to a different cloud provider). In the Interim Report, Ofcom conceptualises the term 'lock-in' broadly to encompass a range of scenarios including technical barriers to switching and a mere unwillingness to switch.

(86) This is not unique to cloud computing but is a characteristic of all IT solutions. For example, where a customer chooses to use Linux over Windows, they will be "locked-in" to the Linux platform. As the Gartner Report notes:[88]

[85] Interim Report, paras 6.13 and headings on pp. 178-179 (paras 6.5-6.6 and 6.8-6.9).

[86] Interim Report Annex, para A7.39.

[87] Interim Report Annex, para A7.39.

[88] Gartner, A Guidance Framework for Managing Vendor Lock-In Risks in Cloud IaaS

In general, IT systems obey the principle of conservation of lock-in — the amount of lock-in in the overall system remains approximately the same regardless of the architecture; however, the location of the lock-in will differ, based on the architecture.

The goal of your (cloud) architecture should not be to eliminate lock-in — that's impossible. Instead, your goal should be to identify points of lock-in, so that related risks can be appropriately addressed and maximal value is obtained in return for the lock-in that you choose to accept.

Thus, all possible technology choices involve some sort of lock-in. Attempting to avoid lock-in can limit the variety and quality of the technology available, increase risk, and increase projects' financial and technical overhead.

(87) Gartner[89] elsewhere identifies other sources of lock-in such as employee skills, processes and tools, and technical capabilities. These factors apply in a range of IT solutions and exist (and are often more significant) with on premises computing as well.[90] Moreover, lock-in can happen at a more granular level, such as product lock-in or version lock-in.[91]

(88) Lock-in in cloud specifically may arise due to technical barriers where there are differences in coding language between cloud environments, customer's dependencies on proprietary cloud APIs which arise from differentiated features and functions, and data gravity. In each case, the customer cannot switch without incurring some cost or some loss of functionality.

(89) In such cases, the customer makes a decision that the cloud solution costs and/or performance/functionality gains are incremental, and not outweighed by the costs of switching to another cloud service – i.e. the switching costs are too high. The costs of switching a workload, whether on premises or between cloud providers, depend on the complexity of the application (i.e., the number of services used), the data volume, the types of services being used (e.g., vendor-specific, cloud agnostic, containerised solutions) as well as the underlying architecture (i.e. monolithic vs. microservices).

(90) Cloud customers are sophisticated buyers with specialised knowledge, procurement teams and/or resources, and are sometimes advised by third-party IT consultants. Customers deliberately evaluate vendor dependence and lock-in while designing their cloud architecture and choosing services. Many such risk evaluations also discuss the benefits of choosing vendor specific solutions[92]. Customers also design their own cloud solutions, and test and incrementally scale their commitment to a specific cloud infrastructure.

(91) Overall, establishing "existing customer lock-in" in cloud – where this is an endemic feature of all IT solutions markets – is not profound and does not advance the analysis of perceived competition problems and solutions. Before showing that Azure existing customers are not

---

[89] Gartner, Cloud Governance Best Practices: Managing Vendor Lock-In Risks in Public Cloud IaaS and PaaS.

[90] For example, Gartner note that avoiding lock-in by building internally locks one into the internal IT team, which may pose greater risks. Gartner, Cloud Governance Best Practices: Managing Vendor Lock-In Risks in Public Cloud IaaS and PaaS. See also: Embrace lock-in in the Age of Cloud | Accenture.

[91] Martin Fowler, Don't get locked up into avoiding lock-in.

[92] Gartner advise that cloud provider customers identify points of lock-in and assess the related risks and impacts, such as business continuity (cloud provider outages and security breaches), governance challenges (multi-cloud management), and future proofing (access to innovation, broad choice of technologies and skilled employees). For example, regarding technological innovation, customers should consider the likelihood that their provider may not be able to rapidly deliver best-in class capabilities and weigh this against the costs of switching and benefits they receive from their provider. Source: Gartner, Cloud Governance Best Practices: Managing Vendor Lock-In Risks in Public Cloud IaaS and PaaS. *See also*: Getting the most from cloud services and containers | McKinsey; Understanding The Potential Impact Of Vendor Lock-In On Your Business (forbes.com); The cloud's lock-in vs. agility trade-off | Deloitte; Why a "Cloud Exit Strategy" is essential to enable the future | KPMG.

exploited relative to new customers (evidenced in Section **6**) it is important, however, to explain as to why, *notwithstanding lock-in*, this should not be a surprising result.

5 Competition for the "next workload" is enduring – and explains the lack of Azure's exploitation of customer lock-in

(92) Competition in cloud is about driving volume to utilise capacity to the maximum – this means winning as many customer workloads as possible. Azure is not choosy about the "origin" or "maturity" of the workload – all volume is good volume and all revenue is good revenue – and does not adapt its strategy or pricing based on the maturity of the cloud customer or workload in question.

(93) In practice, the supposed dichotomy between "**new"** workloads of customers just migrating to the cloud (Ofcom says: strong competition for these) and "**old"** workloads of existing cloud customers (Ofcom says: weak competition for these) merges into competition for the **next** workload. However, for ease of exposition this section starts with the sheer volumes of "new" or "inaugural" workloads that cannot be construed as the roll-over of "existing" cloud workloads.

5.1 The sheer size of the cloud migration opportunity drives competition

(94) Gartner forecasts global spending on public cloud services in 2023 will be around £477.7 billion, up from £395.7 billion in 2022.[93] While this annual forecast growth is high at 20.7%, the 2023 figure still represents only 13% of total IT spending projected to total £3.6 trillion in 2023.[94] As depicted in Figure 3, this leaves 87%, or £3.1 trillion, of non-cloud IT spend. Reports indicate that between 50%-75% of all IT workloads are suited to be migrated to the public cloud.[95] According to Gartner, by 2026 cloud is expected to account for 17% of total IT spend. Overall, public cloud spending is expected to reach $600 billion in 2023,[96] $1 trillion in 2026[97] and $2 trillion by 2028 globally.[98]

Figure 3: Gartner 2023 forecasts cloud and non-cloud IT spending

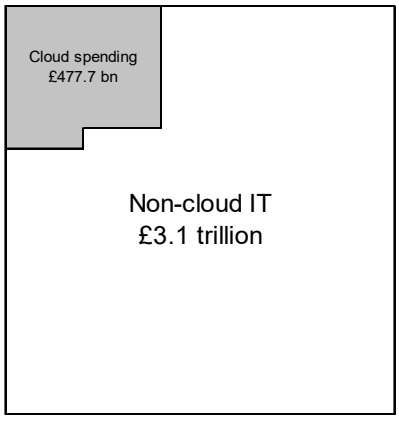

Source: Gartner data cited by Ofcom Interim Report

[93] Interim Report, para 3.7.

[94] Interim Report, para 3.7.

[95] The case for a $2 trillion addressable public cloud market (iot-analytics.com).

[96] Gartner Forecasts Worldwide Public Cloud End-User Spending to Reach Nearly $600 Billion in 2023.

[97] Public Cloud Market To Hit $1 Trillion By 2026: Here's Why | CRN.

[98] The case for a $2 trillion addressable public cloud market (iot-analytics.com).

(95) The biggest limitation to moving to the cloud or switching between clouds is the capacity that exists within an organisation's IT department. Given limited overall bandwidth, organisations deploy this scarce resource in ways to get the best overall return on investment for the organization. At this stage of the industry's development, this normally means moving high cost on premises solutions to much lower cost cloud providers.

(96) Moreover, as Figure 4 below shows, [✂].

Figure 4: Average monthly revenue generated by Microsoft Azure customers by contract

[✂]

(97) In sum, therefore, competition between suppliers for the next cloud workload will dominate irrespective of perceived "low" churn or switching rates of intra-cloud churn, or "existing/old/mature" cloud workloads. This is because the sheer scale of the opportunity for new business dwarfs the size of the opportunity, all else being equal, for switching business between suppliers.

(98) Microsoft's focus on new workloads for customers is illustrated by its investment programmes, where Microsoft compensates users (at least partially) for the costs they incur in migrating a workload to Azure to best take advantage of the services that Azure has to offer. [✂]. This is also clear from the current pipeline of workloads that Microsoft has contracted with various customers. In at least [✂] of the workloads currently under negotiation, Microsoft is aware of at least one primary direct competitor also attempting to obtain the workload. In more than [✂] of workloads, Microsoft believes it is directly competing for the workload [✂]. Importantly, these shares are, by definition, lower bounds.[99]

5.2 Competition for "new" or "pipeline" workloads is strong even for existing cloud customers and for dynamic opportunities as to "old" existing cloud workloads

(99) Massive competition to attract businesses to move their IT needs into the cloud may be undeniable, but it is not "turned off" as soon as Azure discerns that it already has a relationship with the customer. The Azure business does not recognise such a dichotomy and the data (see Section 6) bears this out. In other words, competition is not limited to attracting organisations to migrate to the cloud, but rather to win new workloads, irrespective of the client's past cloud purchases, not least given multi-cloud as the new paradigm (see Section 5.3) which means that they already at least in part familiar with another rival's cloud environment and that that rival is also an "incumbent" supplier.

(100) Nor is competition turned off for an already-serviced workload of a cloud customer. Even customers that have migrated to the cloud do not have static or stable demand for unchanging cloud products and services. Historical IT demand, even on premises, is in a perpetual state of upgrades, replacements, and renewal as: (i) the overall IT sector releases new innovations such as AI; and (ii) individual enterprises and sectors undergo change, most prominently in the move within all industries transforming their business processes and operations, in at least part, to become "digital" and to collect and to utilise more and more data. Companies will generate new demand because the tools necessary to how the business runs will need to change to stay competitive. In the same way that companies historically needed to replace existing IT infrastructure as technology advanced, they will need to update and create new workflows for data that has already migrated to the cloud.

[99] As it reflects only the imperfect competitive intelligence information available to Microsoft.

(101) In other words, even "migrated" cloud workloads are likely to continue to expand and modernise as technological capabilities improve, generating new contestable demand.[100] Such modernisation provides another plane of competition for the next "upgraded" contestable workload between cloud providers. For instance, investments are being aggressively made in powerful supercomputers capable of quintillions of operations per second.[101] Moreover, competition between applications will also mean cloud solutions are not static, and customers are likely to "shop around" for the best solution.

(102) Further, new future cloud-native workloads are also likely to drive competition for workloads. To take just one topical example, the growth of AI will drive additional innovation in cloud infrastructure as well as services. AI has enabled new use cases, AI applications and AI tools in almost every industry vertical, as well as a race to infuse existing applications with AI tools. IDC estimates that by 2026, 85% of enterprises will combine human expertise with AI, ML, natural language processing (NLP), and pattern recognition to help improve productivity.[102] IDC forecasts that by 2025, nearly 50% of all accelerated infrastructure for performance-intensive computing (including AI and HPC) will be cloud-based.[103] AI can also be used to improve cloud performance and the efficiency of cloud computing. AI-powered data management tools can assist organisations in effectively managing and utilising their data; lowering the cost of cloud storage;[104] or optimise cloud performance.[105]

(103) Customer workload demand is therefore dynamic over time and customers do not become "inert" simply because they have migrated some (or even "all" conceivable) workloads to the cloud, particularly as Azure's customers are often sophisticated buyers who negotiate their own terms of purchase (see Section 8). It may make no sense for a customer to switch a given workload from one cloud provider to another if everything is working well, because the incremental gains do not outweigh the costs. But it does not follow that all of a customer's workloads already in the cloud are "locked in" and non-contestable over time. That assumes the underlying IT demand of the enterprise is fixed for a series of workloads. This generalised proposition is incorrect.

5.3 Multi-cloud is now the competitive frontier, driving ever greater competition

(104) Ofcom's survey found that half of the respondents already multi-cloud and a vast majority (86%) of respondents not currently using multi-cloud are open to using multi-cloud in the future[106]. This growing adoption of multi-cloud strategies is driven by several factors, most notably, data sovereignty, cost optimisation and avoiding vendor lock-in. Third-party evidence supports Microsoft's view that multi-cloud is now "*the de facto standard among organisations*".[107] For example:

(i) the **Flexera 2023** State of the Cloud report suggests that 87% of companies and organisations they surveyed (n=750) were using multi-cloud in 2022 compared to 89%

[100] Industry reports indicate that companies that adopted public cloud early (in 2014 and 2015) are undergoing modernization cycles to upgrade to container-based and serverless architectures to compete with more modern cloud architectures.

[101] Introducing Meta's Next-Gen AI Supercomputer | Meta (fb.com).

[102] 40-plus cloud computing stats and trends to know in 2023 | Google Cloud Blog.

[103] 'Do more with less': Why public cloud services are key for AI and HPC in an uncertain 2023 | VentureBeat.

[104] Amazon's Glacier storage service, for example, uses machine learning to identify and remove duplicate data, which can reduce storage costs by up to 50%.Top 5 Ways Artificial Intelligence Impacts Cloud Computing (analyticsinsight.net).

[105] AI can also automatically schedule and manage cloud resources, optimize workloads, and make recommendations to improve cloud performance (see Top 5 Ways Artificial Intelligence Impacts Cloud Computing (analyticsinsight.net)).

[106] Interim Report Annex, para A7.7. See also Annex 3 for further discussion on such conclusions drawn by Ofcom from its survey.

[107] Cloud Computing Trends: Flexera 2022 State of the Cloud Report; Cloud Computing Trends: Flexera 2022 State of the Cloud Report.

in 2021 (out of which 72% were a hybrid cloud[108] and over half were using multiple public clouds, 2% were multiple private clouds and 13% were multiple public clouds).[109]

(ii) A 2023 report commissioned by **Oracle** suggests that 98% of respondents (n=1500) were using (or planning to use within six months) more than one Infrastructure as a Service (IaaS)/Platform as a Service (PaaS) provider.[110]

Figure 5: Use cases for multi-cloud[111]

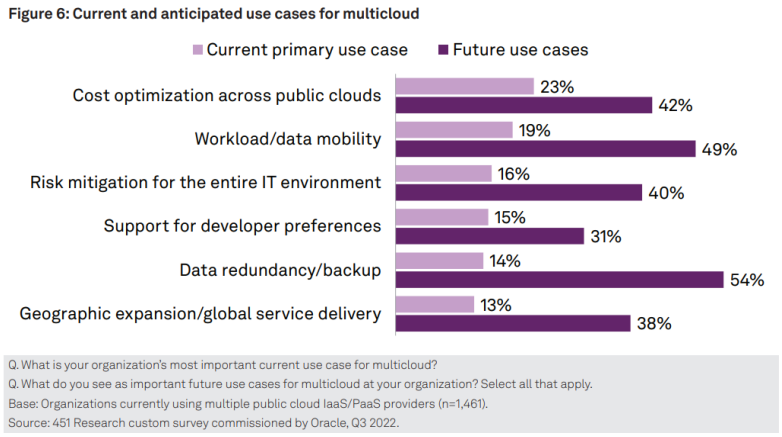

(iii) **Foundry's Cloud Computing Study 2022** reports that in their sample, 16% of organisations rely on a single cloud provider for their public cloud deployments (19% for small and medium-size businesses and 13% for "enterprises", i.e. respondents with more than 1,000 employees).[112]

(105) Meeting customer demand for multi-cloud architectures is a focus area for Azure. All major providers have offerings that simplify the management of multi-cloud infrastructures. These have been developed and offered by hyperscalers such as AWS Systems Manager, Azure Arc, Google Anthos, as well as other vendors such as Scalr, Flexera Cloud Management Platform, IBM Multicloud Manager, and Nutanix Beam.[113] For example, C3 AI markets itself as an Enterprise AI software provider for accelerating digital transformation. It offers vertical-focused services in higher layers of the stack, while operating across multiple public cloud layers. They have a multi-cloud focus, with their key product running on Azure, AWS, and GCP, private clouds and customers' data centres.

(106) ISVs, such as Snowflake.ai, also offer solutions designed to work across multiple cloud providers. ISVs typically offer cross-cloud compatibility and increasing levels of "abstraction" from the underlying infrastructure. Other tools that facilitate multi-cloud include containers, Kubernetes, and cloud-to-cloud migration tools. Such tools are provided not just by cloud providers but by several other companies.[114] This abstraction allows them to sell to customers who already have implementations on any of the large IaaS providers and enables them to keep their customers should those customers choose to switch to a different IaaS provider. These

[108] Flexera describes hybrid as a combination of public and private clouds.

[109] The Flexera report suggests that apps siloed on different clouds and failover between clouds are the top two multi-cloud implementations. See Flexera: State of the Cloud Report 2023.pdf Flexera: State of the Cloud Report 2023.

[110] S&P Global Discovery Report: Multi-cloud in the Mainstream (oracle.com).

[111] S&P Global Discovery Report: Multi-cloud in the Mainstream (oracle.com) S&P Global Discovery Report: Multi-cloud in the Mainstream (oracle.com).

[112] Cloud computing study 2022: Foundry Cloud computing study 2022: Foundry.

[113] Best Cloud Management (CM) Platforms.

[114] Cloud Management Tooling Reviews 2023 | Gartner Peer Insights; and Annex 4 for further details.

platforms provide unique services and capabilities to their end-customers, often driven by improvements from working hand-in-hand with their customers, making it difficult for hyperscalers to displace them.

(107) Cloud providers are increasingly partnering with software developers and ISVs to provide customers with broader access to advanced platform services.[115] Analysts noted that GCP's partnerships for "up the stack" offerings demonstrated its reliance on partnerships to compete with other hyperscalers.[116]

(108) Further, as shown in Annex 5, there are many concrete examples of UK and European companies using more than one cloud provider (among AWS, GCP, Azure, OCI and others).

5.4 Conclusion on competition for the next workload

(109) Cloud is ultimately a volume driven business, given very high fixed costs and low marginal costs. The sheer size of the cloud migration opportunity drives "inaugural" new workload competition, but competition is not turned down when the customer has a cloud provider relationship ready (not least as multi-cloud means it will have or be developing more than one such relationship). As such, there is scope for competition for an existing workload at the point of renewal, replacement, and upgrade (an endemic feature of IT procurement). This means that a sizeable share of "old" or "existing" cloud workloads are in principle contestable and subject to constant review if the current technology becomes outdated or the switching benefits otherwise outweigh the switching costs.

(110) While the above is how the Azure business firmly believes it considers it is competing, it is important for robustness to test this business perception against hard data, which is summarised next.

6 Pricing evidence contradicts the exploitation hypothesis

(111) Microsoft has instructed the Brattle Group ("**Brattle**") to analyse Azure pricing data (on an empirical and a conceptual basis), as discussed with Ofcom on 19 May 2023. The following section summarises this analysis, while Annex 6 sets out the full analysis.

(112) At Microsoft's request, Brattle analysed Microsoft's actual pricing to consider whether prices increased upon contract renewal (which might be indicative of price effects due to lock-in). Brattle concludes that this is not the case.

(113) [✂]. Moreover, website prices tend to fall, especially in the first years after products have been introduced. The facts that website prices have been declining and that discounts do not depend on the purchasing history necessarily means that prices have been decreasing across the board for small customers.

(114) Second, [✂]. How prices change upon contract renewal is for these customers hence an empirical question. Looking at the data, Brattle finds that for such customers discounts have predominantly increased upon contract renewal, and prices have predominantly decreased.

6.1 Microsoft's pricing practices are not designed to lock customers in

(115) [✂].

[115] See for example, Azure's Differentiate with Solutions Partner designations.

[116] "Materne's Software Catalyst: Software Takeaways from Google Cloud Next '21", Evercore ISI, 16 October 2021.

(116) While certain customer groups may receive discounts, or benefit from "Price Protection" (see below), the starting point for all customer groups are the published website prices. Website prices for each cloud services are typically set as follows:

(i) For **established products**, [✂].

(ii) For **new product categories** [✂].

(117) Microsoft's uniform pricing structure contradicts the notion of customers being subject to a lock-in effect, given that prices paid by existing customers are equal to prices paid by new customers. Importantly:

(i) all "pay as you go" ("**PAYG**") customers acquire services monthly and pay the same prices, independent of their history of using Azure services.

(ii) [✂].

(iii) [✂].

(118) These pricing policies are inconsistent with the notion that Microsoft seeks to exploit locked-in customers. If Microsoft Azure were free of competitive constraints for the "next workload" and its pricing designed to in some way exploit customer lock in, it would not compete aggressively on price with its competitors. If this were the case, the prices for existing customers would be considerably higher than the prices offered to new customers or workloads.

6.2 Website pricing over time does not support the lock-in exploitation hypothesis

(119) Brattle's analysis indicates that Azure's prices for both IaaS and PaaS typically decrease over the lifecycle of the cloud products.

(120) For IaaS products, Azure website prices may [✂].

Figure 6: Revenue weighted average price index and number of price changes (IaaS)

[✂]

(121) For PaaS products, Azure website prices tend to [✂].

Figure 7: Revenue weighted average price index and number of price changes (PaaS)

[✂]

(122) The data show that the vast majority of Azure products have actually decreased in price over time, suggesting no evidence of a lock-in effect.

(123) Moreover, Microsoft prices its Azure products centrally in the United States. Shortly before the Ofcom survey was conducted in March 2023,[117] Microsoft adjusted the prices for Azure products in the UK in order to reflect the loss in value of the GBP-USD exchange rate, resulting in a nominal price increase in GBP for Azure products in the UK. Such a price increase is reflecting of external factors, not of price lock-in of existing Azure customers by Microsoft. As set out above, cloud services are characterised by significant product heterogeneity (in particular, in relation to PaaS) and Microsoft is constantly introducing new and improved products to address the expanding needs of customers. Many of the new products will be close substitutes for already existing products, but with a greater performance (e.g. higher computing power) or

[117] The latest adjustment to Azure products came into effect on 1 April 2023, with prices being communicated to customers in January 2023.

additional product differentiation (e.g. new features or innovations). As customers can choose more effective products, or products fitting more precisely to their requirements, the introduction of new products can be understood as an implicit quality-adjusted price decrease (which is not accounted for in the above analysis). While the expansion of cloud service consumption and spending may explain customers' sentiment that the cost of such service tends to increase over time, what matters for the analysis is whether prices for identical or similar products have increased.

6.3 Individually-negotiated prices do not support the lock-in exploitation hypothesis

(124) Survey evidence has formed the basis for a suggestion in the Interim Report that Azure customers who experienced a price increase when they renegotiated their agreement had an average price rise of 19%.[118]

(125) This appears to mischaracterise the findings of the survey. Ofcom asked respondents to confirm if they experienced a price rise for "*some or all services"* and then to quantify the increase**.** The reported increase may therefore apply to a single service rather than to the totality of services a customer acquires on Azure.

(126) Brattle has, on Microsoft's instruction, investigated whether the hypothesis that Microsoft lowered the discount for existing MACC customers is correct by comparing how discounts and prices have evolved before and after a contract renegotiation.

(127) First, the actual experience of MACC customers shows that the discount level remains stable or increases on renewal. For example, the discount for [✂]increased from [✂], and then to [✂], and for [✂], the discount increased from [✂]. The data show that only [✂] of renegotiations led to a customer having their discount decreased. [✂].

(128) In Figure 8 below, the blue bar shows the discount percentage in the most recent contract of the Azure customer in the UK, the red diamonds and crosses show discounts in preceding contracts. The figure reveals that discounts have [✂].

Figure 8: Discount levels of UK customers with successive contracts

[✂]

(129) Second, the effective prices paid for individual products by customers on a MACC tend to be lower than under their previous contract. According to Brattle's analysis of the Azure Consumed Revenue data ("**ACR**") (which provides client-level monthly revenues and consumed units for individual products), IaaS product-customer combinations representing [✂] of revenues and PaaS product-customer combinations representing [✂] of revenues are lower on renewal.[119]

(130) In summary, there is no indication that Azure customers renegotiating contracts experienced price increases (and even more so not by the 20% average level claimed by Ofcom in the Interim Report).[120] The anecdotal survey results therefore lack the appropriate factual basis and may be driven by behavioural biases, e.g. customers interpreting increases in their expenditure (driven by higher consumption) as price increases in the context of the survey.

[118] Interim Report, para 4.33.

[119] [✂]. See Annex 6 for further details on Brattle's methodology and analysis.

[120] Interim Report, para 1.30. For further detail on Brattle's pricing analysis and findings, please refer to Annex 6.

7 The narrative of "weak" competition signified by high concentration, "low" switching, and "high" profitability is at odds with the best available evidence

(131) As demonstrated above, Azure does not exploit "locked in" customers while it competes for new ones, not least because this dichotomy is false. IT vendor/customer relationships endure are over time and are not "one shot games". As set out above, Azure is perpetually competing for the next workload as evidenced by the pricing analysis above. In lieu of analysis of actual demand-side i.e. customer outcomes, the Interim Report focuses on inferences from some (but not all) survey results and less direct supply-side indicia, which are now discussed below.[121]

7.1 High concentration

(132) The Interim Report and Consultation in the key introductory and summary paragraphs paint a picture of stable high concentration and market share stability for the "hyperscalers":

> *"the hyperscalers account for around 70-80% of total UK revenues generated from IaaS and PaaS … AWS and Microsoft Azure account for around 60% to 70% in 2021, with Google at … 5% to 10%, a ratio that has remained broadly stable since 2019. These market shares indicate that by far the majority of customers … will use hyperscaler services …"* [122]

(133) As demonstrated above, while high concentration may be a good reason for scrutiny in certain markets, in this case the evidence on market outcomes does not in and of itself raise concerns given the high fixed costs and need for scale in the global IaaS market.

(134) Microsoft also notes the following:

(i) UK shares of supply by revenue are not meaningful market shares. The supply side of the market is global, not national. Of course, UK shares of supply may be relevant as a comparison with global market shares, but it is not particularly meaningful for *substantive* purposes to zoom straight into the UK (notwithstanding that *jurisdictionally*, the UK is the relevant frame of reference and the obvious focus).[123]

(ii) The label of "stable" shares is technically correct for the three-year period 2019-21 but to the extent this serves a narrative of stability, this snapshot is incorrect for all the reasons outlined above in this Response.

(iii) IaaS and PaaS are distinct sets of services with different considerations as to scale, entry and expansion (see Section 3.3) and shares data and analysis needs to consider them separately.

7.2 "Low" switching

(135) "Higher" switching is not an end in and of itself. Low switching levels can in the right circumstances serve as an indicator of a lack of competitive intensity, for example if barriers to switching permit supplier exploitation of (or market power over) captive or inert consumers. However, exploitation should be measured with respect to outcomes, not switching levels.

(136) This is not least the case where the vague allegation that the switching level is "low" has no relative comparison as to what switching level should be ideal or expected. It is difficult to make

[121] See also Annex 3 for further discussion on the methodology and conclusions drawn by Ofcom from its survey.

[122] Ofcom's Consultation Document, paras. 3.17-18; cf. also Interim Report paras. 1.6, 3.47 et al.

[123] For the sake of completeness, market shares for IaaS and PaaS services in the UK and worldwide are provided at Annex 7.

a meaningful assessment of supposedly "high" switching cost without asking the question: high relative to what? It would be instructive to compare (and ask stakeholders to confirm) that, relatively speaking, it is generally easier for customers to move between clouds than it is to move from on premises IT to cloud in the first place, or how the on premises IT industry worked before the cloud. All the evidence of which Microsoft is aware points to the fact that, all else being equal, cloud services have made switching less difficult than in other IT contexts.

(137) For instance, if customers are receiving a good deal because of ongoing competition for the next workload, a "low" level of switching is to be expected, because the gains from switching (an incrementally better deal where competitive prices are all similar) do not outweigh the cost of switching (actual costs, risks). As already shown, the clear evidence on market outcomes shows that the exploitation of enterprise customers is not made out, at least in relation to Azure.

(138) Even if it were correct to label switching levels as low without a benchmark, low switching, therefore, is not indicative of weak competition in the cloud computing market. It is likely the consequence of customers who are carefully weighing up the switching costs that are inherent against the benefits they receive from their cloud provider in a differentiated market.

7.2.1 This is a dynamic enterprise market where competition policy orthodoxy from certain other markets cannot be read across

(139) Implicit analogies to other certain markets subject to intensive competition scrutiny are misleading and unhelpful.

(i) **Not an end-consumer retail market.** Cloud is characterised by a high percentage (in revenue terms) of sophisticated enterprise customers (for Azure this includes [✂], amongst many others), often with dedicated staff with specialised knowledge and resources (including access to third-party IT consultants with expertise in choosing and building cloud solutions – including Deloitte and Capgemini), which are more likely to have the ability to switch cloud infrastructure providers if it makes sense to do so. There are no vulnerable or inert customers, which is the paradigm "loyalty penalty" case for competition for new customers alongside market power over (or exploitation of) existing customers.[124] This is a wholesale market where longer-run contracts may be expected.[125]

(ii) **Not a mature and stable market.** The cloud market is also not comparable to mature and stable markets such as retail banking, where there may be concerns around long run incumbency,[126] or Ofcom's other regulated sectors – such as mobile telephony and broadband internet – where there now exists a very low proportion of new customers and opportunities entering the market for the first time relative to existing customers active in the market.[127]

[124] Cf vulnerable customers who were the focus of the CMA's attention in a number of other market studies/investigations. Retail banking market investigation final report (publishing.service.gov.uk), para 67. Energy market investigation: Final report (publishing.service.gov.uk), para 276. Retail banking market investigation final report (publishing.service.gov.uk), para 67. Energy market investigation: Final report (publishing.service.gov.uk), para 276. As regards lack of internet access, see para 9.161.

[125] See for instance Contracts for mobile and fixed wholesale access for mobile virtual network operators (MVNOs), like Virgin, or Sky for fixed infrastructure from BT, run for 5yrs.

[126] For instance, the CMA's *Retail Banking* market investigation also found that a reason for low switching in retail banking was that that personal current accounts are relatively low-cost products which customers view as little differentiated and, as such, many customers believed that there is not much to be gained from switching. Retail banking market investigation final report (publishing.service.gov.uk), para 6.57.Retail banking market investigation final report (publishing.service.gov.uk), para 6.57.

[127] For example, the number of new mobile phone subscriptions in the UK was only 2.3% in 2021 (relative to the total number of existing mobile subscriptions in the UK) (Communications Market Report 2022 – Interactive data - Ofcom), and the number of new fixed broadband connections was only 1.5% in 2021 (relative to the total number of existing fixed broadband connections in the UK) (Communications Market Report 2022 (ofcom.org.uk)).

(iii) **Not a digital multi-sided platform market with network effects.** Nor is this a digital market characterised by significant network effects[128] or ecosystem lock-in[129] in which some consumers on one side of the platform do not pay for a service.[130] In *Mobile Ecosystems*, the CMA found that 80% of users only use one smartphone, i.e. typically customers have either an Android or an IOS mobile device, and users rarely switch between iOS and Android devices.[131] Instead, Ofcom's own conclusions demonstrate that multi-homing/multi-cloud is prevalent amongst customers in the cloud infrastructure services market. Due to the very nature of IaaS/PaaS/ SaaS being deployed as a service to customers, if a customer wishes to switch cloud infrastructure service provider they do not need to buy new cloud infrastructure themselves. Switching costs between cloud providers depends on the nature of the solution and are largely in the customer's control.

Cloud has no direct network effects given that the size of each cloud providers' user base provides no benefit to other users. In fact, each incremental user may reduce the overall experience of other users given users acquire the same computing power, storage and network bandwidth. Any indirect network effects generated through the cloud marketplace are weak given the availability of multi-cloud solutions and pervasive multi-homing.

7.2.2 A sense check from the CMA's conclusions in music streaming

(140) To the extent that analogies with other markets are at all useful, far and away the best one that Microsoft can identify is the CMA's recent *Music Streaming* market study as a sense check, where the CMA chose not to make a MIR for reasons which have direct read-across to the cloud market.

(141) First, the CMA found that switching by existing users was limited as suppliers primarily focused on attracting new users (i.e. those who do not already use a paid-for streaming service) given the expanding nature of the market.[132] Ultimately, the CMA concluded that "*while there may be some barriers to switching, they are not currently a major problem to the functioning of the market given the focus of competition on new users*."[133] This proposition does have some read-across for the reasons sets out in Section 5 above.

(142) Secondly, one of the reasons reached for limited switching was the competitive pricing of the main music streaming service providers. The CMA concluded that this was not a concern as long as it meant that long-standing/less-active customers who had not switched did not end up paying significantly more than new customers.[134] This has read-across in light of Brattle's analysis of uniform existing vs. new customer pricing for Azure, above at Section 6.

(143) Thirdly, the CMA found in *Music Streaming* that customers may be reluctant to switch providers when they have curated their own playlists or their streaming service has developed a good

[128] The CMA found in the *Online platforms and Digital Advertising* market study that same-sided networks effects restricts consumer's ability to switch away from (for example) social media platforms. Online platforms and digital advertising-Market study final report-1 July 2020, para 3.205. Online platforms and digital advertising-Market study final report-1 July 2020, para 3.205.

[129] Mobile ecosystems-Market study final report, para 4.209; Mobile ecosystems-Market study final report, para 4.209.

[130] In *Online platforms and Digital Advertising,* the CMA also found that the fact that consumers do not pay a price for certain online platform services (e.g. for social media services) limits their incentives to switch providers for price reasons. Online platforms and digital advertising-Market study final report-1 July 2020, box 2.2. Online platforms and digital advertising-Market study final report-1 July 2020, box 2.2.

[131] Mobile ecosystems-Market study final report, pg. 28 and para 3.41. Mobile ecosystems-Market study final report, pg. 28 and para 3.41.

[132] Music and streaming-Final report, para 4.70.

[133] Music and streaming-Final report, para 4.76.

[134] Music and streaming-Final report, para 4.71.

understanding of the music they like, resulting in highly-valued personalised recommendations which a different streaming service may not be able to replicate easily.[135] This is somewhat analogous to the analytics of user data and the integrated and customised / tailored products offered by cloud infrastructure services.

(144) Fourth, the CMA found very good consumer outcomes in the market (a vast selection of music available for monthly prices that had failed to keep up with inflation in recent years). This again has strong similarities with the cloud market and it was one of the key reasons behind the CMA's decision not to launch an MIR.

7.2.3 The Ofcom survey does not support the proposition that supplier conduct is primarily responsible for any lock-in

(145) The degree of switching in this market is not caused by contractual restrictions as shown by the Ofcom Survey shown in Figure 9 below.

Figure 9: Findings from the Ofcom Survey

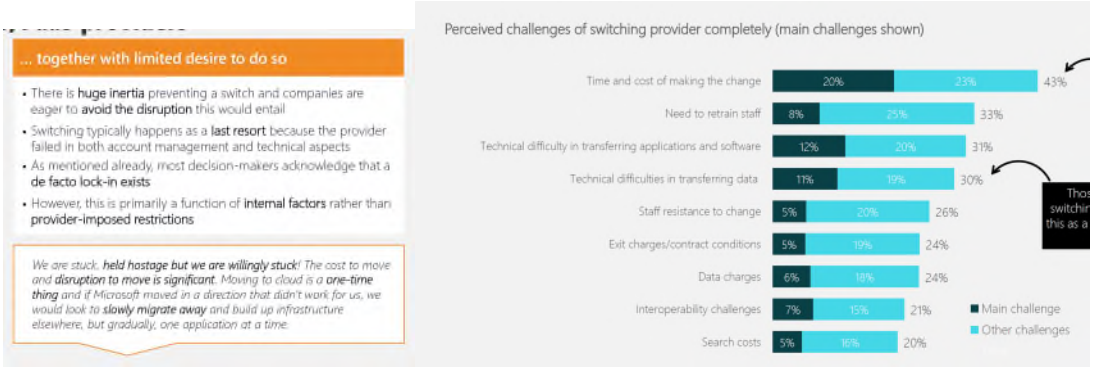

Source: Ofcom and Context Consulting, Cloud Services Market Research, Summary of Findings, March 2023, p. 110 and p. 121.

(146) The Ofcom Survey rather finds that 'low' switching is "*primarily a function of internal factors rather than provider-imposed restrictions*.[136] For firms to switch between providers they must achieve a benefit that compensates for the cost of switching providers as Ofcom also finds in its survey.[137] If there is no benefit to switching (because there is no exploitation of locked-in customers), firms will choose to remain with their current provider – without this being an indication of the reduced competitiveness. In the current, competitive market, there is no advantage to invest considerable time and money in a mostly identical product at a similar price. Indeed, looking at the right side of Figure 9 above, the top five challenges of switching listed by respondents to the Ofcom Survey are of technical nature or self-imposed. The left hand side of Figure 9 clearly indicates that respondents do not think this is a concern, as "*(…) we are willingly stuck (…) and if Microsoft move in a direction that didn't work for us, we would look to slowly*

[135] Music and streaming-Final report, para 4.72.

[136] Ofcom and Context Consulting, Cloud Services Market Research, Summary of Findings, March 2023, p. 110.

[137] The survey results indicate that improved service quality and lower price are the top two triggers that would prompt switching among respondents. Ofcom and Context Consulting, Cloud Services Market Research, Summary of Findings, March 2023, p. 1108.

migrate away (…)"[138] – again highlighting that competition is for the next workload, rather than to entice moves of existing workloads.[139]

(147) Any perceived low level of switching in the cloud services market for existing workloads would therefore not be unexpected and also not unique to this market, as the same phenomenon is also observed in other industries.[140] For example, if Microsoft were to try to increase the price of an existing service for a customer, this customer would have incentives not to migrate additional workloads to Microsoft but rather to a competitor or keep the services on premises. The potential punishment for "unfair" pricing behaviour keeps cloud providers in line, as it does in other industries. There are in fact many industries in which there are long-run supply contracts, but these input markets are still considered to be highly competitive.

(148) Considering the technical difficulties associated with switching cloud service providers for certain customers, the fact that there is observable fierce competition for new workloads (from existing clients and rivals' clients), as well as some level of switching of existing workloads, is therefore evidence of intense competition rather than the opposite. This is supported by the Ofcom Survey results, which indicate substantial instances of both actual and potential switching, as illustrated in Figure 10 below. Specifically, 18% of respondents have actively switched providers for all existing workloads, while 35% have switched to another provider for new workloads (which also makes investments necessary to facilitate switching of workloads between cloud service providers). Another 35% are considering switching provider and only 23% of respondents have never considered switching. For an industry alleged by Ofcom to possess high switching costs, the level of actual and potential switching is surprisingly high, and enough to maintain a high level of competition, which is driven by the behaviour of marginal customers who are willing to switch provider.

[138] Ofcom and Context Consulting, Cloud Services Market Research, Summary of Findings, March 2023, p. 110.

[139] Firms that have not considered switching are the ones that associate switching with higher burdens. Of firms that have not considered switching, 55% consider the time and cost of making the change to be a challenge in switching. This number significantly reduces to 34% when limiting the sample to those that have looked into switching in more detail. Similarly, technical difficulties in transferring data are over-exaggerated by customers that have not considered switching. Customers that did, consider this to be far less of a challenge (39% as opposed to 22%).

[140] For example, in the automotive sector, original equipment manufacturers (OEMs) compete for the opportunity to deliver a certain volume of parts to car manufacturers, and once chosen, switching to a different OEM involves substantial costs on both sides. Nonetheless, competition among OEMs for the initial contract is fierce. But prices also do not increase when the parties are locked into the contract due to switching costs. Being in a long-term relationship means that there is a general rule of fair sharing of benefits obtained, for example, from cost reductions through the price set. In a dispute, there is typically a resourcing of the product to establish a (fair) market price. This tender is often won by the incumbent. Such long-term relationships are generally referred to as relational or implicit contracts, in which a fairness rule in sharing rents is enforced by the threat of contract breakdown. This works because the value of future business opportunities far outweighs the current contracts by several orders of magnitude, creating an incentive to build a long-term relationship between the parties.

Figure 10 : Switching behaviour of UK cloud service customers

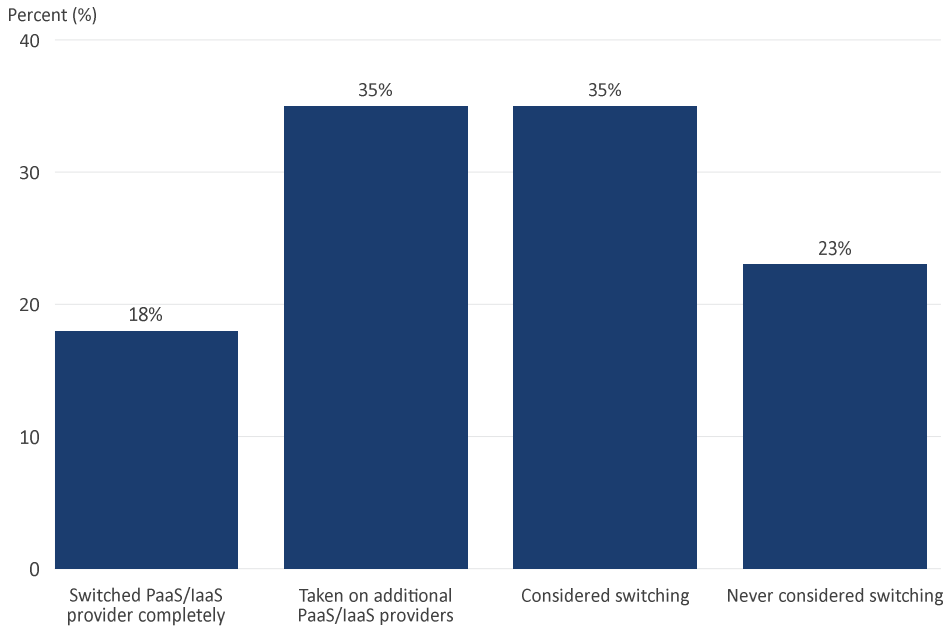

Source: Ofcom and Context Consulting, Cloud Services Market Research, Summary of Findings, March 2023, p.105. Note: Based on the question: Q47. Have you ever done the following? Base: All cloud users (n=889).

7.3 The "high profitability" narrative

(149) In the Interim Report, Ofcom reports "*evidence of high profitability among AWS and Microsoft*" and considers this as an indication that competition may not be working as well as it should in the cloud infrastructure market.[141] Ofcom acknowledges that *"(a) finding that ROCE is higher than the WACC in any particular period is not necessarily indicative of a competition problem*".[142] Besides this limitation acknowledged by Ofcom, Microsoft considers that there are fundamental issues with Ofcom's approach to assessing profitability in cloud services that invalidate Ofcom's subsequent conclusions regarding level of competition.

(150) In order to provide meaningful information given the scope and consequence of Ofcom's analysis, the metrics considered by Ofcom (ROCE ("**return on capital employed**") and WACC ("**weighted average cost of capital**")) should accurately be calculated taking into account the high levels of investment and long return period that are characteristic of cloud computing investments. Important considerations, such as the value of intangibles relevant to Microsoft's cloud business and asset premia due to capacity constraints, are aspects that may be complex to quantify accurately but they have a considerable impact on the accuracy of the resulting estimates.

(151) Importantly, as set out above, Microsoft and other providers have been investing heavily to create and expand the cloud services market. Due to the nature of cloud services, observing Earnings Before Interest and Taxes (**EBIT**) and ROCE over a one- or five-year period will not capture the dynamic nature of the investment's profitability (e.g. because the cloud business requires a very large upfront investment on which a return is earned over a long period of time), and it will ignore early periods where the business was significantly less profitable. Ofcom's

[141] Interim Report, para 6.13(d).

[142] Interim Report Annex, para A6.43.

comparison of ROCE to WACC uses a static value for WACC and only a five-year period of ROCE metrics. Both measures are dynamic which limits the value of point in time comparisons, particularly given that market conditions and business conditions can change considerably after the decision to invest (as is the case in the cloud market). This is also why Ofcom's comparison of the profitability of AWS to more recent entrants such as Google provides little insight from an economic perspective.

(152) Even more fundamentally, the guiding thesis of Ofcom's profitability analysis is incorrect. The Interim Report compares estimates of realised ROCE to estimates of WACC, for AWS, Microsoft Cloud and Microsoft, respectively, with the expectation that ROCE will converge to WACC over the time period considered. Although companies may consider WACC as a general comparative benchmark for ROCE or other measures of profitability, the expectation imposed by Ofcom that ROCE should *converge to equal* WACC is misguided. Expecting ROCE to equal WACC would constrain the returns on the company's use of capital to equal the costs of procuring such capital; in fact, returns can and should consistently exceed the cost of capital owing to pro-competitive reasons such as successful innovation, successful brand management, and operating efficiency. Publicly traded companies in competitive industries frequently have ROCE greater than WACC over sustained periods. Considering the calculation of each metric further illuminates the issues with Ofcom's thesis: WACC has been developed as a discount rate for after-tax cash flows in a valuation model, while ROCE employs EBIT as a measure of profits, which does not include certain cash costs and is a pre-tax quantity.

(153) Ofcom's comparison of realised ROCE to estimates of WACC is thus comparing apples to oranges and is not informative as to competitiveness in the cloud services market.

Part C: Specific concerns do not justify an MIR at this time

8 Committed spend discounts

8.1 Discounts are the essence of how Azure competes on price for larger users

(154) The Interim Report acknowledges that "*(d)iscounts are generally a positive feature of markets, leading to lower prices and indicating competition between providers on price*".[143] Nonetheless, the Interim Report is concerned that committed spend discounts offered by cloud providers risk distorting competition.

(155) Like in many other markets, committed spend discounts have been introduced to provide benefits to customers who agree to a minimum spend. Cloud providers need to make significant investments in capacity which comes with risk. Volume discounts are common in largely fixed-cost industries such as cloud infrastructure, where the marginal cost of supplying an additional unit is relatively low compared to the initial fixed costs.

(156) Under committed spend contracts, customers determine their consumption plan based on specific projected workloads. While this is similar to planning on premises infrastructure investments, customers retain greater flexibility to change their cloud strategy. Such contracts help providers to both plan capacity and de-risk the significant capex investments required to provide cloud infrastructure, facilitating investment and allowing for more efficient capacity utilisation. In a fiercely competitive marketplace, these efficiencies are then passed on to customers as discounts (which, for Azure in the UK, includes customers such as [✂]).

(157) Committed spend discounts therefore allow for better risk allocation between the parties. Rather than distorting competition, they are indicative of intense price competition in the sector. Any interference with firms' pricing could reduce competition and increase prices for UK enterprises.

8.2 MACCs: an overview of key terms

(158) Cloud infrastructure services are generally charged on a standardised pay-as-you-go basis. Microsoft Azure Consumption Commitment contracts ("**MACC**") offer customers individually negotiated discounts in return for a minimum committed spend. MACCs are essentially volume discounts, since discount rates vary according to the committed spend levels.

(159) The typical features of MACCs[144] are summarised below:

(i) MACCs are offered to customers who are willing to commit to a total spend of [✂][145] and the contracts are typically [✂].

(ii) Microsoft works with the customers to ensure it has a projected consumption plan for Azure that underpins the committed spend level. Microsoft has no interest in agreements that are not realistic and works hard to make sure that the customer's commitment is realistic. Typically, the committed spend is set at [✂] of the projected consumption plan. Projected consumption will be based on expected needs for specific

[143] Interim Report, para 5.125.

[144] Additionally, as part of a MACC contract, Azure offers four types of "concessions" or benefits to customers. These are: (i) **End Customer Investment Funds (ECIF)** – funding set aside to pay for services that support the end customer with demonstrations, in deploying their workload on Azure and usage of Azure solutions. If a customer significantly falls behind its committed spend, Azure will use ECIF funds to further support the customer's workload migration to Azure. (ii) **Azure Credit Offer (ACO)** – Azure provides free credits to mitigate the costs of migrating from a competitor or from on premises infrastructure. Often, a customer has to dual-run a workload during migration. ACOs help mitigate such duplicate costs. (iii) **Consumption-based discounts** - discounts against PAYG list prices for the duration of the contract. They are also applicable to spending above the committed level. (iv) Other programmatic discounts such as on Savings Plans, Reserved Instances.

[145] [✂].

workloads identified as part of the planning, and where relevant, historical consumption – i.e. to ensure acceptance by large and sophisticated customers, MACCs are not set as "stretch targets", but rather specifically designed so that customers will easily meet them. This allows the customer to retain the choice of cloud providers for incremental workloads.

(iii) MACC customers have price protection and pay the lower of the website price at the time of contract and the current website price (net of contract discounts).

(iv) Customers retain complete flexibility on *how* to draw down the committed spend. MACC commitments are typically not specific to a particular service and can be used on any Azure services, from any data centre, in any geography, at any time during the contract.

(v) Committed spend can also be drawn down on third-party ISV purchases dollar-for-dollar on the Azure marketplace[146]. As a result, ISVs whose products are available on Azure Marketplace and eligible for MACC drawdown do not compete with committed spend on Azure to win customers. They can access the business of customers who have already agreed to committed spends.[147]

(vi) Customers have flexibility on *when* to draw down their committed spend over the full term of the contract with no specific monthly or annual commitments. Customers are billed monthly based on their usage and the agreed discounted prices i.e., they do not need to have met their commitment to start receiving the discounts.

(vii) If a customer does not meet an Azure consumption commitment during the negotiated period, Microsoft [✂]. Data from recent UK MACC contracts that concluded in 2022 show that as many as [✂] of customers meet at least [✂] of their committed spending during the term of their contracts, with some [✂].[148]

8.3 The rationale for committed spend discounts

8.3.1 Demand side: competitive pricing, predictability on spend and flexibility in usage

(160) At a high level, Microsoft's strategy is to ensure its prices are competitive [✂]. We understand committed spend contracts are now common across cloud providers and also offered by GCP, Oracle and IBM.[149]

(161) Ofcom identifies customer challenges in predicting cloud spend and lack of price transparency.[150] A committed spend contract helps customers address this challenge. While they are invoiced every month based on usage, at an aggregate level, they have visibility of their spend with a cloud provider over the term of the contract. Azure offers functionality on the Azure portal that allows customers to track contract start and end dates, invoiced spend, remaining commitment, and eligible spend.[151]

---

[146] Eligible third-party services draw on Azure services; Microsoft receives a commission of 3% on sales through the marketplace.

[147] A report by Tackle.io (cited by Ofcom[147]) on the state of cloud marketplaces found that 40% of surveyed buyers had an enterprise agreement with marketplace spend included, and that 48% of buyers cited drawing down committed spend as the most important cloud marketplace buyer benefit. Moreover, unlocking co-sell opportunities with cloud providers is the number one reason surveyed sellers choose to list their products on marketplaces (cited by 74%). See Tackle.io. Learnings, predictions, and insights into the state of Cloud Marketplaces in 2022. https://tackle.io/resources/reports/state-of-cloud-marketplaces/, accessed 2nd May 2023. Cited by Ofcom, Interim Report, para 5.150.

[148] See Microsoft's response to Ofcom's RFI of January 2023, Q11.

[149] See Committed use discounts | Documentation | Google Cloud; Account types | IBM Cloud Docs; Cloud Cost Savings – Savings Plans – Amazon Web Services; Oracle PaaS and IaaS Universal Credits Service Descriptions.

[150] Interim Report, paras 5.168 and 5.177.

[151] Track your Microsoft Azure Consumption Commitment (MACC) - Microsoft Cost Management | Microsoft Learn.

(162) As set out above, customers enjoy significant flexibility in how they reach their committed spend (including through third-party ISV purchases on the Azure marketplace). Microsoft's MACCs are not designed or imposed to capture the entirety of a customer's demand across cloud (particularly given many companies' preference for multi-cloud). Evidence shows that customers multi-cloud even when they have a MACC with Microsoft.[152] [153]

(163) Unlike textbook loyalty-inducing discount schemes prohibited under competition law, spend thresholds are generally planned by customers based on specific workloads in Azure only.

8.3.2 Supply side: capacity planning and investment forecasts

(164) Fundamentally, spend commitments allow Microsoft to better forecast, plan and manage data centre and server capacity and capex.[154] On average, depending on the geography, it takes Microsoft [✂] to build a data centre and about [✂] to set up new servers. Increasing the predictability of near- and medium-term future demand allows Microsoft to undertake the construction and operation of larger and more efficient data centres than would be the case without such detailed insight into customer demand.

(165) Further, Microsoft (and other cloud providers) need to make a series of customer-specific investments and are therefore assuming risk that they may not realise the returns on their investments (a classic "hold-up" problem[155]). Committed spend contracts like MACC help overcome this "problem" (and customers receive the benefit of their commitment in the form of discounts and concessions).

(166) Without the signal from customer contracts, albeit noisy and imperfect, the demand planning horizons would be longer as shown in Figure 11 below. In a highly innovative sector like IaaS where servers quickly become obsolete, longer demand planning horizons would lead to inefficient capacity. Moreover, without the ability to forecast with confidence, there may be massive over- or under- utilisation. As a result of competition in the cloud infrastructure market, these benefits are passed on to customers through contractual discounts.

Figure 11: Illustrative impact of committed spend on demand planning

[✂]

8.3.3 Committed spend discounts are therefore mutually beneficial

(167) Committed spend discounts are mutually beneficial: they allow customers to access lower-cost services and better predict cloud costs and usage; at the same time, committed spend discounts assist cloud providers with long-term capacity planning, which is of vital importance given the need for high-fixed cost infrastructure investments that are location-critical.

8.4 Customers do not view discounts as an impediment to switching or multi-sourcing

[152] [✂]

[153] [✂]

[154] See Microsoft's response to Ofcom's RFI of January 2023, Q11.

[155] Hold-up arises when part of the return on an agent's relationship-specific investments is ex-post appropriable by his trading partner. Investments are often geared towards a particular trading relationship, in which case the returns on them within the relationship exceed those outside it. Once such an investment is sunk, the investor has to share the gross returns with her trading partner. See, for example, Patrick W. Schmitz, "Incomplete contracts, the hold-up problem, and asymmetric information", Economics Letters, Volume 99, Issue 1, 2008.

(168) Feedback from customers shows that they don't consider committed spend discounts to be a significant barrier to switching or to the adoption of multi-cloud.

(i) Committed spend discounts do not appear to be a significant barrier to switching. Ofcom's survey found that 18% of respondents switched IaaS/PaaS providers.[156] When asked about challenges to switching, contract conditions were only mentioned by 24% of respondents, and within this, only 5% of respondents identified them as a *main* challenge to switching. It is not listed in the top five reasons for any breakdown of the respondents, including those who have switched and those who have not. This is also echoed in Ofcom's qualitative study, in which participants did not mention the loss of discounts (i.e. price increase) or contractual conditions as being a relevant challenge.

(ii) Committed spend discounts are also not a significant barrier to multi-cloud adoption. Ofcom's survey found that 52% of respondents use more than one IaaS/PaaS provider.[157] Contractual issues were only mentioned by 27% of respondents as a challenge to multi-cloud, which ranked sixth among the challenges selected.[158] This suggests that even if contractual issues are perceived as a barrier, other challenges become more important in experience (in this case, technological difficulties, interoperability, and accountability). Ofcom's qualitative study finds similar results, with commercial practices being the least important reason highlighted by interviewees.[159]

(169) Additionally, industry reports such as Oracle's 2023 survey report[160] and Foundry's 2022 EMEA survey[161] find that contractual issues (and a loss of discounts / resulting price increases) were not a significant obstacle to multi-cloud.

(170) These findings are consistent with intuition: given that committed spend levels are set based on projected consumption needs for one or more specific workloads, they present no barrier whatsoever to using a different could provider for other workloads.

8.5 Discounts are not "so large" as to somehow be adverse to customers' interests

(171) Ofcom acknowledges that it has little evidence on how committed discounts are structured and negotiated in practice.[162] Azure's customers are sophisticated buyers with significant buyer power. Large customers have dedicated IT procurement teams that commonly negotiate terms, including discounts, committed spend levels, contract length and the technical support offered under MACCs. Customers also often tend to play cloud providers against each other and demand parity with competing offers.

(172) As previously submitted,[163] Microsoft [✂]. Most customers receive discounts of between [✂] compared to PAYG pricing, with the average committed spend discount for UK customers being [✂] in 2022. The following chart shows the proportion of UK customers against the committed spend discounts in the ranges provided.

[156] Interim Report, para 5.41.

[157] Interim Report, para 3.26.

[158] Interim Report Annex, figure A7.4.

[159] Context Consulting, March 2023. Cloud services market research: summary of findings. Slide 80.

[160] S&P Global Market Intelligence, February 2023, Multicloud in the Mainstream: Making IT work 'as advertised.' Commissioned by Oracle.

[161] Foundry, Cloud Computing Study. Executive summary outlining the 2022 EMEA research findings.

[162] Interim Report, para 5.148.

[163] See Microsoft's response to Ofcom's RFI of January 2023, Q9.

Table 1 - Discount range of UK MACC contacts

[✂]

8.6 Conclusion

(173) MACCs and other volume discounts offered by cloud providers create efficiencies that generate directly lower prices for public sector and enterprise customers (such as [✂], amongst many others) and allow for more certain forecasting of demand and lower prices for suppliers.

(174) The Interim Report tentatively suggests a blunt tool of prohibiting or restricting discounts but rightly recognises the significant downsides with this course of action. Microsoft agrees with these concerns. Any intervention would undermine Azure's and other cloud providers' ability to forecast their demand to make expensive investments and could lead to higher prices for UK cloud customers.

(175) Consequently, Microsoft respectfully submits that this concern simply cannot serve as a reasonable basis for an MIR.

9 Supplier-imposed restrictions to interoperability and portability

9.1 Introduction

(176) The Interim Report finds material barriers to switching cloud providers or developing multi-cloud architectures due to a lack of interoperability and portability between cloud services, and concludes that these factors are due to: (1) "inherent technical differences" between cloud platforms, which increases the time and cost of switching cloud providers; and (2) the "practices of hyperscalers" including developing first-party services that are less compatible with open standards and not openly sharing information such as APIs and protocols.

(177) Microsoft's view, in brief, is that the Interim Report:

(i) is correct as to factor (1): there are inherent technical features that create some "lock in" (i.e. switching costs);

(ii) is incorrect, at least with respect to Microsoft, on factor (2): Azure's efforts on interoperability mitigate switching costs – and in any event do not aggravate them; and

(iii) fails to reflect that any costs are not holding customers back from switching between cloud providers or the deploying multi-cloud solutions.

(178) Differentiation between providers is the result of an active competitive process, including through services based on open source software and technologies such as Open Stack, Kubernetes and MongoDB. This is part of the lifeblood of innovation in cloud services that has delivered so many benefits to customers; to seek to limit such development whether through mandated standardisation, required equivalency or restrictions on open source-based services would inevitably limit innovation and differentiation. Sophisticated cloud customers understand this dynamic and make strategic choices about how to structure their cloud environments. Given the dynamic nature of competition, any intervention in this area is far more likely to chill investment, to stifle innovation and to commoditise customer choice, rather than to make an innovation-heavy market "work better." Microsoft's introductory comments in response to the Interim Report are as follows:

(179) First, while the Interim Report is entirely correct to note that IT environments have inherent technical differences that generate "lock in" by increasing switching costs, it lacks any comparison or counterfactual for "high" costs of switching in cloud IT relative to the general level

of IT switching costs in other IT contexts. The purchase of any product or service includes some amount of "lock in" to that purchase. It is difficult to make a meaningful assessment of supposedly "high" switching cost without asking the question: high relative to what? It would be instructive to compare (and ask stakeholders to confirm) that, relatively speaking, it is generally easier for customers to move between clouds than it is to move from on premises IT to cloud in the first place, or how the on premises IT industry worked before the cloud. All the evidence of which Microsoft is aware points to the fact that, all else being equal, cloud services have made switching less difficult than in other IT contexts.

(180) Second, while cloud is a relative mitigant for inherent IT lock-in, the degree of lock-in is fundamentally a choice made by the **customer** (most of whom are sophisticated in IT procurement), who weighs between:

(i) optimising a cloud environment for interoperability and portability to mitigate lock-in: **costs are upfront**, and the trade-off is generally lower value extraction from their "primary" or "initial" cloud environment they choose first; and

(ii) optimising for cloud environment customisation and integration: **costs are deferred** and the trade-off is **greater value extraction** from their primary cloud environment but then **greater switching costs** (relative to optimising for interoperability/portability) when moving to a new or secondary cloud environment.

(181) Third, the Interim Report, and the third-party commentary on which it based, are wrong to attribute (at least as subtext) "the practices" of Azure an anti-competitive motive or effect with respect to interoperability and portability, such as (deliberately) developing services that are "less compatible" or choosing not to "openly shar(e)… APIs and protocols" in order to enhance then exploit customer lock-in. Refuting this allegation is not premised on altruism or "doing the right thing". Azure's commercial incentive, which has underpinned its "challenger to AWS" growth strategy, is to provide best-in-class portability and interoperability and to support a wide range of proprietary and open source technologies to facilitate migration to, and use of, Azure. Any changes to make services less interoperable or undercut open source technologies would act equally as a barrier to moving to Azure or accessing its services in the first place: it is simply not possible to design an "asymmetrically portable" system to be both very seamless to switch into and very hard to switch out of. Fundamentally, Azure generates revenue every time a user calls and uses an Azure-hosted service, including from another cloud architecture; to deliberately degrade interoperability or portability would cut off these revenues and not be in Azure's commercial interests.

(182) Fourth, the Interim Report is confused about the distinct concepts of interoperability and portability which may partially explain its stance on interoperability as a barrier to multi-cloud. It makes numerous vague generalisations across hundreds or more IaaS/PaaS/SaaS services, and across cloud competitors. Consequently, the Interim Report misunderstands Azure users' ability to switch to other cloud providers or deploy multi-cloud solutions.

(i) Interoperability is the degree to which diverse systems can work together successfully. Interoperability in the context of the cloud relates to the APIs available to call a specific service or functionality. APIs for Azure Services are fully documented and available online.[164] Alternatively, developers can use the Azure Services SDK to programmatically access Azure services from application code. The SDK supports a wide variety of programming languages and is freely available under an open-source license on

[164] Azure REST API reference documentation | Microsoft Learn.

GitHub.[165] However, interoperability is a question of degree, not a binary state. For example, to the extent that APIs differ between two (or more) cloud providers, this may involve more or less work for the developer to develop the code to call those APIs. If a developer currently relies on Google Clouds Vertex AI service, switching to call Azure's Cognitive Service APIs would require changes to the application to match the functionality provided by Azure and to call its specific APIs.

(ii) Interoperability must be distinguished from portability (i.e. the ability to move an application or data from one cloud to a different environment). For example, instead of calling Azure Cognitive Services from GCP, a customer may want to move the underlying data from GCP cloud storage to Azure. This requires the ability to access the data through APIs and to understand its structure (if any) to move it into a different data storage solution in Azure. Unstructured data can be moved easily and structured data is also often highly portable between cloud providers, particularly where the customer uses a database that is widely available across different providers, such as MongoDB. Similarly, moving an application, depending on the underlying technologies on which it is built, could involve simply pushing a button or it could involve significant refactoring of the application to call into the services of the new cloud. This is similar across all types of technologies. For example, mobile developers will build an application for iOS and then engage in some modifications of that application to make it run on Android.

(183) The Interim Report refers to application and data portability, in each case referring to something being portable only where it can be moved with no refactoring (i.e. where it is automatically or instantly compatible without more). Ofcom's taxonomy therefore implies that workloads are not portable if they require refactoring to run effectively on another cloud architecture, but this is not accurate. For a "one off" move (i.e. switch) applications or data can be refactored – the complexity of this will depend on the application or data in question and will self-evidently be more practical or some workloads than others, the fact an application requires some refactoring to run effectively on another cloud architecture does not imply an absence of portability. This confusion partially explains the contradiction between the Interim Report's view that multi-cloud is rare and industry studies that find it to be the norm.

(184) To map these concepts to the three kinds of multi-cloud referred to in the Interim Report: (i) only interoperability is relevant to considering integrated multi-cloud; while portability is most relevant to (ii) switching between cloud providers but is also critical to (iii) cloud duplication (as it requires an application and data to be compatible with both the original and duplicate cloud environments). Given Ofcom elides "portability" with "portability without refactoring", this renders "non-portable" many workloads which in fact, can be ported, with some switching costs that customers can weigh in their cost/benefit assessment. "Holding" providers to a zero-refactoring portability standard would effectively freeze innovation by forcing homogenisation of cloud providers' offerings (the cloud equivalent of obliging a harmonisation of iOS and Android to avoid any developer modification effort to make an app written for Android work on iOS).

Case study 1 – Azure's innovative proprietary storage solution

The evolution of Azure Storage is the direct result of innovation competition: Microsoft developed this system as an alternative to S3 with superior functionality. For example, Azure Storage allows overwriting of just one part of a stored file, rather than the entire file. To give a practical example, if a two-hour movie was stored on S3 and the user wanted to replace only five minutes, S3 would require the entire two hours to be overwritten, whereas Azure

[165] Connect your app to Azure Services | Microsoft Learn; GitHub – Azure sdk.

Case study 1 – Azure's innovative proprietary storage solution

Storage allows only the relevant five minutes to be overwritten, obviously reducing the resources that need to be consumed to perform the task. Microsoft could not offer this functionality on the S3 standard: Azure Storage differentiated, by design, to make it a more compelling competitive proposition.

Given that Microsoft and AWS have built their storage products with fundamentally different design architectures, a layer of abstraction may be required to use either in isolation (where data will be sent from a cloud storage service to another cloud service) or in tandem (where data will be drawn by a cloud service from both cloud storage services simultaneously). However, the same fundamental data portability issues apply to all cloud providers, regardless of whether they are using the same base format. While there may be some surface-level portability between (for example) GCP and AWS, full functionality offered by each is typically only compatible with native storage and computing services. Google would be required to do additional work to call Amazon storage from BigQuery (and this would need to be done for each cloud provider nominally supporting S3 API architecture). To mitigate these issues, most cloud providers have tools to enable data to be moved from one data source into their cloud.

(185) The first part of this section seeks to unpack factors that drive friction in implementing multi-cloud solutions and the customer view, which will always exist in any technology market, that "things don't work as well" as they potentially could between clouds. In the second part of this section, we give detail on the efforts Azure has made to be best-in-class in terms of interoperability and portability (which as above is the product of a deliberate strategic decision to do so as the historical "challenger") that must make it as easy as possible for AWS customers to utilise Azure and addresses the examples relating to Azure highlighted in the Interim Report. We then explain that Ofcom's own survey, as well as Azure's interactions with customers, suggest that the barriers to multi-cloud are not as significant as presented by Ofcom given the high level of multi-cloud adoption. Finally, we conclude by assessing the proposed interventions highlighted by Ofcom and explain the challenges they would entail and the potential unintended consequences that could result.

9.2 Factors driving multi-cloud friction

(186) As noted above, all users of cloud systems are "locked-in" at one level of the IT stack, but this is the result of deliberate choices made by highly sophisticated customers about how to best optimise their architecture. While it is unsurprising that, if asked, customers would express a preference to have everything work together seamlessly, but the fiction in having different systems work together is not unique to cloud markets, and is inevitable across IT systems more broadly. The existence of such frictions however should not be equated with a lack of customer choice and is driven by the following mix of factors.

9.2.1 Latency in data transfer constrains performance

(187) The overall performance of workloads running across clouds will depend on many factors, including the latency achieved as data is transferred between systems and locations. Self-evidently, latency is more important to some use cases than others – to take the example of using cognitive services to caption television, a delay of one to two seconds would be unimportant for a streaming service like Netflix (which generates the captions once, in advance of them being "launched" to users) than for live TV.

(188) For use cases that require fast response times, customers may prefer to house their data, computing power and networking in a single place to minimise latency. A reference to "place"

in this context is not just to a specific cloud provider, but to the physical place where the data centre is located. Customers can limit the impact of latency on their cloud to a certain extent by using private networks or concentrating their data and services in a single region. However, it is impossible to eliminate latency when deploying workloads over multiple clouds and/or locations as this is fundamentally an issue of physics – it is a "speed of light" issue.

(189) To give a practical example, if a customer stores their enterprise data in a data centre located in Milton Keynes, accessing that data from India will be slower than accessing the data from London – i.e. there will be higher latency, regardless of whether the application seeking to access the data in India is hosted on the same, or a different, cloud. There are technical solutions that exist to try to overcome this. For example, Azure provides ExpressRoute, a private connection between Azure datacentres and a customer's network or another cloud provider's equivalent private connection. This can increase speed and reliability, and reduce costs.[166] Further, Azure interconnects with the Cloudflare service (competitors have similar offerings) that replicates the relevant content all over the world and customers will purchase this service where they have global demands and latency is an important consideration.

(190) These technical solutions can only mitigate latency to a certain degree (rather than eradicate it). Some customers may for good reason decide that it is not worth the cost to duplicate the entire storage location of their data in two different clouds where the data is stored in one cloud but accessed from another, simply to minimise latency, which cannot be lowered below what physics would allow. This will drive single-homing of data storage for some customers and particular applications.

9.2.2 The complexity of the cloud environment

(191) Multi-cloud architectures tend to be inherently more complex. Cloud solutions with a larger number of interconnected applications, that have significant data flowing between service providers, are likely to be less efficient than a cloud environment relying on a single cloud with native storage and processing solutions. Data-centric workloads are likely to suffer worse latency and efficiency if they are implemented with interconnections with other cloud environments, as network considerations are crucial to overall performance (as discussed above). Further, introducing multi-cloud functionality may create single points of failure that can lower system resiliency. It may also create complexities in managing procurement and costs, compliance with regulatory requirements, security, and more.

9.2.3 The quality of coding of each of the applications and platforms, and the lag between cloud provider updates and customer deployment

(192) The quality of coding of cloud applications (for ISVs) and cloud environments (for customers) varies depending on the level of expertise and resources available to an organisation. Microsoft, and other cloud providers, are continuously updating their services with the latest security enhancements and innovations. An ISV or third-party cloud provider may not be able to fully test and update their applications until some time after release. As a result, there may be a reduction in interoperability between applications updated at different times, and leading to lower performance or product quality (i.e. certain features do not work correctly).

9.2.4 Differentiation and IaaS/PaaS/SaaS

(193) As the Interim Report acknowledges, the ease with which customers deploy multi-cloud solutions differs between the various levels of cloud infrastructure. Customer solutions that rely only on base level IaaS tends to be more portable and compatible than solutions that rely on

[166] Connectivity to other cloud providers - Cloud Adoption Framework | Microsoft Learn.

more differentiated IaaS or PaaS services, which might require further effort to ensure the application can run on a different cloud (i.e. is compatible). To run applications or workloads across clouds may require more technical effort, because solutions in one cloud may not be optimised for the architecture of another cloud. Most cloud services, however, can be reconfigured in a way that is fully functional from an end-user point of view.

9.2.5 Sophisticated customers make deliberate trade-offs

(194) The consequence of the above factors, as recognised by Ofcom, is that services from different cloud providers are not as "tightly integrated" as two comparable cloud services from a single provider, which can create an additional technical cost for customers.[167] There are inherent trade-offs in building for interoperability, portability, or cross-cloud compatibility because some more advanced features and more complex engineering can simply be more difficult to move/replicate in another cloud. This has implications for the ease with which it is workloads can be moved between clouds (i.e. portability), how well applications can interact with one another across different cloud providers (i.e. interoperability), and how seamlessly a single application can run in multiple different clouds (i.e. application portability).

(195) Given these inherent limitations, customers face a choice to either design their cloud architecture for maximum interoperability or portability or to optimise for value extraction from a single cloud environment. These decisions are taken on a workload-by-workload basis meaning that, critically, the decision to have one set of workloads on a particular cloud does not preclude later having another set of workloads on another cloud. Not all a customer's IT systems will need to be able to "talk" to one another – e.g. payroll and employee management systems do not require extensive interoperability with modern workplace tools and therefore even a non-portable, non-interoperable architecture would allow a single customer to multi-cloud (i.e. what Ofcom describes as "siloed multi-cloud"). A practical example of how Vodafone does this is set out below as a case study.

Case study 2 – Vodafone Multi-cloud

Vodafone has adopted a multi-cloud strategy, sourcing IaaS and PaaS services from a range of cloud providers and ISVs. Vodafone's approach allows it to pick and choose the "best of breed" services from each cloud provider to meet its needs. Vodafone were targeting putting 65% of their applications in the cloud by 2022.

Vodafone's cloud services include: scaling its infrastructure and running applications in AWS; managing its front-end systems in Azure (which it procures on a project-to-project basis); storing and analysing its data with GCP (with plans to move its big data and business intelligence workloads to a new platform on GCP); hosting business and operations support systems on premises with Oracle; and hosting a variety of applications, including internal billing and support systems, on premises with VMware.

Vodafone has historically managed its own data centres but is now investing in the cloud for its scalability and synergy with 5G technology. Vodafone's new multi-cloud strategy exemplifies the approach in the wider telecommunications sector. Even while expanding into the cloud, telecoms operators will need to maintain multiple clouds and on premises infrastructure in order to comply with resilience and national security regulatory duties under the Telecommunications (Security) Act 2021. The telecoms sector is therefore highly competitive for cloud providers and no telecoms operator is 'all in' on a single cloud provider.

[167] Interim Report, para 5.24.

(196) As a result of the above, customers make a range of choices which affect the nature and scale of the "lock-in" to a particular cloud architecture.

9.2.6 Customers can choose to optimise a cloud environment for interoperability and portability

(197) As acknowledged by in the Interim Report, it is largely within customers' power to develop cloud agnostic environments that allow communication between cloud stacks and the movement of applications and data between cloud environments.[168]

(i) Interoperability

(198) Cloud providers, including Microsoft, facilitate interoperability and multi-cloud solutions by making their APIs open and providing information for customers to allow them to connect their applications hosted on different clouds. For example, Microsoft publishes the vast majority of its APIs and technical documentation on its website and GitHub to allow third-party services to communicate with its cloud services.[169] To utilise these services, customers merely need to 'call' the API by altering the code in their cloud solution to access the functionality exposed end-points. This allows the customer to send data from the cloud service hosting their primary workload to Microsoft for processing, without needing to acquire multiple IaaS or PaaS products. In providing access to its proprietary first-party services, Microsoft does not prioritise traffic from its own IaaS or PaaS users. [170]

(ii) Portability

(199) As the cloud market matures, ongoing innovation is improving customers' ability to switch between providers in a multi-cloud environment. ISVs are designing their services with portability as a first principle, and cross-platform offerings are now widely available to customers. This reflects the preference of ISVs to develop software that has out-of-the-box cross-cloud compatibility to maximise the potential number of customers and to align with market expectations relating to multi-cloud strategies.

(200) There is a growing ecosystem of ISV applications such as Snowflake and Databricks (both of which are supported by Microsoft) and multiple Kubernetes offerings (Microsoft offers two: Azure Kubernetes Service and Azure Red Hat Open Shift) that are built on open-source projects.[171] These applications offer a layer of "abstraction" from the underlying infrastructure, allowing them to be run on each cloud providers IaaS offerings (i.e. they are compatible with different providers). Consequently, customers can more easily move their data and applications between cloud providers if necessary, or interact with their data and applications across multiple clouds using a unified interface. This benefits ISVs by allowing them to retain more customers, wherever they move across the cloud, and places competitive pressure on cloud providers to develop their own first-party applications that are multi-cloud enabled.

(201) For example, Kubernetes allows the development of containerised applications that are decoupled from the underlying infrastructure. By deploying application components in containers, the application components interact with Kubernetes, which then communicates with the chosen cloud environment. This architecture is cloud-agnostic and can be moved between different cloud providers with relative ease. Kubernetes is becoming increasingly

[168] Interim Report, para 5.47.

[169] Azure documentation | Microsoft Learn; Connecting Azure to public clouds | Microsoft Learn.

[170] See Annex 4 for further details on the way in which Microsoft facilitates the interoperability of its cloud products / services.

[171] Deloitte Insights Open for business (deloitte.com).

popular and, going forward, we expect more developers will deliver their cloud applications in a containerised format to maximise portability.

(202) By way of further example, Snowflake is a rapidly growing service that provides customers with a cloud vendor agnostic data platform. It is available on all major clouds, including AWS (from 2014), Azure (from 2018), and GCP (from 2019). Snowflake is a fully portable solution, in the sense that customers can deploy the application and load data from AWS, Azure, or GCP and switch between providers while retaining full functionality.[172] In addition, Snowflake is cross-cloud compatible and can run a single workload over multiple cloud infrastructures, allowing users to move data from their applications across cloud providers and across regions.[173]

(203) From an end-user standpoint, adopting a portable or multi-cloud environment can often be done without degrading the overall functionality of cloud services. However, the process of setting up an architecture to be portable is inherently more complex than developing for a single environment. Customers will make the choice to do this only where that offers a compelling value proposition for their needs. We give a case study below of [✂] investment in a fully portable cloud architecture.

Case study 3 – [✂] – Kubernetes containerised solution

[✂] developed a fully containerised cloud environment to enable it to switch to Azure from AWS. On Azure, [✂] was able to realise lower operating costs and greater developer efficiency through standardised [✂] Dev/Sec/Ops capability and IP.

[✂] multi-cloud strategy focused on ensuring portability as a priority to enable it to 'lift and shift' its architecture at any time (at a higher price). [✂] emphasis on portability (while at also extends to its contracting arrangements, [✂]. This approach maximises [✂] ability to port its applications to another cloud. It reflects the fact, however, that the customer is ultimately in control of balancing portability versus taking advantage of the full power of any particular cloud.

9.2.7 Customers can choose to optimise for value extraction from a given cloud environment

(204) For some organisations, access to cutting-edge functionality, security, and regulatory compliance, without the need to undertake a lengthy procurement process, will be more important than the ability to move capabilities easily across clouds. This focus on breadth and depth of services is reflected in Ofcom's consumer research, which indicates that customers choose Azure due to the number of features (31%), service quality (39%), level of security (34%), and value for money (32%).[174]

(205) Such customers may choose to develop a cloud solution based solely (or largely) on a single provider by acquiring services up the cloud stack (i.e. IaaS, PaaS and SaaS offerings). This provides access to a suite of products that have been specifically designed to run with the infrastructure and services that make up the cloud stack.

(206) The trade-off for reduced portability would include (i) a more cost-effective solution, given that no, or minimal, modifications to the services or middleware are required; (ii) simpler data

[172] Supported Cloud Platforms | Snowflake Documentation.

[173] New Cross-Cloud Capabilities with Snowgrid - Snowflake.

[174] Q25 Ofcom Customer Survey.

governance, security and compliance processes; and (iii) seamless access to more powerful, productivity-enhancing products (and the full range of features these offer).

(207) If a customer wishes to switch cloud providers or wishes to run a workload across clouds simultaneously, they can reconfigure their cloud environment to retroactively achieve portability. In this case, costs are back-loaded and incurred at the time of switching or implementing multi-cloud. We provide a case study below of Brightline's switch from Azure to AWS below.

Case study 4 - Brightline – Full switch

BrightLine migrated its data warehouse, database and application from Azure to AWS in order to improve scalability, increase efficiency, and reduce costs. BrightLine started its migration with its highest-cost service, its data warehouse. BrightLine migrated its data warehouse to Amazon Redshift, and also migrated its .NET application to Amazon EC2 compute, and its database instances to Amazon RDS for SQL Server. After having migrated its data warehouse, BrightLine migrated its database to Amazon DynamoDB. As part of that migration, BrightLine also started using Amazon Simple Storage Service (Amazon S3).

The migration process took Brightline 6 months and reportedly led to monthly cost reductions of over 70%, a more scalable cloud solution and improved customer satisfaction.

9.3 Azure actively facilitates interoperability and portability and thereby mitigates trade-offs

(208) Microsoft recognises that customers are increasingly deploying multi-cloud strategies and no cloud will be competitive if it does not support customers' requirements for flexibility.[175] To this end, Azure has been deliberately designed to be open and accessible, with support for customers and developers in the form of tools to assist switching and develop multi-cloud, information on each of its APIs to facilitate communication between cloud services, and support for cross-platform third party applications (including on the Azure Marketplace).

(209) Microsoft has a strong commercial rationale to make its services widely accessible, both because this allows them to win customers from competing cloud service providers – particularly AWS, which has a significantly larger share of the cloud market, and because making APIs available for developers to build on top of ultimately enhances the offer to Azure customers. [✂]. In addition, as noted above, Microsoft generates revenue every time its APIs are called, even from a different cloud architecture; it would be commercially irrational to limit this.

(210) As the Interim Report recognises, Microsoft offers information and tools to mitigate the technical barriers to switching and multi-cloud, including the use of open-source software and open standards which are available that are available on other cloud providers.[176] We provide further details of various tools Azure supports in Annex 4.

9.4 Microsoft does not have a strategy to frustrate interoperability

(211) The Interim Report relies on some third-party commentary to assert that customers face high costs and substantial technical effort when developing multi-cloud solutions that seek to interoperate with each other due to the specific practices of Microsoft.[177] Specifically, it highlights two Microsoft services which it asserts are not interoperable: 'Azure Stream Analytics'

[175] No one customer is entirely in cloud or entirely in just Azure, either. For Azure to be the world's computer, we must offer a tapestry of solutions that address the mixture of needs that all customers have. See Cloud trends show customers increasing investments in hybrid and multicloud - The Official Microsoft Blog.

[176] Interim Report, paras 5.25 and 5.75.

[177] See, for example, Interim Report, paras 5.47 to 5.48.

and 'IoT Hub', and notes that "extra steps" are required to transfer data from Azure storage to Google's BigQuery. Microsoft does not consider it accurate to characterise these services as not being interoperable. In relation to these:

(i) The example of Azure Stream Analytics

(212) Azure Stream Analytics[178] is a stream processing engine that is designed to analyse and process large volumes of streaming data with sub-millisecond latencies. Given the purpose of the service is to provide end-to-end pipeline, it simply must have a defined source for how the data is being streamed, and these are configured to be Azure services.

(213) Despite this, the service is still extendable to other clouds through connectors to other services for interoperability. For example, one input into Azure Stream is the Event Hub, which can be configured itself to take data from other sources, such that the customer could use Azure Stream Analytics on data it first gets in one of its applications running in a different cloud. Like all its services, Microsoft is continually updating Azure Stream Analytics. At this stage of its lifecycle it [✂]. Microsoft continues to explore possible improvements to enable additional ways to interoperate and plans to support Kafka protocols (a unified, high-throughput, low-latency platform for handling real-time data feeds) in the future to provide even more ways to interact with Azure Stream Analytics from other applications, including those running on other clouds.

(ii) The example of IoT Hub

(214) The IoT Hub[179] provides a cloud-hosted solution to connect edge devices with the cloud. A key aspect of the solution is to offer an SDK for the development of Edge devices to enable them to communicate with a particular secure back end. In the case of Azure IoT Hub, the SDK that Microsoft offers enables customers to direct those devices to send data to Azure. Other providers of IoT services would similarly provide their customers with tools to configure their Edge devices to communicate directly to their clouds.

(215) While it is true that Microsoft did not develop the SDK to enable devices to instead communicate directly to AWS or GCP (indeed Microsoft is not even best positioned to create SDKs for clouds other than Azure), this does not mean that the application is not interoperable with other cloud providers' services. Once the data is in Azure, the customers can choose to send that data across to another cloud for processing if desired. The APIs in Azure Data hub enable customers to connect solutions with it to allow the flow of that data to any other end point of the customer's choosing. Customers are not required to use this application to send data to the Azure cloud and are free to engineer their own solution if desired.

(216) This can also work in the other direction where another cloud collects data from the Edge and then shares it into Azure IoT Hub. For example, Particle[180] allows customers to connect their edge devices to the Particle cloud, which then pushes data onto Azure via the IoT hub. Particle integrates with all major cloud service providers, including AWS, Snowflake, GCP, etc.

(iii) The example of "extra steps" required in using Google's BigQuery

(217) Ofcom notes that customers using Azure storage products may require "extra steps" to load data from Azure into other cloud provider's services such as Google's BigQuery.[181] This is simply not accurate. Google's own documentation lays out the steps necessary for internal and external (Azure Blob Storage and Amazon S3) data transfers to BigQuery, which both take only

[178] Azure Stream Analytics | Microsoft Azure.

[179] IoT Hub | Microsoft Azure.

[180] Particle Cloud – Scalable Hosted Cloud Service for Connected Products.

[181] Interim Report, paras 5.24 and 5.25.

five steps. It also offers the BigQuery Data Transfer Service, which automates data movement into BigQuery on a scheduled, managed basis (including from External cloud storage providers Amazon S3 and Azure Blob Storage and from data warehouses Teradata and Amazon Redshift).[182] As set out in the case study in Section 9.1 above, the additional (albeit minimal) steps required to port data is a direct result of innovation and competition between cloud providers.

9.5 Conclusion on interoperability and portability

(218) Notwithstanding inherent trade-offs, there is clear evidence that interoperability and portability are not substantially holding back switching between cloud environments and multi-cloud deployment by customers. According to Ofcom's survey a majority (60%) of respondents have either switched IaaS/PaaS providers in the past or taken on an additional IaaS/PaaS provider.[183]

(219) This statistic is also supported by a range of industry reports as detailed above and is consistent with Microsoft's own experience. Microsoft does not generally have visibility on whether customers are using multiple clouds on a workflow-by-workflow basis. However, we have observed vast numbers of customers migrating new workflows, taking on secondary providers, and deploying workloads across multiple clouds. We provide a representative list of anecdotal examples of customers doing just this in Annex 5.

(220) Moreover, the drivers of interoperability and portability "issues" are also fundamental to competition between providers. For example, as Ofcom notes, the design of proprietary technology is ultimately beneficial to customers and allows cloud providers to offer new innovative products.[184] Cloud providers invest heavily in innovation to differentiate themselves, which brings inevitable complexity in architecture. While simpler and more open architectures are possible, they will be inherently limited compared to more complex "tightly integrated" services. A tangible illustration of how requiring strict compliance to standards can chill innovation can be illustrated by the Open Document Format. The same concept would apply to cloud IaaS and PaaS services.

Case study 5 – Lessons from the Open Document Format

Microsoft agreed to implement the Open Document Format ("ODF") as one of the file types supported in Microsoft Word under the terms of an interoperability undertaking with the EU.[185] The intention of the new ODF file type was to ensure that customers could easily port files between Microsoft Word and other competing office applications that implemented ODF.

For a file type or application to be portable in practice, the features available on ODF would have to be the same across the relevant applications. While ODF can be read by a wider variety of programs, it can only utilise features available in each program and enabled by the ODF Standard.

Microsoft continued to innovate, adding new functionality to Word, and introduced the "track changes" feature. ODF, however, did not (and still does not fully) support the ability to save changes tracked in a document into the ODF format. For example, there is no mechanism to reconcile markup text in a table.[186] Thus, if Microsoft were limited to only supporting ODF and

[182] What is BigQuery Data Transfer Service? | Google Cloud.

[183] Interim Report Annex, para A7.32, excluding private cloud users.

[184] Interim Report, para 5.62.

[185] Microsoft EU Interoperability Undertaking 2009.

[186] Track changes - The Document Foundation Wiki.

Case study 5 – Lessons from the Open Document Format

its feature set, the innovation offered by tracking changes in Word could not have been offered.

(221) In summary, customers are highly sophisticated and make deliberate choices about how to structure their cloud architecture to optimise it for their own commercial needs. The fact those choices result in a level of lock-in (i.e. give rise to switching costs, incurred either upfront or backloaded) does not indicate the choice is not free, nor that there is a lack of competition.

(222) While customers may have a genuinely reported preference that all systems could work seamlessly together, such a simplistic "desire" belies the feasibility of facilitating "frictionless" operation of disparate systems, which would in fact – as history has shown – inevitably lead to standardisation around the "lowest common denominator".

(223) Given the dynamic nature of competition in this area, Microsoft has serious concerns that any intervention based on woolly concerns and generic feedback will stifle innovation. This would sacrifice (degrees of) dynamic competition and innovation which has served cloud customers so well thus far.

(224) Consequently, while Microsoft is more than willing to discuss very specific concerns to see how they might be specifically addressed if this is not already the case, it does not support an MIR and the premise of ill-tuned intervention based on vague interoperability concerns.

10 Egress fees

10.1 Due to competition, Azure cost recovery for data transfer occurs via egress fees

(225) Microsoft understands why Ofcom has focused on egress fees. In considering this issue, context is critically important.

(226) The "competitive price" for data ingress into cloud environments has been zero since 2011. Historically, AWS charged ingress fees of $0.10 dollars per gigabyte for incoming data transfer when it was the sole credible supplier. Shortly after Microsoft Azure entered the cloud infrastructure services market and reduced ingress fees to zero, AWS also made ingress free to users, almost certainly in direct response to competition from Azure for new workloads.[187] Due to continuing competitive pressure (in particular, in relation to the intense competition for the next workload migrating to or switching cloud), none of the cloud infrastructure providers today charges ingress fees. Moreover, it would not seem rational to charge customers for migrating their workload (and, thus, generating business) to a provider's cloud, disincentivising them to use the cloud service.

(227) At the same time, subject to thresholds, Azure customers pay egress fees to Microsoft for the transfer of their data outside of Microsoft Azure's cloud infrastructure platform to a third-party destination (e.g. to another cloud infrastructure platform, to on premises platform or a customer's in-house solution).[188] Unlike for ingress fees, Microsoft is unable to recover the costs for providing egress services via the price of Azure services / products as customers are, by definition, leaving the Azure platform at the point of receiving such egress services. Again, due

[187] Interim Report, para 5.119 and Dutch cloud market study, pg. 57.

[188] For the avoidance of doubt, references to "egress fees" in this do not include data transfer fees charged for data transfers within Microsoft Azure's cloud platform between different "Availability Zones".

to competitive pressure, Microsoft broadly sets its egress fees to be competitive [✂], and the current egress charges and free tier thresholds for [✂] and Azure[189] are similar:

(i) Microsoft offers a free monthly egress tier (100GB) that allowed some [✂] of Microsoft's active customers in the UK in 2021 to be free of any egress fee charges during any specific month because they were within this limit.[190]

(ii) Microsoft offers discounted egress fees when customers transfer data to other cloud infrastructure providers within the "bandwidth alliance" (which includes members such as Google Cloud, Oracle Cloud, Alibaba Cloud and Tencent Cloud, amongst others).[191]

(228) Microsoft also supports alternative methods for transferring data than via the public internet. For example, Azure ExpressRoute offers outbound data transfer in the UK efficiently and cost effectively (including on a fixed price unlimited data plan) to on premises networks, other cloud providers, and data centres.[192]

(229) In the above context, while competition has kept ingress fees at zero, egress fees at zero below certain volume thresholds (free tiers), and to certain discounting above those egress thresholds, bandwidth for data transfers (both ingress and egress) have real underlying costs for suppliers including significant fixed investments in networking infrastructure, storage solutions, data centres, and middleware (e.g. security, backup). Microsoft globally invests [✂] each year in the infrastructure necessary to carry out ingress/egress data transfers[193] and R&D to improve performance and lower costs for data transfers.

(230) In a market where bandwidth usage via data transfer (in any direction) incurs real costs but where competition has driven suppliers to offer ingress for free, it should not be surprising that cost recovery occurs via egress fees. Self-evidently, any analysis of the levels of bandwidth investment and usage cost recovery (fixed and variable) is a complex one and must be nuanced. The Interim Report suggests that egress fees "*appear only loosely related to cost*" (para. 5.109) and proceeds to consider variable costs of transit fees (5.111-5.113) or only fixed costs "*uniquely attributable to data transfer*" (5.116-7) and that the "*actual incremental cost of providing external data transfer is likely to be well below … total cost*" (5.117) although it acknowledges the need for further analysis for the Final report.

(231) It is correct that, in the ordinary course of business, Azure has seen no value in [✂].

10.2 Azure's egress fee structures are transparent and clear

(232) In some cases, businesses might not be able to predict how their demand for cloud services and data egress might evolve in a given year. To help businesses understand the potential costs, Azure's egress fees are listed clearly on its website, and the only key variable from the customer's perspective is their usage/volume of data.[194]

(233) The total egress fees charged to a given customer depends on a range of factors, including: (i) the volume of data being transferred by the customer (with higher volumes reducing the cost

[189] Prices decrease from ~£0.07/GB to ~£0.04/GB as the volume of data transmitted goes up from 100GB to over 500TB. Microsoft also does not charge customers to put data into or take data out of its Office 365 and Dynamics 365 services. See Pricing - Bandwidth | Microsoft Azure and EC2 On-Demand Instance Pricing – Amazon Web Services.

[190] Microsoft response to question 7 of Ofcom's Follow-up Questions Dated 27 January 2023.

[191] See Interim Report, para 5.110; and Bandwidth Alliance | Reduce Data Transfer Fees | Cloudflare.

[192] Pricing – ExpressRoute | Microsoft Azure; When to Use ExpressRoute Local for Microsoft Azure Private Peering | Megaport; ExpressRoute – Virtual Private Cloud Connections | Microsoft Azure; Connectivity to other cloud providers – Cloud Adoption Framework | Microsoft Learn.

[193] See *Microsoft Response to Call for Inputs: Cloud infrastructure services market study 3 November 2022*, pg. 2.

[194] As acknowledged by Ofcom at para 5.101 of the Interim Report.

per GB);[195] (ii) the type of network being used to transfer the data (e.g. if it is being routed over the Microsoft Premium Global Network or via Routing preference transit ISP network); and (iii) the geographical region/continent between the existing and the destination location of the customer's data.[196]

(234) Microsoft also actively assists customers to forecast and to control their future expenditure on Azure. For example, Microsoft provides 'Azure Cost Management + Billing' (which analyses monthly charges for different Azure features and projected cost over time) and 'Usage and estimated costs' (which provides a listing of monthly charges for different Azure Monitor features and shows costs under different pricing tiers). Microsoft also assists by training customer employees, and by providing cost optimisation/forecasting solutions (e.g. as part of the Azure Migration and Modernization Program[197]).

(235) Importantly, Microsoft's customers (such as [✂], amongst many others) are typically sophisticated customers who will be able to utilise the tools available to them (including the expertise of third-party IT consultants) and that can assess the total cost of ownership across the lifetime of the cloud services (including determining the cost of data ingress and egress).

10.3 Ofcom's theory that egress fees inhibit switching and/or multi-cloud is mostly conjecture and is inconsistent with its other findings

(236) The Interim Report's findings that egress fees present a material barrier to switching[198] bear poor relation to the reality of Microsoft's interactions with its customers. In Ofcom's own survey only 6% of respondents said that egress fees constituted the largest challenge to switching between cloud infrastructure providers.[199]

(237) Egress charges for enterprise customers are a small percentage of their total cloud costs across their cloud journey. In most cases, customers will not pay any egress fees in a given month (as set out in paragraph (227)(i) above). Where a customer wishes to switch cloud providers, the costs are also likely to be low relative to their cloud spend. An IDC Report for the EC, which compared the hypothetical one-time egress fees charged to a company that needed to retrieve all of their data in case of a cloud provider switch. This cost was found to be approximately only 0.30 – 0.35% of the expected annual operating cost using either Azure or AWS for cloud services.[200]

(238) Moreover, the current level of egress fees does not preclude or disincentivise Microsoft Azure customers from using more than one cloud infrastructure provider under a multi-cloud strategy. Most cloud customers purchase cloud infrastructure services from more than one provider and have the ability to ingress an unlimited amount of data and egress up to 100GB for free (per month). This encourages customers to build and run multi-cloud scenarios involving some data ingress and egress without incurring bandwidth charges.

[195] As explained in Microsoft's responses to Ofcom's follow-up questions dated 27 January 2023, Microsoft has lower per-GB egress fees for higher data egress volumes because fewer larger egress workloads are easier to manage than a larger number of smaller bandwidth usage workloads. Like many volume discounts, the price difference is driven by the higher cost of managing more customers.

[196] As explained in Microsoft's responses to Ofcom's follow-up questions dated 27 January 2023, Microsoft sets its prices for different data transfers for different regions based on a number of considerations, [✂]. For example, egress from South America to any other destination outside of South America is typically more expensive than egress from North America or Europe to any other destination. See Microsoft's website for further details: Pricing – Bandwidth | Microsoft Azure.

[197] Azure Migration and Modernization Program | Microsoft Azure Azure Migration and Modernization Program | Microsoft Azure.

[198] Interim Report, para 5.82.

[199] Interim Report, para 5.84.

[200] "Switching of Cloud Service Providers", European Commission, 2018, page 50-56.

(239) The consistent feedback from Microsoft's business teams is that customers do not generally refer to egress fees as a significant factor in customer choice between cloud providers and they do not inhibit customers' ability or incentive to switch providers or to multi-cloud strategies (i.e. they do not lock-in customers to Azure). This is in part because customers are unlikely as part of a multi-cloud strategy to design a cloud infrastructure system requiring significant and regular data transfers between different cloud infrastructure platforms (thus potentially incurring egress fees) for various reasons such as network/data latency, data privacy/security and cloud system resilience. Moreover, egress fees are only a small share of the costs incurred in switching a workload to a different cloud provider, and rival cloud providers frequently provide assistance to mitigate customers' costs of switching cloud provider (whether by financial means or by technical assistance or by other means).

(240) As regards Ofcom's findings that egress fees are a concern to customers who would like them to be reduced/removed,[201] it is not a surprise to Microsoft that paying customers of cloud infrastructure services would like a certain aspect of their fees to be reduced or removed but this does not represent a rounded conclusion on cost/benefit and the need for cost recovery in the cloud environment.

10.4 Conclusion

(241) Data storage and transfer has a cost and therefore in Microsoft's view, an element of charging for data transfer (including between Azure regions) is important for the purposes of efficient infrastructure usage. Given the size of untapped opportunity for cloud migration, it is not surprising that competition has driven ingress fees to zero and pushed cost recovery to subsequent stages of the customer cloud lifecycle. Customers consider egress costs in designing the flow and usage of data in their solutions, and thus charging for data aligns the customer and provider incentives to seek the most efficient organisation of such data flow.

(242) An intervention that made data egress/transfer "totally free" would result in excess (suboptimal) usage, while price regulation raises the concerns that Ofcom already identifies in its Interim Report.

[201] Interim Report, paras 1.21 and 5.82.