

I. Introduction

Oracle appreciates the opportunity to offer our perspective on the cloud services market in the UK. Oracle offers the following consolidated commentary to the CMA's Updated Issues Statement (published 6 June 2024), its working papers (published 23 May 2024 and 6 June 2024), and the Market investigation qualitative customer research final report (published 23 May 2024) (hereinafter the "Jigsaw Report").

- a. Egress fees are profit-driven barriers to both switching and multi-cloud.

In its Updated Issues Statement, the CMA recognizes the importance of egress fees in creating barriers to switching and multi-cloud.

Egress fees represent a classic example of extractive profits. The egress fees charged by the largest hyperscale cloud providers do not correspond to costs incurred, as evidenced by a comparison to Oracle's approach, based on a cost-recovery pricing model. Oracle's pricing strategy permits Oracle to offer highly competitive services while ensuring customer flexibility. As the CMA notes, "egress fees may not be reflective of the costs of providing egress for some cloud providers." (Updated issues statement ¶153)

More specifically, "egress fees" as charged by CSPs represent a blended estimate of essentially two variables: (i) cost of moving the data from one place to another, plus (ii) profit margin. Because, as a technical matter, it is impossible for a CSP to know where the data will end up – a CSP can only know that the data is leaving its environment, not the ultimate terminus – the CSP cannot determine with 100% accuracy how much cost the CSP itself will incur for each data transfer in real time. The fixed rate of "egress fees" is therefore a calculated assessment of how much it costs the CSP in terms of capital investment, intermediary charges (i.e., to an ISP or other network provider), plus the amount of profit a CSP seeks to make on the transaction added on top.

By pursuing a cost-recovery pricing model, Oracle has made the business decision that it will not extract *unreasonable* profits from its customers for moving their own data – which also demonstrates that this is a business choice, not a technical one.

Take peering as an example. When one CSP exchanges data with another directly, that is called "peering." Peering incurs lower costs to the transferor CSP because these data flows bypass ISPs and their associated costs entirely. (CSPs typically – but not always – must pay an ISP to use their networks to transfer data between point A and point B.)



As the CMA has noted, “Cloud providers say they incur different costs based on location/peering provider for transfers and routing.” (Egress fees working paper, Appendix A ¶A.8)

Yet high egress fees are still charged to customers moving data between CSPs using peering networks, and such cost savings are not passed on to the customer in reflecting the lower cost of data transfer that the CSP itself incurs. This suggests that other CSPs’ business decision to charge a much higher egress fee is based on a blended price estimate that includes a significantly greater profit margin variable.

The CMA acknowledges part of this in its Egress fees working paper when it states that “peering charges are typically variable (eg \$/mbps) whereas asset costs are fixed.” (Egress fees working paper ¶3.13) But OVHcloud’s position that “the only technical justification for egress fees is the cost for using internet bandwidth from ISPs or from other operators” also misses market realities. (Egress fees working paper ¶3.14) A provider like OVHcloud – which may not have built its own global infrastructure – does not account for capital expenditures on physical cables and server racks that are appropriate to include in the estimated costs to a CSP.

If egress fees for all CSPs were based on a cost-recovery model (as with Oracle), the costs of egress fees should be variable based on the CSP’s actual underlying costs. Instead, excessive egress fees are used by some CSPs as an easy way to extract unreasonable profit margin while also serving as a gating factor, helping to lock in their customers.

For example, the CMA found that Microsoft does not pay any ISPs for data transfers in the UK, (e.g., Egress fees working paper ¶3.60), and Microsoft itself submitted that it only incurs “incremental” interconnection costs for transfers over the open internet. (Egress fees working paper Appendix A, ¶A.10) Likewise, the CMA is correct to highlight that even when AWS’s indirect costs such as sales and marketing were excluded, AWS’s average revenues and list prices for egress still exceeded its costs. (Egress fees working paper ¶¶3.46, 3.56) The CMA should continue to investigate whether CSPs account for such savings when calculating egress fees charged to customers.

The CMA should take steps to establish standardized terminology across CSPs with respect to more than egress fees, to ensure future compliance is more than in name only (i.e., a CSP should not be able to say it is compliant by not charging “egress fees” but instead charging “data mobility fees”).

In this, we agree with IBM's position that "if the costs incurred by a cloud provider are not recovered in one form, they will likely be recovered elsewhere, in order to allow the provider to achieve an economic return." (Egress fees working paper ¶4.23) The CMA should pay special attention to this practical insight: while CSPs can and should make a profit, CSPs should not use regulators' attempts to enforce free and fair competition as an exercise in creative word choice. For this reason, we fully support the CMA's proposed remedy of requiring CSPs to "adopt consistent terminology in relation to egress fees" *and* to "display egress fee prices prominently on webpages and/or in contracts." (Egress fees working paper ¶4.91)

Finally, in focusing on egress fees as a percentage of customer spend, **the CMA underestimates the importance of the absolute value of egress fees to the customer.** This fails to account for "sticker shock" at the highest levels and among the largest cloud computing customers: [0-5%] of spend sounds much smaller than, for example, \$10M just to switch providers. And no customer negotiates on egress fees against their total cloud spend; egress fees are *additional* costs above the budgeted total cloud spend. From the customer's perspective, their cloud spend typically increases gradually over time. By the time a company's cloud computing strategy has matured, a CIO will be hard-pressed to recommend to the CFO that the company should spend millions on egress fees in hopes of receiving better returns and cost savings in the long term. As a general matter, we note that large customers with whom Oracle has negotiated for switching workloads to OCI from a competing CSP have cited egress fees of their legacy cloud provider as a concern.

- b. The technical barriers to both multi-cloud and to switching can be overcome.

First, as evidenced by Oracle's decisions to enter into interconnect agreements with both Microsoft¹ and Google², technical barriers to multi-cloud are surmountable. They require engineering resources, to be sure, but we emphasize that it is a rational business decision to ensure customers have access to multi-cloud solutions. The CMA is correct to identify that, as of today, "customers must invest extra effort and resources to mitigate lock-in," (Updated issues statement ¶41) though CSPs can certainly make choices to facilitate and ease that additional burden. The CMA's conclusion that "customer or provider-led mitigations may not effectively support efficient multi-cloud and switching for customers" (Updated issues statement ¶41) may not fully account for the efforts Oracle (and others) are making to enable customer choice above all. Some CSPs are actively making it easier for customers to

¹ <https://www.oracle.com/cloud/azure/interconnect/>

² <https://www.oracle.com/news/announcement/oracle-and-google-cloud-announce-groundbreaking-multicloud-partnership-2024-06-11/>

support multi-cloud architectures with technical and business agreements designed to reduce customer pain points. Moreover, as more businesses move to the cloud, they may not need or be able to afford a multi-cloud architectures at the outset. Once their understanding of their cloud computing needs and strategy reaches maturity, however, they may be too entrenched with their first provider to justify spending on multi-cloud (thereby increasing single vendor lock-in). It is critical to ensure multi-cloud is accessible and free from anticompetitive barriers at all stages of a company's IT modernization timeline.

On the other hand, some CSPs use their dominant position to lock customers into using a single provider. For example, AWS's recent announcement purporting to "waive" some egress fees would not apply to a multi-cloud architecture, and *only* applies to switching costs – AWS will only grant credits for "data transfer out to the internet (DTO) charges when you want to move outside of AWS"³. AWS's offer is therefore irrelevant to a customer seeking to multi-cloud under the CMA's framework.⁴ Further, a customer cannot avoid egress fees if they seek to migrate only some services or workloads from AWS to a competitor.⁵

As for switching, some CSPs have done little to make this a viable business decision for their customers. Even if a customer qualifies for AWS's offering and provides the required notice to AWS, the 100GB threshold that AWS has established may operate as an additional driver of lock-in for the largest and most profitable of AWS's customers.⁶ If a customer is large enough to be moving more than 100GB per month out of AWS and to a competitor, the customer is necessarily moving a great deal of data – and is likely complex enough to require a multi-cloud strategy. AWS's largest customers may not be able to justify the proportionally higher switching costs of moving such large datasets, particularly if the credits are far from assured. The CMA should be particularly careful not to accept at face value the "voluntary commitments" some CSPs may be championing in other jurisdictions (Updated issues statement ¶154), as they are crafted to cause minimal impact on their most lucrative anticompetitive practices.

³ <https://aws.amazon.com/blogs/aws/free-data-transfer-out-to-internet-when-moving-out-of-aws/>

⁴ AWS's new policy would likely not apply to multi-clouding under any of the CMA's three formulations: (i) cloud duplication; (ii) integrated multi-cloud; or (iii) siloed multi-cloud. (Competitive landscape working paper ¶3.18)

⁵ AWS's FAQ also states: "If you only want to move your total usage of a single service, but not everything, contact AWS Customer Support." There are no further details available about whether AWS would ever waive egress fees in the case of a customer seeking to multi-cloud.

⁶ We note that Oracle offers up to 1TB per month of free data mobility, without requiring customers to go through an onerous and opaque application process. See, e.g., Kevin Bogusch, *Cloud Data Egress Costs: What They Are & How to Reduce Them* (Jan. 24, 2024), <https://www.oracle.com/cloud/data-egress-costs/>.

- c. AI is the next frontier for hyperscalers to raise barriers to entry and expansion.

The CMA should continue to investigate the effects of CSPs’ “significant further investment . . . in accelerator chips (eg GPUs), to meet growing demand for AI services.” (Updated issues statement ¶129(b)) GPUs are a technology that continue to evolve rapidly, and customers deserve the flexibility to move their data to providers with the specific technology they require. Lowering barriers to interoperability (such as egress fees) will ensure increased competition and mitigate the risk of creating more walled gardens.

As noted in the Jigsaw Report, one satisfied Azure customer is quoted as saying that “if one cloud provider came up with an [AI] capability that we wanted to exploit that would have a business benefit, we would act. But we could just have a multi-cloud strategy, right ?” (Jigsaw Report p.40) This forecasted choice – to adopt a multi-cloud strategy for AI in particular – does not exist in the current market, but it should. This customer recognizes that to take advantage of best-in-class AI solutions, they require a multi-cloud architecture. Further, as CSPs have differentiated offerings, FM developers recognize that multi-clouding is key to competition and innovation.⁷

We note that the importance of egress fees will only increase with the growth of AI, as running AI in the cloud both requires and produces a voluminous amount of data. Moving those enormous datasets between providers may be cost-prohibitive due to egress fees. The CMA should continue to pursue remedies to lower the existing barriers to multi-cloud.

We also agree with the CMA that “[p]artnerships between the large cloud providers and FM model developers are extensive and are likely to play an important role going forwards.” (Updated issues statement ¶130) But not all “partnerships” are created equal. Most of the existing CSPs have entered into exclusive deals, or have acquired, or have all-but-acquired, at least one of the newest and most innovative companies pursuing FM development. Microsoft’s relationship to OpenAI and InflectionAI, or AWS’s relationship to Anthropic, or Google’s acquisition of DeepMind are all examples of “partnerships” with some level of exclusivity that ties up access to these models to the use of the CSP’s cloud service. Whatever their legal status, these “partnerships” undoubtedly threaten to further entrench the biggest hyperscalers at the expense of smaller would-be competitors. Oracle, by contrast, has prioritized customer choice by entering into non-exclusive agreements with multiple FM model developers, from

⁷ E.g., Oracle, *OpenAI Selects Oracle Cloud Infrastructure to Extend Microsoft Azure AI Platform*, <https://www.oracle.com/news/announcement/openai-selects-oracle-cloud-infrastructure-to-extend-microsoft-azure-ai-platform-2024-06-11/>.

Cohere⁸ to xAI⁹ and many more. Oracle believes that its customers benefit most when they have a choice to use any and all FM models in the market. Increased competition among FM developers and between associated CSPs also incentivizes continued innovation in this exciting new technology.

II. Conclusion and Recommendations.

a. Scope of the CMA's Analysis

We thank the CMA for its comprehensive and thoughtful analysis of the UK cloud computing market thus far.

We understand from industry press that the public sector was excluded from the CMA's research phase. Regardless, we trust that the CMA will include the public sector and central government in its analysis of the cloud sector, as these entities represent a significant source of revenue for cloud providers in the UK. (Ofcom Cloud Services Market Study, Final Report, ¶3.87) This is important given Ofcom's preliminary findings in relation to the public sector. In particular, Ofcom noted that "the leading positions of AWS and Microsoft on G-Cloud Framework appear similar to our [Ofcom's] findings for the wider market." (Ofcom Cloud Services Market Study, Final Report, ¶3.91) We look forward to learning more about what points of friction public sector entities have experienced in their efforts to build best-in-class, efficient cloud computing systems.

b. Remedies

As the CMA well knows from previous enforcement actions, transparency is necessary but not sufficient to deter or remedy anticompetitive conduct. When combined with toothsome enforcement, however, a regulator can enforce standards, transparency, and pro-competitive policies to mitigate the risk that Big Tech will cannibalize yet another emerging market.

There is also a very real skills gap identified throughout the CMA's reports, as well as other regulators' inquiries (e.g., the European Commission¹⁰). The EU has also identified a skills gap in cloud computing adoption as part of its Digital Decade targets. In a well-functioning market, one might expect a robust set of facilitator companies to

⁸ <https://www.oracle.com/customers/cohere/>

⁹ <https://www.datacenterdynamics.com/en/news/elon-musks-xai-to-spend-10bn-on-oracle-ai-cloud-servers-report/>

¹⁰ E.g., European Commission, *Realising Europe's Digital Decade* (21 June 2024), https://ec.europa.eu/commission/presscorner/detail/en/FS_24_1392.

grow up around bridging the gap between CSP offerings on behalf of customers. For example, as the Jigsaw Report noted, several customers identified as a “Microsoft shop,” and therefore did not see the benefit of entertaining other CSPs’ services. (Jigsaw Report p.29) This secondary, theoretical consulting market – where technically advanced intermediaries help design and integrate multi-cloud architectures to improve service, lower customer cost, and mitigate the risk of lock-in for customers – does not exist. This strongly suggests a market failure; customers tend to choose a single CSP from the beginning of their transition to cloud, and face barriers such as high egress fees that limit business incentives to switch or to multi-cloud.

And finally, we reiterate that egress fees that do not reflect actual costs to a CSP of moving data from one place to another represent a clear example of extractive profits that exacerbate barriers to entry and to switching – and are not merely a profitable business decision. The CMA should continue to carefully investigate and standardize language across CSPs to ensure that “egress fees”-by-any-other-name do not become a market reality.

c. Closing

Oracle appreciates the CMA’s efforts to investigate this important and growing market in the UK. Oracle reiterates its willingness to assist the CMA in its ongoing efforts, and we look forward to the next stage of the CMA’s work.