Economic response to the CMA's Competitive Landscape Working Paper

CMA cloud services market investigation

Prepared for

Amazon Web Services, Clifford Chance

Prepared by

CRA

Charles River Associates
8 Finsbury Circus
London EC2M 7EA

Date: 2 August 2024

CRA Charles River Associates

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

1. INTRODUCTION AND EXECUTIVE SUMMARY

1. This report sets out our economic response to the CMA's Competitive Landscape Working Paper dated 23 May 2024 ("CL WP").[1] We appreciate that the CL WP presents the CMA's emerging views based on its ongoing evidence gathering process, and that the CMA has not yet reached any provisional findings. Against this backdrop, our response provides the CMA with further evidence to demonstrate that competition for the supply of IT services, including cloud services, is in fact functioning well, both globally and in the UK.

2. The economic analysis presented in this response is complementary to our previous submissions and, as such, should be read in conjunction with the following: [✂].

3. We do not repeat our previous submissions or evidence in this report. However, we consider that our previous findings remain valid and are further enhanced by the additional analyses presented in this report.

4. The CMA's CL WP investigates the supply of public cloud infrastructure services (cloud services) in the UK. The CL WP analysis focuses on (i) market definition, (ii) shares of supply, (iii) customers' ability to switch and multi-cloud, (iv) barriers to entry, (v) innovation, and (vi) profitability. The CL WP's preliminary conclusions are that:

- Traditional IT services (i.e. on-prem) are unlikely to exert competitive pressure on cloud solutions.[2]
- The cloud services segment is concentrated, and concentration is increasing over time.[3]
- Multi-cloud is limited, mostly found amongst larger customers. Switching is also infrequent.[4]
- Barriers to entry are significant. Large cloud service providers benefit from economies of scale, which represent a barrier to entry and expansion.[5]

---

1 CMA (2024). Cloud services market investigation: Competitive landscape working paper.

2 "*[…] traditional IT should be treated as separate from the markets for IaaS, PaaS and SaaS […]*" (CMA's CL WP, paras 4.78).

3 "*The overall cloud services sector is concentrated, and concentration is increasing over time.*" (CMA's CL WP, para. 5.21).

4 *"However, the evidence suggests that, while there is some degree of multi-cloud use, it may be quite limited in scope and mostly found amongst larger customers."* (CMA's CL WP, para 9.6).

5 *"[…] large cloud providers benefit from economies of scale. […] when considered together they [the factors generating scale effects] represent material barriers to entry and expansion.*" (CMA's CL WP, para. 7.43).

- The evidence on innovation is inconclusive.[6] The fact that AWS and Azure keep investing significantly raises barriers to entry for rivals, who will also need to invest.[7]
- AWS and Azure have been generating returns from their cloud services above their cost of capital.[8]

5. We find that the preliminary conclusions of the CL WP are incorrect and not substantiated for the following reasons:

- **On-prem is exerting competitive pressure on cloud services providers.** For many workloads, customers consider on-prem solutions as a competitive alternative to cloud services. The competitive pressure exerted by on-prem providers can be captured by two metrics: (i) the share of customers *switching workloads away from a cloud service provider's public cloud* that have moved workloads to on-prem ("diversion ratios"), and (ii) the share of *all customers* choosing on-prem over cloud ("cloud repatriation rate") for workloads. Estimated diversion ratios from cloud to on-prem are between 30% and 65% while estimated rates of cloud repatriation are between 24% and 36%. The high magnitude of these ratios shows that it is incorrect to exclude on-prem from the competitive assessment. **(Section 2)**
- **If we ignore on-prem and consider the narrow "cloud infrastructure" segment defined in the CL WP, we find that concentration has been decreasing**. Cloud infrastructure alone is not the relevant market definition because it ignores on-prem solutions which exert competitive pressure on cloud services as many customers consider them to be viable alternatives. Moreover, even when ignoring on-prem solutions, the concentration index (as measured by HHI) has been decreasing in all three cloud segments considered by the CL WP. [✂]. This is also in line with recent survey evidence, in which the majority of IaaS and PaaS customers stated that there is a lot of competition in the cloud segment. **(Section 3)**

---

6 *"[…] it is difficult to assess the exact level of quality and innovation in cloud services (both over time and currently), and we do not currently consider that the evidence we have reviewed to date allows us to accurately compare quality and innovation across cloud providers"* (CMA's CL WP, para 6.22).

7 *"[…] cloud providers make significant investments in research and development.", "[…] while new entrants may be able to enter the markets without making significant investments in research and development initially, the scale of investment in research and development required to remain competitive may represent an ongoing barrier."* (CMA's CL WP, paras 7.40-7.41).

8 *"[…] AWS and Microsoft have been generating returns from their cloud services above their cost of capital, and that this could be expected to continue in the future*" (CMA's CL WP, para. 9.17).

- **Multi-cloud is widespread, suggesting there are no barriers preventing customers who would benefit from multi-clouding from doing so**. **Customers also show strong ability to move workloads.** The CL WP's results suggesting that multi-cloud is uncommon are likely severely biased, as its analysis is restricted to only three cloud service providers. Moreover, the CL WP has not provided any benchmark or threshold to determine when multi-cloud rates are low. Even when taken at face value, CL WP's own analysis shows material multi-cloud rates [✂]. There is no economic basis for such rates (i.e., of [✂]%) being lower than what would be expected in a well-functioning market. Moreover, our estimates suggest that if errors and biases in the CL WP's analyses are fixed, multi-cloud rates could be as high as [✂]%. These multi-cloud numbers are in line with various survey evidence and demonstrate the likely magnitude of the bias when restricting the analysis to only three cloud service providers. Further, the CL WP's conclusions on switching being limited are flawed as the CL WP does not provide a benchmark for a "healthy" switching level and bases its conclusions only on sparse interviews. Meanwhile, customers' ability to switch is confirmed by both AWS's data showing significant switching rates and third-party studies and surveys. **(Section 4)**
- **Evidence does not support that entry or growth of smaller cloud providers is impeded by barriers to entry or scale economies.** Evidence shows that cloud service providers can and do enter and expand at moderate costs, with many smaller providers growing faster than larger ones. The high number of entrants and growth of smaller providers in past years suggests that existing entry costs are unlikely to substantially affect the functioning of the cloud segment. Moreover, scale economies appear to be only of limited magnitude, such that they are unlikely to limit smaller providers' ability to compete in the segment. **(Section 5)**
- **The cloud segment is highly innovative, consistent with a healthy competitive landscape.** The cloud segment is characterised by constant innovation and investment, as shown by the growth in the number of services, increased functionalities, and patent activity. AWS itself has released a large number of services and features every year. AWS's innovative efforts also translate into lower costs, which are passed on to customers through lower prices. Observing that firms continue to invest in ways that benefit customers through lower prices and higher quality, as also found by the CL WP,[9] should not raise competitive concerns. On the contrary, investments are a sign of a healthy, functioning competitive landscape. **(Section 6)**
- [✂].

6. Finally, we note that certain data and calculations related to the shares of supply and profitability have not been disclosed to AWS, so it is not possible to exactly replicate the CL WP's analysis or comment on its accuracy. For example, the CMA has not shared the material related to 60 interviews conducted by the Jigsaw market research, which is often mentioned as qualitative evidence throughout its working papers- including the CL WP. Therefore, to respond to the issues raised in the CL WP, we have assessed the data available at hand, which supports the findings set out in this report.

9 CMA's CL WP, para 7.37.

2. ON-PREM EXERTS SIGNIFICANT COMPETITIVE PRESSURE ON CLOUD SERVICE PROVIDERS

7. The CL WP erroneously fails to consider traditional IT services (i.e. on-prem) in assessing the competitive landscape for the supply of cloud services and in doing so does not acknowledge that on-prem exerts significant competitive pressures on cloud services. Rather, the CL WP notes that: "*traditional IT is unlikely to be a close substitute for public cloud for a high proportion of customers*".[10] Moreover, the CL WP supposes that "*while public cloud may be a credible alternative for traditional IT customers, the evidence we have received suggests that the reverse is generally not true*".[11] However, the CL WP provides no reliable evidence to support these claims. Only anecdotal interviews are mentioned without further detail. The CL WP does not provide any empirical analysis showing that customers do not consider or use on-prem as an alternative to cloud services. For example, no diversion ratio from cloud services to on-prem services is estimated.

8. In this section, we show that on-prem is an effective alternative to cloud services. Competition often takes place at the workload level rather than on a customer level, and on-prem solutions may offer a more cost-effective and secure option for many workloads.[12] We present growing evidence that customers switch back from cloud to on-prem. In this context, the following two metrics highlight the substitutability of cloud with on-prem:

- **Diversion ratios from cloud to on-prem:** the proportion of customers that move their workloads to on-prem *out of all customers that move any workloads away from their cloud provider*.
- **"Cloud repatriation" rates:** the proportion of customers that move some or all workloads to on-prem out *of all cloud customers*.

9. In the following, we provide empirical evidence for both metrics.

2.1. Evidence for high diversion from cloud to on-prem

10. **High diversion ratios are particularly informative on the competitive pressure exerted by on-prem services.** A diversion ratio analysis, which looks at where customers or workloads go when they leave their current cloud service provider, is a relevant benchmark for at least two reasons. *First*, cloud repatriation rates in terms of switching back from cloud to on-prem (as discussed further below) can still understate the competitive significance of on-prem. For example, if few customers switch to on-prem, it may indicate low willingness to switch or high satisfaction with their current cloud provider. As such, it is important to also consider the diversion ratio, which indicates how many of those companies who switch providers choose on-prem. Therefore, it is more informative to consider the proportion of workloads moved to on-prem *when switching occurs*.

10 CMA's CL WP, para 4.73.

11 CMA's CL WP, para 4.72.

12 Forbes (2023), *The Rise Of Cloud Repatriation: Why Companies Are Bringing Data In-House.* (Link)

11. *Second*, there is no need for *all* customers to consider on-prem as an alternative for cloud services to exert significant competitive pressure. It only suffices that on-prem is an alternative for a significant proportion of workloads.

12. As an illustrative thought experiment, one can apply the hypothetical monopolist test (HMT) to argue that on-prem exerts significant constraint on cloud services. The HMT's objective is to determine if an alternative (e.g., on-premises services) imposes a competitive constraint on a product (e.g., cloud services) and should be considered part of the same market. If, following a small (typically 5-10%) price increase by a hypothetical monopolist in the cloud segment, a sufficient *portion* of customers moves some of their workloads to on-premises services, such that the price increase becomes unprofitable, then on-premises services should be deemed part of the relevant market, as sufficiently many workloads would be switched by customers to on-prem to constrain the cloud pricing. It suffices that on-prem is an alternative for *enough workloads* (but not necessarily all workloads) to make a hypothetical monopolist's price increase unprofitable.

Evidence shows that when customers switch away from a public cloud, they frequently choose on-prem

13. According to 451 Research, 65% of organisations that have switched workstreams away from a "hyperscaler"'s public cloud in the past 12 months switched to their own on-prem services (see below). This share represents the diversion ratio from "hyperscalers" to on-prem only and not the diversion between larger cloud providers. [13]

13 The on-prem diversion ratio is computed by considering the clients going from hyperscale public cloud to their own datacenters out of all those that moved out from public cloud.

Figure 1: Destination of workstreams switching away from "hyperscaler" public clouds in the past 12 months

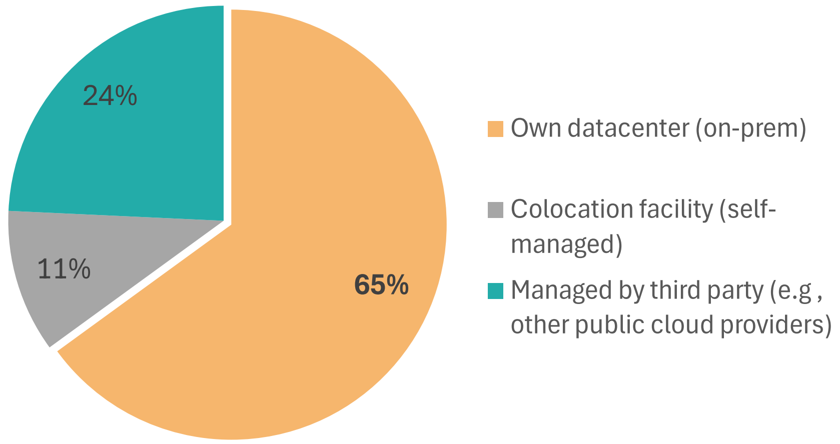

Source: CRA analysis of 451 Research data, Figure 1 (Link)

Note: Figures have been normalised to total 100%; source figures sum to greater than 100% as respondents have migrated multiple workstreams in the past 12 months.
Example: 87.3% of respondent moved workloads to self-managed infrastructure and 28% moved workloads to infrastructure managed by third parties. 87.3%+28% = 115.3%. Therefore, normalised self-managed = 87.3/115.3 = 75.7%.
It follows, own datacentre = 75.7% x 85.8% = 64.9%

14. The high diversion ratio demonstrates that a high proportion of customers who are using the cloud, and thus have already incurred any switching and startup costs to do so, are *choosing to switch to* on-prem solutions instead of running those workloads in the cloud. It is also worth noting that the pool of respondents considered by 451 research (643 datacentre/colocation users) mainly includes large enterprises. This has two important implications that make this result even more relevant. First, large firms are expected to behave rationally, thus survey outcomes are not likely affected by behavioural biases. Second, large firms account for the largest portion of cloud services demand. Hence, their business strategies are the main drivers of market outcomes. This is important because, even if not *all* customers view on-prem as an alternative to cloud services, a significant proportion of larger firms viewing it as one would be enough for on-prem to constrain cloud providers.

15. Additional evidence for high diversion ratios can further be found from a survey conducted by Public First (on behalf of CCIA) among 1,001 current or potential users of IaaS/PaaS/SaaS in the UK. According to the Public First survey,[14] of the UK organisations that have ever switched cloud infrastructure providers, 29% switched to on-prem services.

16. This survey demonstrates that businesses, incl. cloud customers, are using on-prem solutions or are willing to consider them in future despite the possibility of running the same workloads in the cloud. Thus, the competitive constraint of on-prem providers on cloud service providers can be at least as strong as of any individual cloud service provider incl.

14 Public First (2023), *Public First Poll for CCIA (Cloud Users)*, Question 31. (Link)

AWS/Azure/GCP – in other words, a customer is at least as likely to switch to on-prem as they are to any given cloud service provider.

Figure 2: Destination of customers switching cloud infrastructure providers

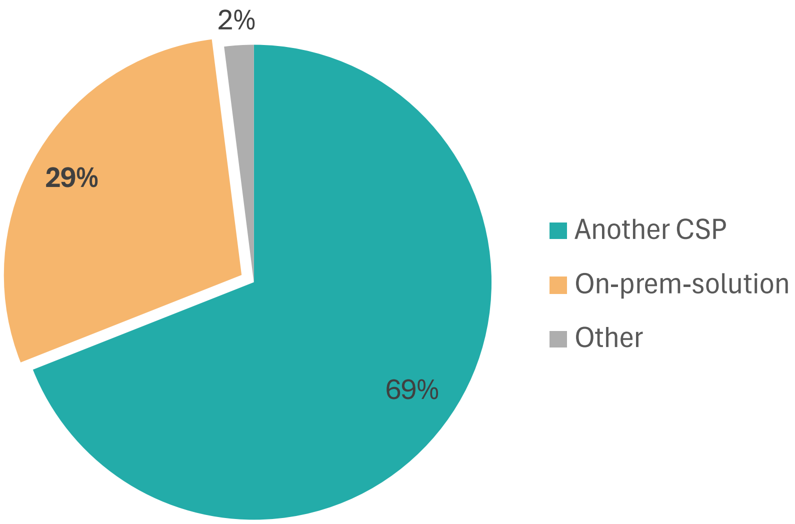

Source: CRA analysis of Public First survey data (Link), Question 31

17. Further evidence from this survey suggests that this pattern of high diversion is expected to even increase in the future. In fact, 47% to 50% of all survey respondents noted that they would consider using on-prem solutions for these services in the future even though alternative cloud options are broadly available.[15] Further, 46% to 53% of all survey respondents currently use on-prem servers for storage, database management, running applications, backup, and disaster recovery.[16]

18. The CL WP provides no compelling justification for placing no "evidential weight" on the Public First survey**.** The CL WP cites a lack of clarity on the type of sampling, methodology and the type of switching captured as reasons. It also highlights Amazon's membership in CCIA.[17] The resulting inference that numbers may have been doctored to favourably represent AWS is not an acceptable critique. Furthermore, results on the adoption of on-prem services from the Public First survey are in line with those of other surveys (see also next section), which suggests a certain reliability in terms of both methodology and outcomes. Lastly, CCIA has offered to follow up on "*any topics where [the CMA] would appreciate further evidence*", meaning the CMA has sufficient opportunity to address the concerns listed above.[18]

15 Ibid, Questions 3, 5, 7, and 9.

16 Ibid, Questions 2, 4, 6, and 8.

17 CMA's CL WP – Appendix A, paras A.39 – A.42.
CMA's CL WP, Fn. 501.

18 CCIA response to the CMA's CL WP (Link).

19. To see the potential of on-prem as a substitute to cloud in this case, it is further instructive to look at the size of the on-prem segment. Cloud currently represents less than 15% of IT spend,[19] whereas more than 70% of IT spend is represented by on-prem solutions, which warrants a thorough economic analysis of its competitive pressure on cloud service providers.[20]

20. The high diversion ratios provide strong evidence that customers see on-prem as an alternative to cloud services. When users switch from a cloud supplier, they often switch to on-prem. A 30-65% diversion ratio demonstrates that on-prem exerts significant competitive pressure on cloud providers. The CL WP does not provide empirical evidence to the contrary.

21. Overall, the significant competitive constraint imposed by investments and innovation undertaken by on-premises IT providers cannot be ignored. Many on-premises IT providers are continuing to invest in order to retain existing customers and attract customers away from cloud services providers.

- For instance, Penta Infra have acquired the Nexus data centre campus in Belgium, a clear sign of the expectation that they will be supporting on-premises workloads on a long-term basis, as they have invested heavily in the provision of a data centre for existing on-premises customers.[21]
- In addition, HPE and Danfoss have announced a collaboration to deliver off the shelf heat recovery modules that allow data centres to be cooled more efficiently while recovering and reusing excess heat. This demonstrates that on-premises solutions providers are continuing to develop new products that make on-premises IT solutions more attractive and cost efficient for customers.[22]

22. More generally, infrastructure incumbents such as Dell, VMware and HPE are "pivoting toward cloud-like 'as a service' operating models", demonstrating strong supply side substitutability with cloud providers.[23]

23. Another indication for the continuing and future importance of on-prem solutions as a competitive constraint on cloud services providers is that they constitute relevant compute capacity options for AI companies. On-prem solutions are considered alongside co-located environment, solutions provided by cloud services providers, and hybrid solutions combining these options. Examples of growing on-prem compute capacity solutions related to AI are set out below:

---

19 Gartner for total IT spending (Link) and Gartner for total cloud spending (Link).

20 In order to estimate the value of on-prem services we have subtracted from total IT spend ($5,061bn, link) first the value of those IT categories that are unlikely, at the moment of writing, to contain alternatives to cloud services (i.e. "devices", $688bn). We then removed also the value of cloud spend ($679bn, link). The resulting value ($3,694bn) has been divided by either the total IT spend ($5,061bn) or the total non-cloud spend (($5,061bn-$679bn=$4,382bn) to get the shares reported in the text.

21 Data Center Dynamics (2024), *Penta Infra acquires data center in Brussels, Belgium.* (Link)

22 Hewlett Packard Enterprise (2024), *Hewlett Packard Enterprise and Danfoss partner to curb data center energy consumption and reuse excess heat.* (Link)

23 451 Research (2021), *Cloud repatriation: What it is, what it isn't, and why it's not going away,* (Link)

- In November 2023, Dell announced a collaboration with Hugging Face to offer on-premises deployment of customised large language models ("LLMs") on Dell infrastructure products and services.[24]
- Dell announced a collaboration with Meta in October 2023 to make it easier for Dell customers to deploy Meta's Llama 2 on-premises with Dell's generative AI portfolio of IT infrastructure.[25]
- In March 2024, HP and NVIDIA announced that they are developing new laptops that can build, test, and run LLMs locally (as opposed to on the cloud).[26]

2.2. Evidence of cloud repatriation

24. In addition to looking at diversion ratios to on-prem, which focus only on customers who switch, it is also informative to consider rates of cloud repatriation - *i.e.*, the percentage of all customers who move any workloads back to on-prem. Many companies that moved to the cloud are moving workloads back to on-prem. Survey evidence shows that the share of cloud customers that have moved workloads back to on-prem or plan to do so is significant, between 24% and 36%, confirming that traditional services are in fact a real alternative to cloud ones. See also Figure 3 below. In a similar vein, Barclays' survey of CIOs found that 83% of enterprise CIOs that responded to the survey plan to repatriate at least some IT workloads from the cloud to private cloud or on-premises in 2024.[27]

25. These results underline that switching to on-prem is a genuine option, is prevalent and puts competitive constraints on cloud service providers.

---

24 Dell (2023), *Dell Technologies and Hugging Face to Simplify Generative AI with On-Premises IT.* (Link)

25 Dell (2023), *Dell and Meta Collaborate to Drive Generative AI Innovation.* (Link)

26 NVIDIA (2024), *LLMs Land on Laptops: NVIDIA, HP CEOs Celebrate AI PCs.* (Link)

27 The Human Factor (2024), *Why CIOs are Reverting to On-Prem and What it Means for Enterprise AI.* (Link)

Figure 3: Share of firms that have migrated or will migrate workloads from cloud to on-prem according to different surveys

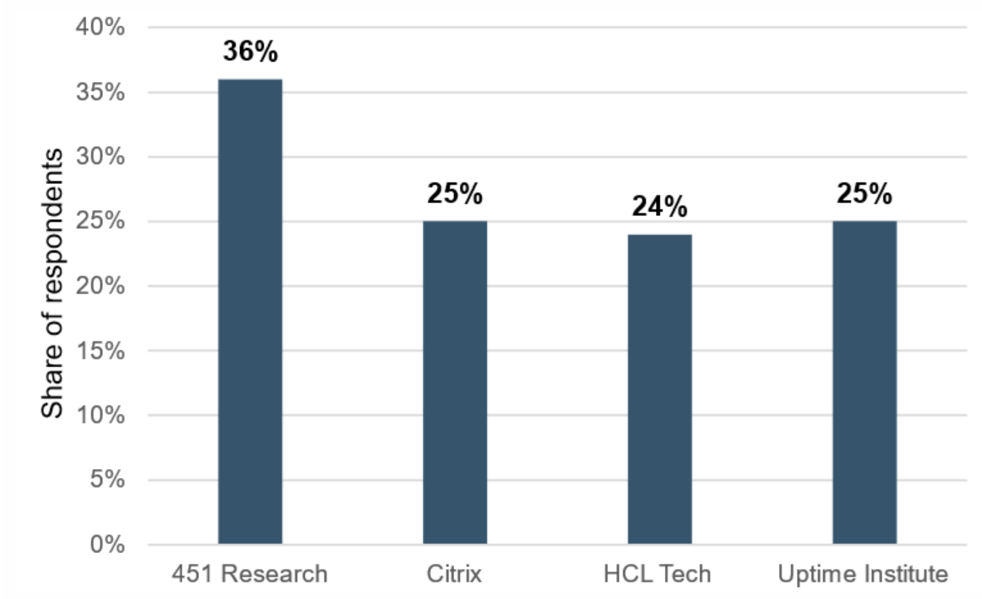

Source: 451 Research (Link), Citrix (Link), Uptime Institute (Link) and HCLTech (Link)
Note: In its survey, HCLTech queried company executives about their intentions to migrate at least a portion of their workload back to on-prem systems. Similarly, 451 Research posed this question to datacentres/colocation and IaaS/PaaS users. Citrix was cited by a publisher (Infoworld), which stated the percentage of organizations that have moved half or more of their cloud-based workloads in the UK. Lastly, the Uptime Institute surveyed organizations, asking whether they had moved any production applications from a public cloud to their own datacentres.

26. The significant competitive constraint imposed by on-premises IT solutions is also illustrated by specific evidence of existing customers of cloud services providers having either switched away from or are considering a switch from the cloud to on-premises or adopting a hybrid cloud architecture.[28] Such examples of cloud repatriation, or more generally moves from cloud to on-prem, include:

- Dropbox transitioning from cloud services back to its own infrastructure in 2015, leading to cost reductions. Starting in 2015 it began moving users away from AWS's S3 onto its own custom-designed infrastructure and software. It is reported that Dropbox saved ~ $75 million over the two years following the transition by building its own tech infrastructure.[29]
- Ahrefs, one of the leaders in SEO tools, decided that on-prem was a better option because it was more compatible with their business, so more profitable. Ahrefs claims to have saved ~ $400 million by keeping a significant share of their workloads on-prem (having compared that option against the costs of using AWS).[30]

28 [✂].

29 GeekWire (2018), *Dropbox saved almost $75 million over two years by building its own tech infrastructure*. (Link)

30 Mediium (2023), *How Ahrefs Saved US$400M in 3 Years by NOT Going to the Cloud*. (Link)

- Plitch, a software developer in gaming, moved their workloads from the cloud to on-premises because of their needs for security over their Research and Development ("**R&D**") data and code, and because Plitch considered that it could obtain more processing power for their AI needs on-prem than on the cloud.[31]
- 37Signals moved their workloads from the cloud to on-prem as they felt that on-premises solutions were more cost-effective and reliable.[32]
- [✂].[33]

3. THERE IS NO EVIDENCE OF CONSOLIDATION; INSTEAD, EVIDENCE AT HAND SHOWS THAT CONCENTRATION IN THE ARTIFICIALLY NARROW CLOUD SEGMENT IS DECREASING

27. The CL WP aims to estimate the shares of cloud services supply, based on three measures: (i) revenues, (ii) datacentre capacity, and (iii) flow of new business.[34] Based on its estimates of revenue shares of supply, the CL WP states that the "*overall cloud services sector is concentrated, and concentration is increasing over time*", citing AWS and Azure as the main players.[35] Based on its estimates of datacentre capacity shares of supply, the CL WP concludes that AWS and Azure are the main datacentre capacity providers, which the CL WP uses to strengthen its previous conclusion.[36] Based on its estimates of shares of supply based on flow of new business, the CL WP concludes that Azure is winning customers in the UK at a significantly higher rate than other cloud providers, whereas AWS's share of new customers has been around [✂]% by count and [✂]% by revenue.[37]

28. However, the CL WP's emerging views on the level of concentration suffer from the fundamental flaw of excluding on-prem solutions, which, as shown above and contrary to statements in the CL WP, exert significant competitive pressure on cloud services. Moreover, the CL WP focuses only on cloud infrastructure, excluding even the so called "SaaS" cloud segment. Therefore, "cloud infrastructure" is a narrower segment than the one that should be considered to meaningfully assess concentration. Nevertheless, even when restricting the concentration analysis to the cloud infrastructure segment, we observe a *decreasing* concentration index over time.

29. The data and calculations related to (i) – (iii) above have not been disclosed to AWS, so it is not possible to exactly replicate the CL WP's analysis or comment on its accuracy. However, to respond to the issues raised in the CL WP, we have assessed the data available at hand. The results are as follows:

---

31 BBC (2024), *Are rainy days ahead for cloud computing?* (Link)

32 Ibid.

33 [✂].

34 CMA's CL WP, para. 5.4.

35 CMA's CL WP, para. 5.21 and 5.22.

36 CMA's CL WP, paras. 5.33-5.35.

37 CMA's CL WP, paras 5.50-5.53.

- **Omission of on-prem artificially narrows the market, biasing any concentration analysis.** To properly conduct a concentration analysis, and assess whether cloud services are competitively provided, one cannot ignore the competitive pressure exerted by on-prem solutions on cloud services. As shown in the previous section, on-prem services are a serious alternative to cloud for many workloads. The CL WP artificially narrows down the relevant providers by ignoring on-prem and restricting itself to measuring cloud concentration alone. Subsequent results should be considered based on this important omission.
- **Our results based on [✂] show a decreasing cloud concentration index.** [✂].
- **The CL WP** [✂].[38]
- **The CL WP wrongly assesses the flow of new business based on the number of customers, rather than workloads.** The notion of new versus existing customers is not the relevant metric for cloud providers and is not reflective of how competition works, as competition primarily takes place for incremental workloads rather than new customers. Existing customers therefore seek the best place to run their workloads, such that IT providers compete fiercely on a workload-by-workload basis. Therefore, the CL WP's analysis likely underestimates the share of new workloads for which AWS competed / won.

3.1. Contrary to CL WP's claims, the concentration index in the artificially narrow cloud segment is decreasing

30. The CL WP combines actual data on revenues from AWS and other cloud service providers with third-party datasets by IDC and Synergy to estimate shares of cloud services supply. Based on these estimates, the CL WP concludes that, for IaaS and IaaS/PaaS: (i) the market is concentrated and (ii) concentration is increasing over time. Moreover, the CL WP states that AWS's share has remained broadly stable in 2019-2022.

31. However, based on the most recent [✂] data we find that concentration indices (HHIs)[39] decreased in all segments considered by the CL WP. [✂]. This is shown in Figure 4 below.

32. The most recent [✂] data also shows that AWS's shares decreased in each of the three segments considered by the CL WP (i.e., IaaS, PaaS, and IaaS+PaaS). [✂]. This is shown in Figure 5 below.

33. [✂] thus contradicts the assertion that concentration as measured by HHI is increasing in the supply of cloud infrastructure in the UK. We further note that the CL WP's analysis omits 2023, which corresponds to the most significant decrease in concentration and in AWS's shares.[40] [41]

38 CMA's CL WP, para. 5.29.

39 [✂].

40 [✂].

41 We note that there are differences in our data vs the CL WP's data sources such that if CL WP were to include 2023 in its analysis, it might lead to somewhat different figures.

Figure 4:

[✂].

Figure 5:

[✂].

3.2. AWS has a small datacentre capacity share, both globally and in the UK

34. The CL WP calculates shares of supply by datacentre capacity using data from AWS, Microsoft, Google, IBM, and Oracle (in MW) within the UK, UK + EEA and globally.[42] The CL WP concludes that AWS and Microsoft are the largest providers.[43] The CL WP notes that the current shares are overestimated and that, subject to availability, it plans to gather data from other "IaaS" providers that serve UK customers.[44]

35. Given that the CL WP recognised that its shares are being overestimated due to the inclusion of only a few suppliers, we have estimated more accurate supply shares using (i) AWS's actual global and UK data centre capacity in MWs (for the numerator) and (ii) third-party data on global and UK total datacentre power capacities in MWs (for the denominator).[45] The resulting volume-based supply shares for compute services are shown in the Table 1 below. According to all estimates, AWS's share is below [✂]% both globally and in the UK.

Table 1: [✂].

4. CUSTOMERS FREQUENTLY MULTI-CLOUD AND SWITCH

4.1. The CL WP's multi-clouding analysis is flawed and does not support its assertion that multi-cloud is uncommon

4.1.1. Description and flaws of the CL WP's multi-clouding analysis

Description of the CL WP's multi-clouding analysis

36. The CMA's CL WP criticises and dismisses multi-clouding evidence from both (i) the quantitative surveys that measure multi-cloud[46]; and (ii) AWS's win/loss data analysis. The CL WP also appears to dismiss submissions from other cloud service providers.[47] It then builds its own estimate of multi-cloud prevalence based on customer datasets from AWS, Microsoft, and Google. The CL WP estimates the unweighted (by count) and weighted (by spend) share of customers that use at least two clouds among these three providers only.

42 CMA's CL WP, para 5.29.

43 CMA's CL WP, paras 5.33, 5.34, and 5.35.

44 CMA's CL WP, para 5.29.

45 [✂].

46 CMA's CL WP, para 3.58-3.62.

47 CMA's CL WP, para 3.48-3.57.

The CL WP's results should be interpreted with caution as they are subject to several limitations as acknowledged by the CL WP itself.[48]

37. From its analysis, the CL WP claims that:

a. Multi-cloud is uncommon: "[c]*ustomers using multiple clouds is uncommon, but […] large customers are more likely to multi-cloud*".[49]

a. Integrated multi-cloud is rare: *"the evidence indicates that an integrated multi-cloud architecture is rare"*.[50]

38. In the below, we show that:

- The CL WP's analysis suffers from several conceptual flaws, including the absence of any benchmark against which to assess when multi-cloud is low, and a narrow focus on only 3 cloud service providers that likely underestimates the multi-cloud rates severely.
- Considering the relevant spend-weighted multi-cloud rate reveals that even when taking CL WP's analysis at face value, i.e. without correcting for errors and likely biases, there are material multi-cloud rates of around [✂]%. There is no economic basis for such rates being lower than what one would expect in a well-functioning market.
- Correcting for errors and accounting for biases in the CL WP's analysis reveals that multi-cloud rates could be as high as [✂]% among UK customers. This result is consistent with several third-party studies that confirm a wide adoption of multi-clouding by enterprises and shows the likely magnitude of the bias due to the CL WP only including 3 cloud service providers in its analysis.
- We discuss in more detail weaknesses of the CL WP's analyses, such as arbitrary spend thresholds used to define customers who multi-cloud.

Flaws of the CL WP's multi-clouding analysis

39. The CL WP's multi-clouding analysis is conceptually flawed for the following reasons:

- There is **no benchmark** against which to gauge whether observed multi-cloud levels are "high" or "low". Not all customers wish to multi-cloud, as it would not be efficient for some of them to do so.
- CL WP's analysis considers **only three cloud service providers (AWS, Azure, and Google)**, ignoring all other cloud service providers and all on-prem providers. This biases their multi-cloud estimates downward. This is because it ignores all multi-clouding between other cloud service providers with AWS, Azure, and Google, and ignores all multi-clouding between two cloud service providers other than these three. The prevalence of multi-cloud is therefore necessarily underestimated when only a subset of the competitors is considered.

---

48 Note that the CMA's CL WP acknowledges that their analysis is subject to several limitations. (CMA's CL WP, para 3.71).

49 CMA's CL WP, para 3.75.

50 CMA's CL WP, paras 3.79-3.80.

- CL WP's analysis is **distorted by the overreliance on very small customers.** The CL WP gives the same weight to very small customers and larger customers in terms of cloud spend, even though smaller customers are: (i) less likely to benefit from multi-clouding, and (ii) represent a smaller share of spend.
- To assess effects on competition, it is relevant to consider **the share of *spend*** that is subject to multi-cloud, not the share of customers. This is because competitors are foreclosed only if they do not have access to workloads, and therefore cloud spend. They cannot be foreclosed because small customers with few workloads and less than $10k annual spend (that represent [✂]% of customers but [✂]% of total cloud spend) do not multi-cloud.

40. We also note that the CMA is inconsistent in the way it considers that the share of spend – as opposed to the share of customers – is relevant for the competitive assessment. For example, in its working papers on discounts and licensing practices, the CMA considers that it is more appropriate to consider the share of spend and the impact on the larger customers *"whose choice of cloud provider has a greater impact on competition"*[51] (which we agree with). However, in its multi-cloud analysis, the CMA places a lot of weight on the share of customers instead (which is less relevant than share of spend).

41. Below we discuss in more detail various flaws in the CL WP's multi-clouding analysis. As we show further below, correcting for associated errors demonstrates that multi-cloud rates are high.

4.1.2. The CL WP's own analysis shows material rates of multi-clouding when weighted by spend

42. The CL WP's own analysis shows substantial rates of multi-clouding when weighted by spend. The CL WP's analysis (fixed for a few coding errors) does not indicate that multi-cloud is "low": [✂]% of customers multi-cloud (weighted by spend). This is also demonstrated in the below table, which shows our replication of the CL WP's prevalence of multi-clouding analysis. In particular:

- A multi-cloud level of [✂]% is arguably not "low", particularly given this estimate is known to be biased downward because the CL WP considers only three cloud service providers, ignoring other cloud service providers and on-prem providers.
- While there is no benchmark to tell whether the optimal level of multi-clouding is higher or lower than [✂]%, the corrected CL WP's results already demonstrate the high prevalence of multi-clouding. These figures, and the even higher multi-cloud rates when accounting for additional shortcomings in CL WP's analysis, show that there is no indication of any barriers that prevent customers from multi-clouding. Put simply, customers can multi-cloud if they want to.

43. As customers <u>can</u> multi-cloud, they will do so when it is beneficial for them, which is not 100% of the time. If customers can multi-cloud, they can decide for themselves when to do it. As for the degree of multi-cloud, it will be up to customers' choice and preferences, trading off the benefits and costs of a multi-cloud strategy.

51 [✂].

Table 2: [✂]

4.1.3. CL WP ignores multi-clouding with smaller cloud service providers

44. The CL WP's failure to collect and use data from smaller cloud service providers biases its multi-clouding estimates. The CL WP's analysis is based on the revenue datasets of only three cloud service providers (Microsoft, AWS, and Google), such that any customer that multi-clouds with other smaller cloud service providers would not be identified as a multi-clouding customer. This approach is inconsistent with the CL WP's apparent concern about smaller providers not being able to compete with AWS, Azure, and Google. The data collected by the CMA does not allow us to observe the very type of multi-cloud that the CMA is interested in assessing.

45. The CL WP itself acknowledges that the focus on Microsoft, AWS, and Google is a limitation of its analysis.[52] The scale of the resulting bias is likely significant considering that, by omitting other cloud service providers, the [✂] in the UK for 2020-2022 (which corresponds to the cloud spend definition of the CL WP's dataset)[53].

46. To estimate the extent of the bias, we implement an adjustment to account for the presence of other cloud service providers. This adjustment is aimed at showing the extent of the possible bias in the CL WP analysis. It is done by approximation, relying on cloud shares of other providers[54]. As a result, we find that the multi-cloud rate [✂] when this adjustment is implemented. This demonstrates the extent of the likely bias in the CL WP's analysis restricted to three cloud service providers only.

47. We are not arguing that our adjustment can replace collecting actual data for other cloud service providers. Our adjustment is only an approximation since actual behaviour is not observed. Our main conclusion holds however: the bias due to the exclusion of most cloud service providers from the analysis is most likely substantial.

4.1.4. Multi-clouding is low only for the smallest of customers

48. As noted above, the prevalence of multi-cloud should be assessed based on the share of *spend*, not share of *customers*. This is because competitors are foreclosed if they do not have access to workloads/cloud spend. They are not foreclosed because small customers with less than $10k spend (that represent [✂]% of total cloud spend but [✂]% of all customers) do not multi-cloud.

49. Moreover, smaller customers are unlikely to multi-cloud for good economic reasons: (i) they don't have the required IT resources to multi-cloud, (ii) nor do they derive material benefits from multi-clouding. Hence, competition for smaller customers is about switching their entire workloads from another cloud service provider or from on-prem; it is not about getting smaller customers to multi-cloud. Focus should be on multi-cloud shares weighted by customer spend.

50. The CL WP also fails to acknowledge that it is questionable whether customers have a specific desire to multi-cloud. The counterfactual the CL WP appears to have in mind is

52 CMA's CL WP, para 3.71(d).

53 [✂].

54 See the Annex for further details.

that, absent the alleged barriers to switching and multi-cloud, a majority, if not all, customers would constantly switch and multi-cloud. There is no evidence that a majority of customers would necessarily be interested in having multiple cloud services providers. While large customers and customers in regulated industries may want to multi-cloud, it is questionable whether mid- and small-sized enterprises, such as OGVA customers, would be better served by multi-clouding.

51. Many companies are happy using one operating system, one accounting software, one internet security solution and one set of productivity tools, and this is not because the software providers impose barriers to switching, but simply because the cost that the companies incur from having multiple solutions are larger than the benefits of achieving a lower price for some specific use case that can be outsourced to another provider. Against this background, it must be understood that the CL WP's multi-cloud analysis (as presented in Figure 3.2 of the CL WP) is fully consistent with an equilibrium where those (predominantly large) customers that do want to multi-cloud do multi-cloud and those (predominantly small) customers that do not want to multi-cloud do not multi-cloud, without any alleged barriers imposed by cloud service providers.

52. As shown in Table 3 below, if we remove customers with just below $10k/year spend, the unweighted share of multi-clouding already [✂]. If we further consider customers with total cloud spend above $30k/year or $50k/year, the share further increases to [✂]% and [✂]%, respectively.

Table 3: [✂]

53. The CL WP's choice of spend buckets in Figure 3.2 of the CL WP is also misleading. The CL WP conveniently groups customers from a wide range of $10k to $1m spend bands into a single bucket demonstrating a low multi-clouding level of [✂]% for this spend bucket. However, as noted above, it is not clear why such small customers of only $10-50k yearly spend *should* multi-cloud. As shown in Figure 6 and Figure 7 below, once we re-adjust buckets to increase the upper threshold of the lowest-spend bucket from $10k to even just $50k per year, the share of multi-clouding customers significantly increases to [✂]% in the $50k-1M bucket too (while staying mostly above [✂]% in subsequent buckets).

54. CL WP's claims of low sample sizes for higher-bucket-spend customers[55] are also unjustified. The CMA seems to have purposefully introduced this to frame it as if results for large customers are unreliable. The results are equally reliable for large and small customers, it is just the nature of the data that there are *fewer* customers with *very* large spend. Irrespective of that, when looking at customers with at least $10-50k yearly cloud spend, we still have sufficiently large sample sizes.

55 CMA's CL WP, para 3.85.

Figure 6: [✂]

Figure 7: [✂]

4.1.5. Implementing adjustments to correct the flaws in CL WP's analysis yields high multi-clouding levels

55. Based on the flaws discussed above, we have implemented the following adjustments to the analysis in the CL WP:

- **Adjustment 1: we recognise the relevance of weighted shares.** Unweighted shares distort the results due to the long tail of small customers that are not expected to multi-cloud (as they are less likely to benefit from multi-clouding) and that are less relevant for the competitive assessment. The share of spend (rather than share of customers) is a more relevant metric to assess the effects on competition because, to be foreclosed, other cloud service providers would need to lack access to a significant share of workloads/cloud spend. They will not be foreclosed due to small customers with less than $10k annual spend (that represent [✂]% of customers but [✂]% of total cloud spend) not multi-clouding.
- **Adjustment 2: we fix coding errors.** There are some coding errors, such as customers being counted twice in the denominator and incorrectly aggregating multiple clients as one, leading to an overall downward bias in the multi-cloud spend share[56]. We provide further details on this in the Annex.
- **Adjustment 3: we account for providers other than AWS, Microsoft, and Google.** While there might be different ways of accounting for other providers, this analysis shows the extent of the possible bias from ignoring most cloud service providers from the analysis.

56. The impact of each adjustment on average multi-cloud share for the period 2020-22 is summarised in Figure 8. Once all adjustments are made, we obtain a multi-cloud share of [✂]% (weighted by spend) and [✂]% (unweighted by customer share), as opposed to the unweighted [✂]% in the CL WP analysis. We note that [✂]% multi-clouding rates are consistent with several third-party studies that support multi-clouding being widely adopted by enterprises, as shown further in Section 4.1.7.

56 We also note that CL WP's fuzzy matching procedure results in imprecise and potentially unreliable results due to the existence of two forces with opposing effects, the net effect of which the CL WP fails to assess despite acknowledging the existence of those forces: existence of 'false negatives', i.e. customers that are erroneously labelled as non-multi-clouders because of a similarity score from the matching below 99%, and 'false positives', i.e. customers that are erroneously labelled as multi-clouders due to falsely being matched with another customer of a different cloud service provider while still achieving a similarity score above 99%. As an example for false negatives, we have noticed that a company "xyz limited" is sometimes labelled under the same name in one dataset, but as "xyz ltd" in another, which is not identified as a match by the CMA's fuzzy match. Further, we find that some customer names in the Microsoft dataset are 'masked', i.e., the names that appear in the dataset are encrypted, possibly for anonymisation. This results in a failure to account for these customers as possible multi-clouders, and therefore likely leads to additional false negatives which cannot be accounted for in the analysis.

Figure 8: [✂]

4.1.6. CL WP's 70/30 multi-cloud split threshold is not economically sound

57. The CL WP further argues that "*Customer architectures with a more even [spend] split may be more likely to have integrations across their clouds*"[57] and thus calculates the proportion of customers that have at least 30% of their spend on their secondary cloud. However, there is no economic or empirical basis for this *arbitrary* threshold. The CL WP hypothesises that only a split of 70/30 or more even would be observed in the presence of integrated multi-cloud.

58. Fundamentally, there is no reason why there should be a relationship between integrated multi-clouding and a more even split of spend. Customers can use different cloud service providers for different purposes in a variety of ways. Both even and uneven splits can be expected with both integrated and non-integrated multi-cloud architectures. For example, a customer might use AWS for its administrative services and Google for its analytics department. Without any loss of generality, we can assume that the analytics department would be more data intensive than the administrative. Despite the two departments being heavily integrated, it would be reasonable to find that a customer has a much lower share with AWS than with Google. The CL WP ignores such a possibility and proposes an arbitrary threshold with no economic evidence, thereby again severely underestimating the multi-cloud share.

4.1.7. Third-party evidence supports a high prevalence of multi-clouding

59. Several third-party studies support our finding that multi-clouding is widely adopted by enterprises. We previously submitted this evidence to the CMA which was disregarded in its CL WP.[58] This was the case for the Flexera report (showing 87% of firms multi-cloud)[59], the Oracle survey (98%),[60] and the Public First survey (71%)[61].

60. Complementary to this, we provide additional evidence from recent third-party studies aligning with the results of our adjustments to the CL WP analysis and substantiating the body of evidence of a high prevalence of multi-cloud amongst UK customers. In particular:

- *First*, Cockroach Labs' review of the state of multi-cloud finds that nearly half of UK companies employ a multi- cloud architecture in 2024.[62] Including dual use with on-prem services, 77% of UK companies use several providers. The results are presented in Figure 9 below.

57 CMA's CL WP, para 3.79.

58 CMA's CL WP, Appendix A.

59 Flexera (2024), *Stata of the Cloud Report*. (Link)

60 Oracle (2023), *Multicloud in the Mainstream, February 2024*, Figure 1. (Link)

61 Public First (2023), *Public First Poll for CCIA (Cloud Users)*. (Link)

62 Cockroach Labs (2024), *State of Multi-Cloud 2024,* p.7. (Link)

Figure 9: UK firms by cloud architecture, 2024

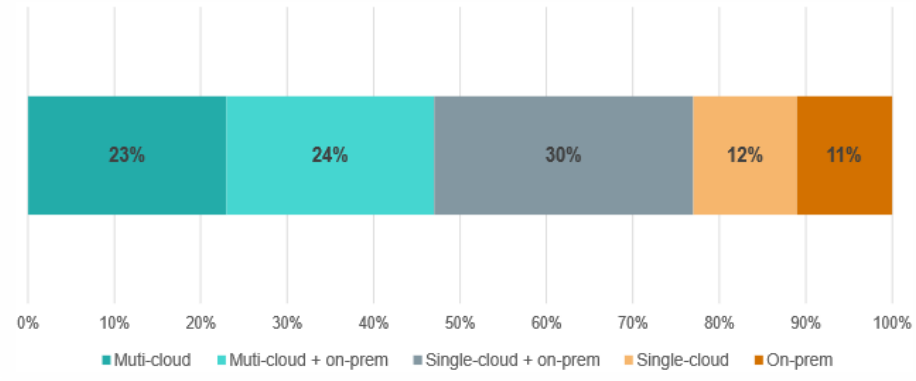

Source: CRA analysis of Cockroach Labs, *"State of Multi-Cloud 2024",* p.7. (Link)

- *Second*, OVHcloud studied over 500 IT decision makers of large organizations in the UK[63], finding: (a) 62% of organisations are currently using a multi-cloud environment; (b) a further 18% of organizations are actively in the process of transitioning to a multi-cloud environment; (c) 64% of firms see their use of multi-cloud increasing in the next two years, and (d) only 1% have no plans to use multi-cloud at all. The results are presented in Figure 10 below.

Figure 10: UK firms state of multi-cloud, 2024

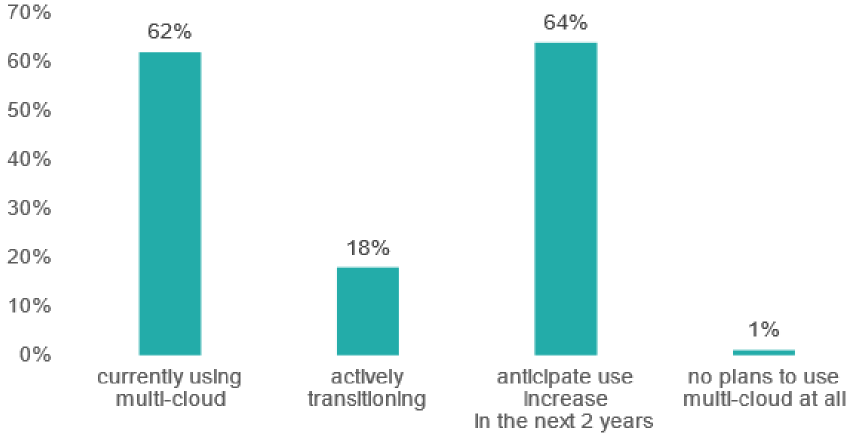

Source: CRA analysis of OVHCloud, "*The State of Multi-Cloud: 64% of organisations see their use of multi-cloud increasing in the next two years*". (Link)

61. One fundamental shortcoming of the CL WP's multi-cloud analysis is its consideration of only three cloud service providers. Upon review, the CMA should have instead opted to conduct its own *quantitative* survey that robustly examines customers' multi-cloud usage among *all* cloud service providers, potentially including on-prem. Given the absence of such

63 OVHCloud (2024), *The State of Multi-Cloud: 64% of organisations see their use of multi-cloud increasing in the next two years*. (Link)

a quantitative survey by the CMA[64], it should give much greater weight to quantitative surveys conducted by third parties (as presented above) than to its own multi-cloud analysis.

4.2. Customers also show strong ability to move workloads

62. **CL WP's preliminary assessment on customer switching.** To understand the prevalence of switching and potential restrictions to it, the CMA gathered information on switching from market participants, including CRA's analysis of AWS's data. The CL WP analyses qualitative and quantitative evidence, but states that evidence to assess customer switching is limited. It nevertheless concludes that switching overall is limited. The CL WP deems the quantitative analysis of CRA showing how switching is prevalent based on AWS's data to be of limited use. It notes limitations of the CRA analysis and states that the findings of CRA's quantitative analysis *"are not inconsistent with finding low switching levels in the cloud infrastructure services market*."[65]

63. In the following, we discuss how the CL WP's preliminary conclusions on switching are (i) not supported by solid evidence and (ii) are contradicted by our results on market outcomes.

4.2.1. The CL WP's analysis is inconclusive

64. The CL WP's analysis is conceptually flawed as it ignores several economic factors. In particular:

- Switching rates will be lower if customers are satisfied with their existing providers—e.g., because they are benefiting from low prices and/or a high-quality service. Therefore, low switching rates can be indicative of a highly competitive environment. For example, according to Digital Ocean, among customers who do not consider switching, the vast majority (80%) are not switching cloud service providers due to satisfaction with current providers[66]. Observing low levels of switching is not inconsistent with a competitive segment, nor does it imply that customers are unable to switch. Indeed, low levels of switching are consistent with high customer satisfaction.
- Although there are some inherent switching costs associated with IT services, these can be offset by other cloud service provider, and IT providers who want to attract new customers.
- Given that customers likely expect to stay with a new provider for several years when switching, the switching costs in proportion to general cloud spend over several years are even lower.

65. Looking at switching levels alone is not meaningful if not compared against some benchmark for what the level of switching should be in presence of effective competition. The CL WP does not provide a clear benchmark against which to measure whether

64 As we explain below, the JIGSAW survey the CMA relied upon was of qualitative nature, using a small sample of customers for each of the 5 cloud service providers it assessed. As such, this survey cannot used to obtain credible insights on customers' multi-cloud rates for either these cloud service providers or for the entire segment.

65 CMA's CL WP, para. 3.103.

66 Digital Ocean (2022), *Currents: Digital Ocean's report on how startups and small-to-medium businesses are faring today,* p.12. (Link)

empirically observed switching rates are high or low. The CL WP's claims that AWS's switching rates are "low" is therefore unsubstantiated because there is no benchmark against which to compare observed switching levels. This is a conceptual shortcoming.

66. In any case, given that (as noted above), there are inherent switching costs associated with IT services, very high switching rates might not constitute an efficient outcome either, since either customers or cloud service providers (or both) would have to bear these additional costs. The relevant question is whether customers want to switch but are prevented from doing so by barriers to switching imposed by providers. To this relevant question, the CL WP provides no evidence. We provided evidence that customers can and do switch when they want to, which the CL WP dismissed, as discussed below. AWS's customers will switch less if they are satisfied with AWS's services.

67. **The CL WP dismissed our quantitative analysis based on AWS's data without valid justifications.** The CL WP's argument that there are some limitations in our quantitative analysis is not compelling. In particular, the CL WP argues that defining customers as churned if the customer reduced its spending by at least 80% for three consecutive months and did not return to 80% of their original spend within six months may only imperfectly capture churn, as spending may increase again after three months or, alternatively, the customer can continue spending.[67] The CL WP also notes that normal seasonal fluctuations can be the drivers of spend decrease rather than customers moving workloads outside AWS's cloud. Moreover, it states that the dataset includes many small customers who only spend a minimal amount on AWS such that removing these small customers could further reduce the churn rate.[68]

68. However, all datasets and quantitative analyses will inherently contain "limitations", as we discuss in the next subsection, which is not an adequate justification for the CL WP excluding relevant information.

69. CL WP's conclusions are also heavily reliant on relatively sparse qualitative evidence from customers and providers. This approach is unlikely to produce robust results:

- First, the JIGSAW survey analysis used by the CL WP relies on a **small number of respondents** per cloud services provider. It excludes all cloud services providers or on-premises IT providers, except for AWS, Azure, Google, IBM, and Oracle.
- Second, the survey analysis relies on unstructured qualitative interviews producing **anecdotal responses**.
- Third, while the CMA is currently considering whether it could measure switching by analysing customer data of AWS, Azure, and Google, we note that, similarly to the CL WP's multi-clouding analysis, this analysis would **ignore all other cloud services providers or on-premises IT providers**. It would also suffer from the same methodological errors and flaws as the ones identified in the CL WP's multi-clouding analysis (as discussed above). The CMA must therefore aim to gather customer revenue data from smaller cloud service providers as well as taking into account various flaws identified in its multi-clouding analysis.

67 CMA's CL WP, para. 3.102.

68 Ibid.

70. We next show how empirical evidence clearly points to customers being able to freely switch.

4.2.2. [✂] shows that customers can and do switch when they want to

71. Anti-competitive concerns are only legitimate when clients cannot switch when they would like to. However, available evidence suggests that switching occurs and is not inhibited. [✂]. Specifically, we show that: (i) [✂]; and (ii) additional third-party evidence shows a high prevalence of switching amongst cloud customers.

[✂] customers can and do switch

72. As explained above, the notion of new versus existing customers is not the relevant metric for cloud providers and is not reflective of how competition works in cloud. In IT services, competition primarily takes place for incremental workloads rather than new customers. IT providers compete fiercely on a workload-by-workload basis. Therefore, the important question is whether customers can move *workloads* between cloud service providers.

73. [✂].[69]

74. Moreover, because larger customers are even less likely to switch entirely, competition should also be assessed by considering the evidence of multi-clouding, which is more prevalent for larger customers. In that respect, the evidence presented in the previous section alongside our previous analysis demonstrated that multi cloud is widespread for larger customers. This also constitutes evidence of effective competition taking place across cloud providers.

Table 4: [✂]

75. In the round, the combined evidence on switching and multi-clouding is indicative of a well-functioning market.

4.2.3. Third-party evidence also shows that customers can and do switch when they wish to do so

76. Third-party survey evidence also provides additional support for switching being generally common and prevalent among companies that wish to move their workloads.

- According to a 2022 Digital Ocean study, half of the surveyed businesses have switched cloud providers in the past. Of those who have not migrated, 77% expressed satisfaction with their current cloud provider and did not feel the need to switch.[70]
- In a 451 Research survey conducted in 2021, 48% of respondents indicated that they had transitioned a workload or application away from some "hyperscaler"'s public cloud to some other venue in the past 12 months.[71]

69 [✂].

70 Digital Ocean (2022), *Currents: Digital Ocean's report on how startups and small-to-medium businesses are faring today,* p.12. (Link)

71 451 Research (2021), *Cloud repatriation: What it is, what it isn't, and why it's not going away,* Figure 1. (Link)

77. These results are also summarised in Figure 11. Overall, our results suggest that there are no credible concerns for insurmountable barriers to switching, and workload switching is indeed highly prevalent.

Figure 11: Organisations that switched from clouds according to surveys (%)

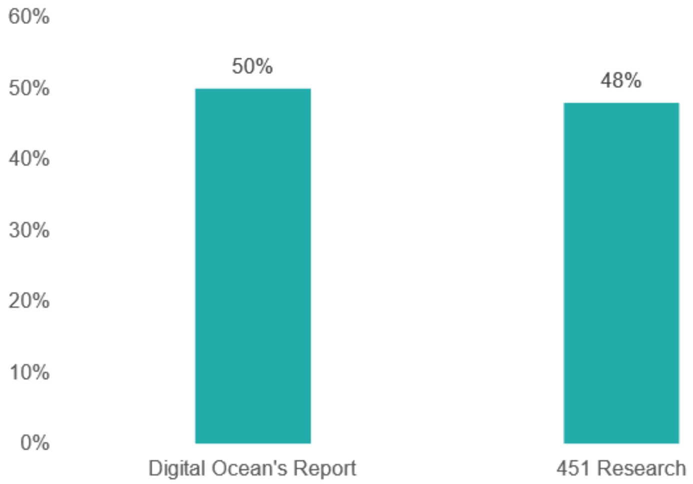

Source: Digital Ocean's Report (Link), 451 Research (Link).

5. THERE IS NO INDICATION OF SUBSTANTIAL BARRIERS TO ENTRY

CL WP's view on barriers to entry

78. The CL WP claims cloud service providers achieve economies of scale through sunk cost of investments in cloud infrastructure, bulk purchasing servers, operating efficiencies, and R&D spending.[72] It voices concern that *"large cloud providers benefit from economies of scale"* and that the factors causing scale effects *"when considered together they represent material barriers to entry and expansion"*.[73]

79. The CL WP hypothesises that while entry in the market might be possible for smaller entrants, scale in investment and R&D may represent an ongoing barrier.[74] For example, significant capital investments related to IaaS would be required. Additional investment in accelerator chips also increases capital investments by new entrants.[75]

80. While the CL WP notes that some effects alone might only be marginal, it concludes that there are material barriers to entry and expansion when combined.[76] The CL WP further mentions the importance of a large portfolio of cloud services as a potential source of

72 CMA's CL WP, para 7.6.

73 CMA's CL WP, para 7.43.

74 CMA's CL WP, para 7.41.

75 CMA's CL WP, para 9.19 b-c.

76 CMA's CL WP, para 7.43.

economies of scope. This supposedly gives certain providers advantages over their competitors.[77]

81. In the below we show that there is no support for barriers to entry inhibiting the functioning of the cloud segment, especially in terms of entry or growth of smaller providers, or that any entry costs are of particularly high magnitude. In this context, the high number of entrants and growth of smaller providers suggests that existing entry costs are unlikely to have a substantial effect on a firm's decision to enter or invest into the cloud segment. Moreover, we provide support for scale economies being of only limited magnitude, without limiting smaller providers' ability to compete in the segment.

82. In Section 5.1, we present additional analysis showing evidence *against* high entry barriers that would negatively affect the growth of smaller providers in the segment, in particular:

- **Smaller cloud service providers can both invest and secure funding**, which allows them to grow alongside the larger providers. Initial financing does not represent a deterrent if firms have access to capital markets (which, as shown below, they do).
- A hypothetical cost breakdown for a potential new entrant shows **costs (excluding datacentre costs) are not prohibitively high.**
- There are limited economies of scale and increasing costs of complexity. In particular, **economies of scale can be limited** because costs need to scale proportionally to size and may need to be duplicated rather than remaining constant, being spread over a larger customer base.

83. Section 5.2 shows that when looking at entry and growth of cloud service providers, the data shows an environment that is becoming even more dynamic and competitive over time, with many entrants obtaining significant scale both globally and in the UK. This speaks against the CL WP's concerns, as if they were true, one would expect to observe the following:

- Limited entry of new players.
- Signs of limited scale for smaller companies, including reduced growth and high sunk costs.
- Disproportional growth of large companies, and evidence of large-scale related cost-savings.

84. However, our previous submissions included a vast amount of economic evidence showing none of these outcomes are observed in IT services (or in the cloud services segment more specifically). We have previously shown that cloud service providers of all sizes are able to enter and gain scale.

85. Section 5.3 presents additional evidence on market outcomes [✂] and third-party studies in support of our previous findings: [✂]

86. We have previously provided ample economic evidence on the absence of barriers to entry and economies of scale that raise competition concerns. The evidence provided below is only complementary to our previous submissions, relying on the most recent data available.

77 CMA's CL WP, para 9.19 d.

5.1. Entry costs are moderate, while larger cloud service providers experience increasing cost of complexity

Smaller providers have good access to capital markets, as evidenced by high funding and investment activity

87. As shown below, smaller cloud service providers can both invest and secure funding that allows them to grow alongside larger providers. Cost of entry does not represent a deterrent if firms have access to capital markets- which, as shown below, they do.

88. The steep demand growth for compute services, driven by the rapid expansion of Generative AI solutions, has pushed cloud service providers to allocate significant investments towards expanding their supply of compute power. This associated compute power growth can be observed across cloud service providers of all sizes:

- For instance, in the last year, cloud service providers Alibaba and Voltage Park have invested ~$400mn and ~$500mn in cloud infrastructure to capture the value generated by the increasing demand for AI services, respectively.
- Further, after raising ~$320mn in funding in February 2024, Lambda Labs is currently in talks to raise another ~$800mn.[78]
- At the same time, even smaller players and startups like Aethir and Skytap have secured fundings of ~$9mn and ~$18mn, respectively, to deploy cloud compute services in the last year.
- Players such as Hive, who is reinventing the cloud from a centralised model that uses expensive physical servers to a distributed cloud infrastructure, aim to disrupt the industry with their innovative technology.[79]

89. Further rounds of investment and funding, by provider, are shown in Figure 12 and Figure 13.

78 Financial Times (2024), *Nvidia partner Lambda Labs seeks $800mn as AI computing demand soars.* (Link)

79 EU-Startups (2024), *Geneva-based Hive raises €12 million Series A to make sustainable distributed cloud computing for the masses*. (Link)

Figure 12: Investment and funding for cloud infrastructure of compute providers with valuation >$1bn

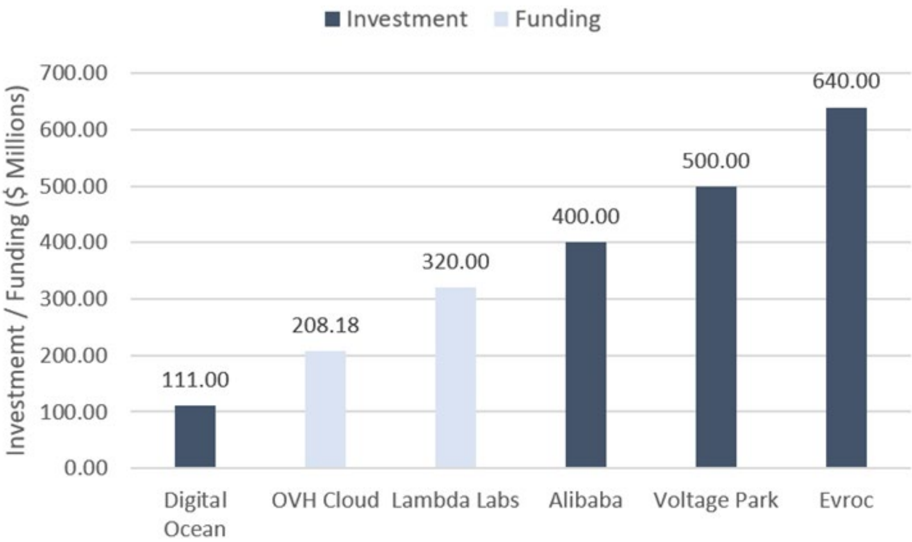

Source: Data used in the charts regarding investments and fundings Digital Ocean; OVHCloud; Lambda Labs; Alibaba; Voltage Park; Evroc

Figure 13: Funding for cloud infrastructure of compute providers with valuation <$1bn

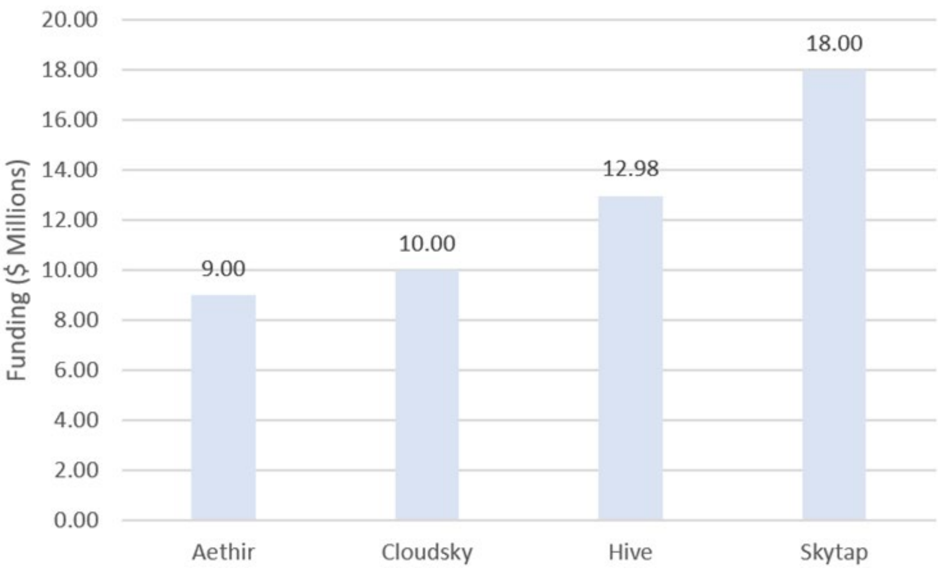

Source: Data used in the charts regarding investments and fundings Aethir; Cloudsky Technologies; Hive; Skytap

A hypothetical cost breakdown for a potential new entrant shows costs (excluding datacentre costs) are not prohibitively high

90. Evidence of increasing entry into the cloud infrastructure segment, presented below, can be explained by the ability of cloud service providers of all sizes to defer fixed costs (e.g., data centres construction spending) to third parties. In doing so, they only incur a cost which is proportional to the size of their business, instead of having to make significant upfront investments in fixed assets. The CL WP acknowledges that building a global network of datacentres represents a high upfront cost, but such an investment is not a necessity for

smaller cloud providers looking to enter the market. [80] This is for two reasons. First, leasing and colocation of datacentres are capital efficient strategies for smaller businesses, allowing them to incur much lower upfront capital costs, scaling capacity as the business grows. Second, in the UK, it is very common for cloud providers to co-locate most, if not all, of their datacentre capacity.

91. There is additional evidence that points towards entry costs not being prohibitively high:

- 2024 monthly datacentre rental prices in London for 250-500kW are between $160 and $195- below average rates worldwide.[81]
- On top of avoiding construction costs, colocation also saves on maintenance and staff costs, resulting in total costs that are 34% of the cost of owning a datacentre over a 3-year period, and 61% over a 10-year period.[82]
- Research has shown that colocation of datacentres can be up to 64% more economical than constructing a Tier 2 datacentre in-house.[83]

Limited economies of scale and increasing costs of complexity

92. **Infrastructure expansion entails significant cost duplication.** This is because of some cost items, such as staff training and maintenance, scale with network size. Moreover, other cost items increase more than proportionally to network size (e.g., security and management of the network) as larger infrastructures are more exposed to outages and harder to monitor or manage. Cloud service providers with smaller networks incur these costs in a smaller measure, such that their overall unit cost can be lower than those of larger ones.

93. **This makes economies of scale absent, or at least more limited than the CL WP suggests**. [✂].[84] In this context, it is important to note that owning many datacentres in different places entails the cost duplication mentioned above. [✂]. This means that increasing the number of datacentres incurs a proportionate cost equal to the building rent.

80 CMA's CL WP, page 129.

81 CBRE (2024), *Global data trends in 2024.* (Link)

82 Expedient, *Data center build vs buy calculator.* (Link)
Data calculated considering the percentage cost of colocation datacentre compared to owning a datacentre and assuming electricity cost of 22 cent per kWh, the avg electricity price in the UK in 2024 (Link), and assuming that 60% of power is consumed (avg).

83 Slice Up (2024), *What is Co-Location Data Center? Benefits and Types Explained.* (Link)

84 Source [✂] and ABI research (Link)

Figure 14: [✂]

5.2. Continuous entry into the cloud segment suggests absence of substantial barriers to entry

Many providers are entering and expanding in the UK cloud segment

94. Many players compete to serve the growing demand for cloud services, as evidenced by (i) the expansion of service offerings[85], and (ii) the entry of new providers. For example:

- Global investment in data centres is predicted to increase at a CAGR of 18% out to 2028.[86]
- Global data centre construction spending by "non-hyperscalers" is projected to increase by more than 50% between 2022-2030.[87]
- The prevalence of startup activity in cloud is evidence of the increasing number of new providers with ~40 start-ups globally in cloud industries within the past year.[88]
- Multiple providers – including Alibaba, T-Systems, Applied Cloud Computing, and Seagate Lyle[89] – have recently expanded into the UK segment.

95. This is further confirmed by our analysis based on the updated IDC data. Results, reported in Table 5 show that the number of active cloud service providers in the UK (excluding AWS, Azure, and GCP) has increased by [✂]% between 2017 and 2023 while total revenues have increased by nearly [✂]%. This is strong evidence that smaller cloud service providers can enter the UK cloud segment.

Table 5: [✂]

Growth in compute capacity is expected to continue for all cloud service providers

96. Further, growth in compute capacity is expected to continue for all cloud service providers. According to SemiAnalysis, total (AI and non-AI) global datacentre power capacity for "non-

85 Tencent (2021), *Tencent Cloud launches four new internet data centers in Bankok, Frankfurt, Hong Kong and Tokyo.* (Link)

Wasabi (2024), *Wasabi is expanding globally.* (Link)

86 SDX Central (2024), *Cloud hyperscalers, AI propelling data center capex growth.* (Link)

87 McKinsey & Company (2023), *Investing in the rising data center economy.* (Link)

88 Crunchbase Hub Profile (2024), *List of top cloud computing companies founded in the last year.* (Link)

89 TechRadar (2018), *Alibaba Cloud expands to UK.* (Link)

DCD (2016), *T-Systems comes to Virtus' London2 data center.* (Link)

The Week (2022), *An emerging leader in the Cloudtech space Applied Cloud Computing expands its global footprint further sets up its operations in the UK region.* (Link)

Intelligent Data Centers (2024), *Seagate Lyve Cloud expands UK operations and channel distribution partnerships.* (Link) hyperscalers"[90] is expected to more than double over the next four years, growing from 34 thousand to over 72 thousand MW. This implies a compound annual growth rate (CAGR) of 18% between 2022 -2028.[91] High growth is also predicted in other cloud segments. Moreover, AI global datacentre power capacity for "non-hyperscalers" is expected to grow by 476% over the next four years, from 7 thousand to 40 thousand MW, with an expected compound annual growth rate (CAGRs) of 79% between 2022 and 2028.[92] These developments are further demonstrated in Figure 15 and Figure 16 below.

Figure 15: "Non-hyperscalers'" total global datacentre critical IT power

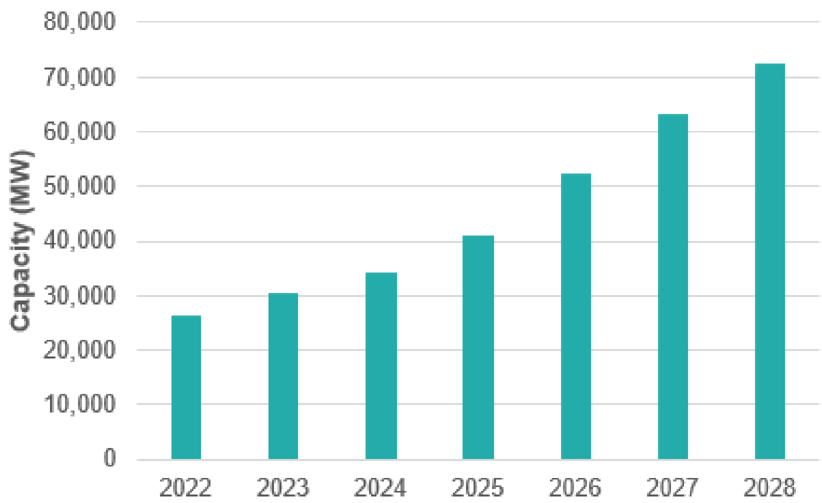

Source: Synergy; semianalysis

---

90 *"Hyperscale research is based on an analysis of the data center footprint and operations of 19 of the world's major cloud and internet service firms, including the largest operators in SaaS, IaaS, PaaS, search, social networking, e-commerce and gaming."* (Link)

91 Obtained by calculating the share of total (AI and non-AI) datacentre power capacity operated by "non-hyperscalers" and on-premise datacentres, using SemiAnalysis data (Link)

92 Obtained by calculating the share of AI-only datacentre power capacity operated by "non-hyperscalers" and on-premise datacentres, using SemiAnalysis data (Link)

Figure 16: "Non-hyperscalers'" AI datacentre critical IT power

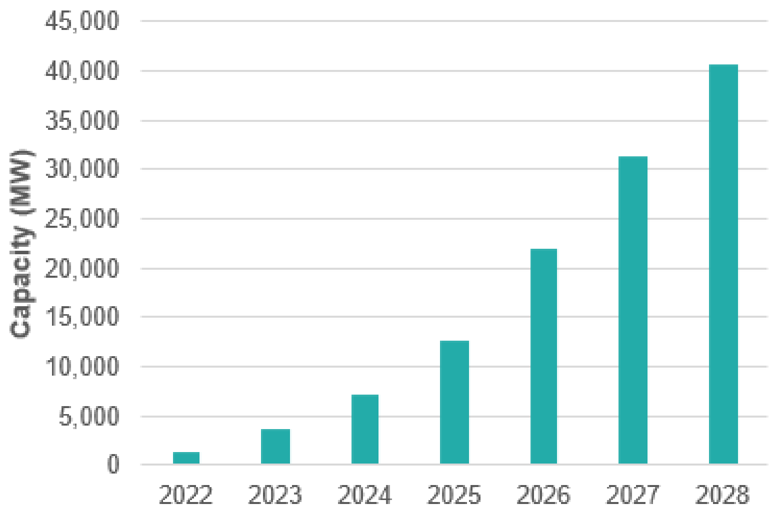

Source: Synergy; semianalysis

Firms active in adjacent segments can and do enter cloud services

97. There is a long history of entry into cloud segment from adjacent segments. This is also consistent with Amazon's own launch of its cloud product. For example:

- **Traditional IT hardware companies:** Dell started offering IaaS in 2011, and datacentre provider IO entered the cloud computing market in 2014.[93]
- **Telecom businesses** have typically entered the cloud segment by leveraging on their network infrastructure. Notable examples are AT&T, Deutsche Telekom (through its subsidiary T-systems), Verizon, and Huawei.[94]
- Similarly, **Google already established multiple datacentres** around the world to support its core business (search engine/online advertising) and then entered the cloud-first in PaaS (2008) and then IaaS (in 2011).[95]
- **Global wholesale marketplace** Alibaba launched its cloud in 2009.[96]

98. Recently, Nvidia has leveraged its position as the main supplier of GPUs to start offering cloud services to customers, further intensifying competition. In the past, cloud service providers were the major buyers of Nvidia's GPUs. These were used by cloud service providers for providing server space with additional power to customers. However,

93 Dell Blog (2012), *Dell Cloud Services: Building an Enterprise Solution Portfolio in 2012.* (Link)

Data Center Knowledge (2014), *IO Launches OpenStack Cloud on Open Compute Hardware.* (Link)

94 AT&T (2024), *Next Level Network.* (Link)

T-Systems, *Open Telekom Cloud: Public Cloud from Europe.* (Link)

Verizon, *Built with Unlimited space for your Digital World* (Link)

Huawei (2023), *Huawei Cloud Named a Leader in China's Government Cloud Market Three Times in a Row* (Link)

95 Datamation (2017), *Google Cloud Platform: History, features & pricing* (Link)

96 Alibaba Cloud (Link)

currently, Nvidia is directly competing with cloud service providers by supplying its GPU-accelerated servers to software developers (customers) in two ways: (i) entering into agreements with cloud service providers for leasing Nvidia-powered servers in cloud service providers’ datacentres and renting the same servers to developers;[97] and (ii) offering solutions in co-located datacentres by relying on non-cloud service provider facilities.[98]

5.3. Smaller cloud service providers have expanded rapidly, while there is evidence that scale economies for larger cloud service providers are limited

Smaller IaaS/PaaS providers are able to gain scale

99. [✂] [99] [100]

100. [✂] confirms that the number of cloud service providers has been growing each year. As demonstrated in Figure 17:

- Between H1 2017 and H2 2023, the number of cloud service providers with semi-annual cloud compute revenues above $5 million increased from [✂] to [✂], and the number of cloud service providers with semi-annual cloud compute revenues above $2 million increased from [✂] to [✂].
- Over the same period, [✂] firms doubled their revenues reaching semi-annual levels above $5m and [✂] firms doubled their revenues reaching levels above $2m in H2 2023.
- Excluding AWS/Azure/GCP, the median revenue of cloud service providers with compute revenues greater than $5 million in H1 2017 has increased by more than [✂]% between H1 2017 and H2 2023. Over the same period, median revenue of cloud service providers with compute revenues greater than $2 million in H1 2017 has increased by more than [✂]%.

Figure 17: [✂]

101. The following examples further illustrate how smaller cloud services providers compete with AWS in particular:

97 The Information (2023), *Nvidia muscles in cloud services, rankling AWS.* (Link)

Nvidia, *Expand horizons with Nvidia in the cloud.* (Link)

98 See Nvidia’s Colocation offerings, *DGX-Ready Colocation Data Centers* (Link)

99 [✂].

100 Similar results also hold for IaaS only and for UK+EEA and globally.

- Splunk, which has been acquired by Cisco, provides log analysis capabilities to customers, among other services,[101] and therefore acts as a competitor to AWS's OpenSearch Serverless.[102] For instance, some customers have stated that they prefer to use Splunk over OpenSearch as they feel that it provides a more secure services, with customers directly comparing the two services in public forums.[103]
- HiveMQ competes directly with AWS and Microsoft's IoT cloud services, with HiveMQ advertising itself to customers by promising unlimited connections, 99.95% uptime SLA and 24/7 support.[104] HiveMQ also promotes itself as the as "*the enterprise standard for reliable, scalable real-time communication in constrained IoT environments*," including on AWS itself.[105]
- Heroku is a direct competitor to AWS and other cloud services providers offering storage and database solutions and handling more than 60 billion requests every day. It is valued by customers for several reasons, including the ease by which new customers can deploy Heroku. [106] [107]
- Wasabi offers storage solutions that are compatible with Amazon S3. Wasabi advertises its services as providing a simpler and more budget friendly storage solution than competitors and also offers specific features to customers that may not be readily available elsewhere, for instance, the ability to make any object or storage bucket immutable to change or deletion.[108]

101 Splunk, *Splunk Observability.* (Link)

102 AWS, *What is Amazon OpenSearch Serverless?* (Link)

103 Kihm, T (2023). *Comparative Analysis: Splunk vs Elasticsearch vs OpenSearch.* (Link)

104 HiveMQ (2020), *Comparison of MQTT Support by IoT Cloud Platforms.* (Link)

105 HiveMQ, *Build Sclalable IoT Systems with MQTT and AWS.* (Link)

106 Heroku*, What is Heroku*? (Link)

107 RedSwitches (2023*), Heroku vs AWS: Decoding the Best Platform for your Deployment Needs*. (Link)

108 Wasabi*, Backup to Wasabi and Recover More Than Your Data.* (Link)

- A variety of cloud services providers offer solutions that compete with AWS's content delivery network service Amazon CloudFront (**"CloudFront"**). These include Akamai, Cloudflare, Fastly, CDN77, Stackpath, CacheFly, KeyCDN, and BunnyNet among others. Each service markets itself in its own way with several comparing themselves explicitly to CloudFront. For example, Cloudflare describes their R2 Storage as much cheaper than CloudFront.[109] Similarly, Akamai have a quarter of a million edge servers, deployed in thousands of locations around the world[110] and offer significant competition to CloudFront.[111] Fastly, by contrast, markets itself as a "real time" content delivery network that provides customers with the ability to instantly purge in 150 milliseconds and have access to real-time performance analytics and logs again offering a competing service to CloudFront.[112]

Cloud service providers other than Azure and AWS are growing rapidly

102. As shown in Figure 18 below, when considering the entire cloud segment (IaaS+PaaS+SaaS), other cloud service providers have earned more revenues than segment leaders, for the entire period H1 2017 – H2 2023. Indeed, AWS and Azure have a combined revenues of $[✂]in H1 2017, reaching $[✂] in H2 2023. Others cloud service providers instead have earned $[✂] in H1 2017 and $[✂] in H2 2023. Overall, [✂], the combined UK share of AWS and Azure in IaaS + PaaS + SaaS segment decreased by [✂]% in the last year (2022-2023).

103. Similar results hold when restricting the analysis to cloud infrastructure (i.e., IaaS+PaaS). As also shown in Figure 18, AWS and Azure have combined revenues of $[✂] in H1 2017, reaching $[✂] in H2 2023. Others cloud service providers instead have earned $[✂] in H1 2017 and $[✂] in H2 2023.

Figure 18: [✂]

Smaller companies often experience larger growth (IaaS/PaaS, UK)

104. Consistent with our previous submission,[113] we have used [✂] to consider the size-growth relationship for the top IaaS/PaaS providers offering compute services to see whether there is any evidence of scale effects in the cloud compute segment. We find that a lower scale does not prevent cloud service providers from competing and gaining customers.

105. The below figure shows the relationship between initial size and growth for the top 15 IaaS/PaaS providers, by UK revenues in 2023, based on [✂]. We find that:

- The compound annual growth rate (CAGR) between 2017 and 2023 is in the [✂]% range for most of the top 10 providers, regardless of their initial scale in 2017. CAGR is even higher for cloud service providers outside the top 10.

---

109 Cloudflare, *Why choose Cloudflare R2 vs Amazon R2?* (Link)

110 SeekingAlpha (2020), *Akami: Cybersecurity Is The Primary Catalyst.* (Link)

111 TechCrunch (2024), *Akami extends its edge-computing platform as it looks to challenge AWS, Azure and GCP.* (Link)

112 Backblaze, *Fastly vs AWS CloudFront: How Do the CDNs Stack Up?* (Link)

113 [✂].

- Many of the top 15 cloud service providers, especially those with lower revenues, have grown faster than the largest ones. For instance: [✂].
- As the chart shows, if anything, there is a negative relationship between size and growth, i.e., smaller cloud service providers are growing at faster rate.

Figure 19: [✂]

6. ABUNDANT EVIDENCE THAT INNOVATION IS OMNIPRESENT IN THE CLOUD SEGMENT

6.1. CL WP considers that the evidence on innovation is inconclusive

106. The CL WP acknowledges the reception of evidence that cloud services are characterised by strong innovation, that cloud service providers are continuing to innovate over time, and that innovation is expected to continue in the future.[114] The evidence relates to the number of patents issued, R&D investments, customer surveys and number of services released. AWS also provided evidence of the increasing number of AWS's services and features released every year between 2013-2023, and an increasing number of new versions/updates released for EC2.

107. Although the evidence clearly shows a persistent and increasing effort in innovation by AWS, to compete with other cloud service providers, the CL WP finds this evidence inconclusive. In particular:

- The CL WP argues this evidence does not tell us whether the quality of these services or the level of innovation reflects what may be expected in a well-functioning market.[115]
- The CL WP finds qualitative evidence from surveys subjective and less reliable. It also finds quantitative innovation measures to be proxies for the "quantity" of innovation rather than the "quality" of innovation.[116]
- The CL WP notes that the evidence received to date "*does not directly tell […] about the quality of or level of innovation in cloud services […] and whether the quality of services or level of innovation reflects what may be expected in a well-functioning market*".[117]

108. Before discussing the evidence in more detail, we note that the CL WP is content with only pointing to the hypothetical possibility that innovation could be even stronger in a hypothetical counterfactual. Such assertions do not prove anything and certainly fall short of establishing an adverse effect on competition.

114 CMA's CL WP, paras 6.12-6.14.

115 CMA's CL WP, para 6.15.

116 CMA's CL WP, para 6.18.

117 CMA's CL WP, para. 6.15. CL WP also argues that "it is difficult to evaluate evidence on quality and innovation" as this evidence is (1) less quantifiable than prices / is subjective.

6.2. Contrary to the CL WP's claim, the evidence is compelling

109. Not only does the CL WP fail to establish an adverse effect on competition, but the CL WP also fails to provide evidence that innovation is lacking.[118]

110. The reality is that there is abundant evidence that innovation is omnipresent in the cloud segment. In particular:

- Cloud service providers are increasingly investing;
- New cloud services and features are constantly being released;
- New patents are filed;
- The quality of cloud services keeps increasing;
- Efficiency keeps improving, allowing prices to decrease (see AWS's pass-on evidence).

111. These facts constitute strong evidence of a well-functioning and competitive market. In a market without competition, we would expect increasing prices, decreasing choice, decreasing quality, no investment, and no innovation. [✂].[119]

112. Similarly, the quality of these innovations (a key concern raised by the CL WP) is evidenced by the fact that the innovations are often directly responsive to customer requests and are aimed at ensuring the quality of a service for customers. For example, AWS, in conjunction with over a dozen security industry leaders and led by Splunk, developed the Open Cybersecurity Schema Framework ("**OCSF**") to try and establish a commonly-agreed-upon schema for data sources among different security-focused IT solutions. AWS [✂].[120]

113. For each of the metrics set out above, we observe the opposite. This underlines how the cloud segment is very competitive.[121] The CL WP should not dismiss tangible facts with speculative claims about how the cloud segment could be in a hypothetical world.

6.3. Complementary evidence on AWS's innovation activity

114. In addition to the previously submitted evidence on innovation, [✂].[122] [123]

---

118 No clear counterfactual to compare outcomes with what they might be in a well-functioning market (paras 6.17-6.19)

119 Section 6 ("Abundant Evidence That Innovation Is Omnipresent In The Cloud Segment") of the CRA CL Response provides further examples on how the high level of customer demand drives innovation in IT services.

120 AWS (2022), *AWS co-announces release of the Open Cybersecurity Schema Framework (OCSF) project.* (Link)

121 Context Consulting market research, slides 132 and 133, *Context Consulting market research, Ofcom: Cloud Services Market Research, Summary of Findings March 2023* (Link)

122 Acs et al. (2002) "Patents and innovation counts as measures of regional production of new knowledge", Research Policy. *Elsevier (2002)* (Link).

123 Marku, E. (2018) "Measuring Innovation Quality: A Patent Analysis", Journal of Business and Management. *IOSR Journal of Business and Management (IOSR-JBM) (2018)* (Link).

Figure 20: [✂]

115. [✂]. Results are in line with AWS continuously introducing innovative services: AWS launched more than [✂] new services and features each year between 2016-2023, offering more than [✂] services in 2023.

Figure 21: [✂]

6.4. CL WP's position with respect to investment decisions

116. While the CL WP acknowledged ongoing innovation and a relationship to quality, it makes several questionable statements in its assessment, from an economic perspective:[124]

- First, the CL WP alleges that cloud service providers have an incentive to innovate to "expand" the market, which would be different from competing with existing rivals. However, cloud service providers have an incentive to innovate to compete with rivals over future customer needs. This is a key sign of healthy competition. "Expanding the market" would in any event be the indication of the value added that is created, which benefits customers. It would also increase output which is another positive outcome typically observed in competitive markets.
- Second, the CL WP suggests that significant investments by large cloud providers may raise barriers to entry (because rivals would also need to invest), which could raise competitive concerns.[125] This does not make economic sense from our perspective and the logic should be reversed. Cloud service providers invest because of competition. Investments are pro-competitive, not anti-competitive. It is when there is no competition that firms no longer need to invest. The fact that AWS invests abundantly should be taken as evidence of healthy competition. The market would not be more competitive if some providers stopped investing. It seems that the CL WP suggests that high innovative activity can be used as a sign against functioning competition, but it does not provide any economic basis for such hypothesis.
- Further, as shown above, smaller competitors are also able to raise funds for the required investments with strong access to capital markets.

7. [✂]

117. [✂]. [126]

118. [✂].

119. [✂]:

- [✂].[127]

124 CMA's CL WP, para. 6.13

125 Ibid. para. 7.42

126 CMA's CL WP, para. 6.23

127 CMA's CL WP [✂].

- The EBIT margins for AWS have consistently been between 25% and 30% for the last 8 financial years.[128]
- [✂].[129]
- [✂].

120. [✂].

121. [✂].

7.1. [✂]

122. [✂].[130] [131]

123. [✂]:

124. [✂].

125. [✂].

7.1.1. [✂]

126. [✂].

127. [✂].

128. [✂].

Figure 22: [✂]

7.1.2. [✂]

129. [✂].

130. [✂].

Figure 23: [✂]

Figure 24: [✂]

7.2. [✂]

131. [✂].

7.2.1. [✂]

132. [✂].

128 CMA's CL WP, 6.40

129 [✂]

130 CMA's CL WP, para. 6.36

131 [✂].

133. Second, we highlight the CL WP's ROCE estimates already show a downward trend which, according to the CL WP, could be indicative of increasing competition.[132]

134. [✂].

[✂]

135. [✂].[133] *[134]*

136. [✂].[135]

137. [✂].

138. [✂].[136]

[✂]

139. [✂]:

140. [✂].

Table 6: [✂]

[✂]

141. [✂].

142. [✂].

143. [✂].

Figure 25: [✂]

7.2.2. [✂]

144. [✂].

145. [✂]:[137]

146. [✂].

132 CMA's CL WP, para. 6.33.

133 See Ofcom's Annex of the Final Report (para A2.119)

134 [✂].

135 [✂].

136 [✂].

137 CMA's CL WP, paras 6.41-6.44 and 9.19(b).

ANNEX

CL WP fails to fully deduplicate the list of unique UK customers

147. The CL WP calculates the proportion of multi-clouding customer by performing the following steps:

- Identification of multi-clouding customers (numerator): The CL WP performs 3 fuzzy matches (each fuzzy match is between 2 cloud service providers). The resulting combined dataset is (A) the list of multi-clouding customers.
- Identification of unique UK customers (denominator): The CL WP takes the full revenue datasets (B) of AWS, Microsoft, Google, and removes from (B) customers with customer names contained in (A), thus making (C) the list of not multi-clouding customers. The CL WP's analysis uses (A) + (C) as its denominator (i.e., all unique customers).

148. We identify two issues with this approach.

149. First, the CL WP fails to fully remove multi-clouding customers from (B). This is because, if a customer multi-clouds between cloud service provider 1 ("CSP#1") and 2 ("CSP#2"), it would appear in dataset (B): once under name #1 in the 1st cloud service provider's data and once under name #2 in the 2nd cloud service provider's data. The CL WP's analysis only removes from (B) entries with name #1.

150. The CL WP's duplication error can be illustrated with the following hypothetical example: [✂]

151. [✂].

152. Second, the CL WP's methodology results in "one-to-many" and "many-to-one" fuzzy matches in (A). [✂].

[✂]

153. [✂]:

154. [✂]