AWS Confidential Response to the CMA's Technical Barriers Working Paper

Index

I.   INTRODUCTION

1.   The CMA's Technical Barriers Working Paper of 6 June 2024 (the "**TB Working Paper**") claims that customers face technical barriers when integrating and operationalising multiple public clouds.[1] It also claims that customers experience technical challenges relating to switching between public clouds, which can affect their willingness to consider switching and the extent to which switching takes place.[2] Furthermore, the TB Working Paper puts forward that existing mitigations to these technical barriers do not fully overcome the challenges of using multiple public clouds and switching between them. It also notes that, overall, cloud services providers face a complex mix of incentives when deciding whether to support multi-clouding and switching.[3] Finally, the TB Working Paper puts forward potential remedies that could be imposed if the CMA were to find any adverse effects on competition ("**AEC**") in relation to technical barriers.[4]

2.   As explained in AWS's public response to the TB Working Paper of 25 June 2024 ("**AWS's Public Response**"),[5] the TB Working Paper has not established that technical barriers prevent customers from switching or multi-clouding in any way that is capable of harming competition, meaning that there is no AEC. To the extent that there are inherent technical barriers, these cannot be resolved through regulatory intervention. The potential remedies to remove the inherent technical barriers considered in the TB Working Paper are therefore unnecessary and unwarranted. They are also counterproductive as their disproportionate nature would severely harm innovation and customer choice.

3.   This confidential response provides a more detailed explanation of the points mentioned above, highlights the mixed, anecdotal, and contradictory nature of the evidence referred to in the TB Working Paper, and clarifies why the evidence relied upon by the CMA does not support regulatory intervention. It also refers to findings by AWS's external economists, Charles River Associates ("**CRA**"), which will be presented in CRA's Response to the CMA's Competitive Landscape Working Paper of 23 May 2024 (the "**CL Working Paper**"), that will be submitted shortly to the CMA ("**CRA's CL Response**").[6]

4.   The introduction of cloud services by AWS and cloud services providers' efforts to support interoperability have made switching and multi-clouding easier than ever before. Prior to the introduction of cloud services by AWS, switching IT services providers was often a prohibitively expensive, multi-year process. Switching IT services providers typically required substantial up-front investment in duplicate hardware from the new IT services provider and involved moving from one data centre to another, physically moving equipment to the new location.

---

[1]   TB Working Paper, paragraphs 1.6-1.7.

[2]   TB Working Paper, paragraphs 1.8-1.11.

[3]   TB Working Paper, paragraphs 1.12-1.16.

[4]   TB Working Paper, paragraphs 1.17-1.18.

[5]   AWS's Public Response, paragraph 56.

[6]   This confidential response also [✂].

Data and software would have to be restored, with failures manually rectified over a long period.[7]

5. When AWS pioneered the provision of cloud services, a key component was providing customers with the flexibility to design their solutions to meet their needs, including moving between, and interoperating across, different IT environments. AWS was the first IT services provider to offer pay-as-you-go pricing, creating an immediate reduction in the cost and burden of switching providers and solutions. Since AWS pioneered the cloud, IT services providers have been incentivised to support a wide range of operating environments, database types, storage types, and containers, to enable customers to replicate existing setups when switching without the level of manual work historically required. In addition, IT services providers have invested in automation and tools to enable easier migration, allowing transformation between virtualisation layers and between operating systems, amongst others, or supporting portable infrastructure-as-code ("**IaC**") deployments. As a result, customers today have an unprecedented degree of flexibility in choosing, combining, and switching between IT services providers to achieve the customer's desired IT solutions.

6. AWS designed its services to be interoperable since its incentive was, and remains, to win customers workload-by-workload, knowing that (i) many workloads would remain on-premises, and (ii) very few customers would ever be "all in" with a single IT services provider. Any suggestion that AWS would then intentionally design artificial technical barriers to prevent switching and multi-clouding ignores the dynamics of the sector and AWS's broader business strategy of prioritising the customer's needs. Moreover, by effectively ignoring switching between cloud services and on-premises and failing to compare it to switching between cloud services providers,[8] the TB Working Paper's emerging views fail to grasp that not only do inherent technical barriers also exist outside of cloud-to-cloud migrations, but that AWS and other cloud services providers have taken significant steps to make switching and multi-clouding significantly simpler.

II.   THE CMA HAS NOT SHOWN THAT TECHNICAL BARRIERS PREVENT SWITCHING AND MULTI-CLOUDING

7. Regulatory intervention is unnecessary, as the CMA has not established that technical barriers prevent customers from switching or multi-clouding in any way that is capable of harming competition, meaning that there is no AEC. To the contrary, customers can, *and do*, switch when they need or want to. The TB Working Paper's emerging views, namely that high technical barriers hinder switching and multi-clouding, are not supported by the evidence presented in the TB Working Paper.

8. First, the TB Working Paper's emerging views are based on the incorrect premise that switching and multi-clouding is limited.[9] As explained in AWS's Public Response,[10] the CMA's analysis uses a flawed methodology and limits the data in ways that leads it to underestimate switching and multi-clouding. In particular, the CMA's analysis (i) does not have a benchmark

---

[7]   See also AWS's Public Response, paragraph 58.

[8]   The TB Working Paper notes that the CMA has not examined migration costs from on-premises to the public cloud or *vice versa* (*see* paragraph 3.12 and Figure 3.1).

[9]   TB Working Paper, paragraph 4.7.

[10]   AWS's Public Response, paragraph 4(b).

against which to gauge whether observed multi-clouding levels are "high" or "low"; (ii) considers only three IT services providers (AWS, Azure, and Google Cloud Platform ("**GCP**")) ignoring other cloud services providers and all on-premises services providers; (iii) is distorted by the overreliance on very small customers, who are unlikely to multi-cloud for good economic reasons; (iv) is inconsistent in the way it considers that share of spend – as opposed to the share of customers – is relevant for the competitive assessment; and (v) relies on unstructured qualitative interviews producing anecdotal responses. When applying the necessary corrections to resolve these methodological errors, the CMA's own analysis shows a high prevalence of multi-clouding.[11] Moreover, an analysis [✂] shows that customers can and do switch when they wish to do so.[12]

9.   Second, the evidence presented in the TB Working Paper on customers' ease, desire, and the benefits of switching and multi-clouding is mixed and contradictory.[13] The TB Working Paper includes but ultimately disregards strong evidence showing that customers can switch and multi-cloud when they need or want to. For example:

   a.   Customers explained that they are integrating between and/or within applications and workloads on different public clouds.[14] Other customers said that multi-clouding is part of their business strategy.[15]

   b.   A subset of customers noted that, while there are challenges to multi-clouding, there are some workarounds such as using third-party tools or building custom solutions to connect services.[16]

   c.   Responses from other customers showed that they experienced minimal barriers to integration across multiple public clouds.[17]

   d.   One customer explained that it is flexible in placing new workloads due to the benefits of using multiple public clouds.[18]

   e.   Some customers said that, while they currently do not have a use-case for integrated multi-cloud, they may do so in the future.[19] The TB Working Paper also highlights evidence which shows that customers tend to put related workloads on the same public cloud to reduce operational complexity and prevent a reduction in resilience.[20] This customer feedback and evidence illustrates that customers' principal concern is

---

[11]   For further details, please see CRA's CL Response.

[12]   For further details, please see CRA's CL Response.

[13]   The TB Working Paper itself acknowledges that the evidence relating to technical costs that customers incur when integrating multiple public clouds is mixed (see paragraph 4.12).

[14]   TB Working Paper, paragraphs 4.8-4.10 and 4.17.

[15]   TB Working Paper, paragraph 4.17.

[16]   TB Working Paper, paragraph 4.13.

[17]   TB Working Paper, paragraph 4.14.

[18]   TB Working Paper, paragraph 4.23.

[19]   TB Working Paper, paragraph 4.17.

[20]   TB Working Paper, paragraph 4.19.

having *the ability* to multi-cloud or switch IT services providers, whenever it makes technical or economic sense to do so. However, neither the need nor the desire to multi-cloud or switch will always be there, as customers are often happy with their IT services provider and therefore see no need to multi-cloud or switch.

10. Reputable third-party surveys and anecdotal customer feedback beyond that collected by the CMA also demonstrate that customers can switch and multi-cloud and are not limited by technical barriers in doing so.[21] However, the CMA places limited evidential weight on these surveys, effectively dismissing any evidence that does not complement its emerging views.

11. Furthermore, the TB Working Paper notes that evidence gathered by the CMA shows that "*most customers face additional technical costs to operate more than one public cloud, whether or not they choose to integrate their use of them*."[22] This observation needs to be nuanced:

   a. First, it is not surprising that customers who multi-cloud incur some degree of higher costs than customers using a single cloud. This is due to the inherent technical costs of integrating multiple IT environments. However, customers may be willing to incur these inherent costs if it allows them to enjoy the perceived benefits of multi-clouding (*e.g.*, in terms of security and resiliency) and to take advantage of the "best in breed" IT services across different IT services providers. For example, a particular customer may perceive that one cloud services provider is best at storage and another is better at data analytics, so that customer may prefer to use both. While using both cloud services providers may come at an additional, inherent technical cost, it allows customers to achieve their desired outcome.

   b. Second, using multiple public clouds does not mean that the customer must pay for equivalent services across each cloud. For example, if customers use Amazon CloudWatch,[23] they do not need to pay for multiple overlapping observability solutions, because they can use Amazon CloudWatch (or a third-party solution) to observe each of their IT environments through one service.[24] If customers use AWS Systems Manager,[25] they do not need to pay for multiple management solutions for patching and upgrades, because they can use AWS Systems Manager (or a third-party solution) to manage each of their IT environments through one service. In other words, if customers use these solutions from AWS or equivalent solutions from other IT services providers, they do not necessarily incur greater costs from duplicating multiple similar solutions across more than one cloud, since they are using one solution, instead of multiple solutions, to manage multiple IT environments. The use of one solution also reduces the need for internal staff to retrain on new systems to observe or manage new IT environments.

---

[21] For further details, please see AWS's Public Response, paragraph 16, and CRA's CL Response.

[22] TB Working Paper, paragraph 4.11.

[23] See: https://aws.amazon.com/cloudwatch/.

[24] See, for example: https://aws.amazon.com/blogs/mt/observe-your-azure-and-aws-workloads-simultaneously-with-amazon-cloudwatch/.

[25] See: https://aws.amazon.com/systems-manager/.

III.   INHERENT TECHNICAL BARRIERS CANNOT BE RESOLVED THROUGH REGULATORY INTERVENTION

12.   The TB Working Paper puts forward that (i) there are very significant technical barriers that reduce customers' willingness to consider multi-clouding;[26] and (ii) there are technical barriers that significantly reduce customers' willingness to consider switching cloud services provider. [27] However, the evidence presented in the TB Working Paper concerning the technical costs associated with multi-clouding and switching is too mixed to raise any competition concerns. As explained in the sub-sections below, for almost every area explored in the TB Working Paper, there is mixed and contradictory customer feedback on whether such technical costs are perceived as an issue.

13.   These mixed and contradictory views illustrate that technical barriers are specific to each customer, as different customers value different aspects of cloud services, have different requirements, and face unique challenges depending on their use case. In other words, the existing technical barriers are not artificially imposed by cloud services providers across customers to prevent them from switching or multi-clouding but are inherent to IT services.[28] In fact, the TB Working Paper has not presented any evidence of cloud services providers imposing artificial technical barriers to prevent customers from switching or multi-clouding.

A.   The technical switching and multi-clouding costs with respect to core services are inherent to IT services and cannot be resolved through regulatory intervention

14.   The perceived concerns around feature differentiation, interface differentiation, and asymmetry of integrations set out in the TB Working Paper [29] are unwarranted. These differences and asymmetries are inherent to IT services and reflective of a healthy level of competition, as they indicate high levels of innovation and customer choice. They should therefore be encouraged rather than prevented through regulatory intervention.

15.   Feature and interface differentiation are the result of competition naturally taking place, as different IT services providers seek – in their view – the best way to offer a service or feature to customers, leading to differentiated products. Innovation via new services and features implies a natural level of differentiation, as new services and features tend to be substantially different to older ones. In other words, feature and interface differentiation are an expression of intense competition, as IT services providers take different innovative approaches to meeting customer needs. For example, most cloud services providers offer data storage services. All these services meet the same basic customer need, namely storing data. However, customers are looking for cost-effective storage options and high availability and data durability, among other needs, so this is where cloud services providers try to differentiate themselves by offering innovative features or strength in a particular type of storage. [✂].

16.   In addition, customer feedback included in the TB Working Paper recognises that technical barriers may naturally result from fundamental differences in how cloud services providers approach cloud services, pointing to differences in approaches, Application Programming

---

26   TB Working Paper, paragraph 4.16.

27   TB Working Paper, paragraph 4.27.

28   See also AWS's Public Response, paragraphs 62-65.

29   TB Working Paper, paragraphs 5.1-5.97.

Interfaces ("**APIs**"), technical implementations, tools, frameworks, methodologies, and best practices.[30] For example, AWS has chosen to adopt three availability zones per region. As a result, AWS has different underlying infrastructure and related APIs to other cloud services providers, a choice driven by AWS's view of how best to support its customers with greater availability and resiliency. While this may lead to some technical burden when switching between IT services providers, such as managing the transformation of software from different virtualisation layers (*e.g.*, Kernel-based Virtual Machine to VMware's ESXI server) and between different operating systems (*e.g.*, UNIX to Linux) or databases (*e.g.*, Postgre to SQL), customers view AWS's approach to cloud services as a key reason for choosing AWS.

i. Feature differentiation

17. The evidence presented in the TB Working Paper on the technical switching and multi-clouding costs related to feature differentiation is, at best, mixed. The TB Working Paper itself notes that "*customers had mixed views on whether there are differences in the features of comparable core services across clouds that make it harder to use multiple clouds or switch.*"[31] Such mixed feedback cannot support a conclusion that feature differentiation results in any AEC.

18. In fact, the TB Working Paper includes several pieces of feedback which illustrate that feature differentiation does not raise any concerns. For example:

    a. Some customers did not note any increase to the technical cost associated with switching or integrating multiple clouds coming from differences in the features of core services. One customer even said that there was a general parity of services between cloud services providers and working with multiple cloud services providers was easy.[32]

    b. Some customers and other organisations did not note significant challenges to switching or multi-clouding in relation to Infrastructure-as-a-Service ("**IaaS**") services.[33]

    c. An organisation said that basic IaaS services and their features are similar across providers and easily portable.[34]

19. Furthermore, when put in the proper context, none of the evidence included in the TB Working Paper points towards a need for regulatory intervention with respect to feature differentiation. For example:

    a. One cloud services provider claimed that some other cloud services providers use proprietary standards and code to make it very difficult for customers to switch certain workloads away from their infrastructure towards that of another cloud

---

[30] TB Working Paper, paragraph 4.3.

[31] TB Working Paper, paragraph 5.6.

[32] TB Working Paper, paragraph 5.8.

[33] TB Working Paper, paragraph 5.9.

[34] TB Working Paper, paragraph 5.12.

services provider or to impose artificial friction on customers' ability to share data between services running in different infrastructure environments.[35] AWS does not impose such artificial friction and does not recognise this concern in its own service design. AWS makes many of its software development kits ("**SDKs**") and APIs publicly available under open-source licences, and uses open protocols, interfaces, APIs, and data formats across services. [36] Furthermore, AWS's infrastructure services are fundamentally interoperable. For example, customers using the Amazon EC2 compute service have the option of running the popular proprietary Microsoft Windows operating system, as well as open-source operating systems including Amazon Linux 2, SUSE Linux, Red Hat Linux, Fedora, CentOS, Ubuntu, Debian, FreeBSD, and Oracle Linux. Each of these operating systems can also be run on almost any other compute platform, including in other IT environments. If this feedback is referring to managed services, AWS does offer managed services for popular open-source software. However, as explained in point (c) below, these services provide customers with more options for running their open-source workloads. While managed open-source services include plug-ins that integrate with other AWS services, changes to the underlying open source are documented (*i.e.*, so that customers can easily identify any relevant changes)[37] and do not prevent customers from easily transitioning to their next solution. The open-source portion of an application almost always requires various integrations and connections to other software that enable a full-featured application, regardless of where customers choose to run their open-source software. From a portability perspective, there is no material difference between integrations in AWS's managed open-source services and integrations that customers would need to implement had they self-managed the open-source service from the outset. In other words, nothing about AWS's managed open-source services raises artificial technical barriers to switching.[38]

b. An independent software vendor ("**ISV**") said that one area where there would be significant reworking is in AWS' DNS service – Route 53. In this ISV's view, Microsoft does not have an analogous service and therefore a move from AWS to Azure would require a substantial re-engineering of code.[39] This ISV feedback cannot credibly be relied upon as evidence of technical barriers to switching. First, based on publicly available information, it appears that Microsoft does have a managed DNS service, namely Azure DNS. [40] Second, even if this ISV feedback was true, cloud services providers should be allowed to have different services and should not be penalised just because their competitors do not offer identical or analogous services. Cloud services are not homogenous commodity goods. There is no singular "IT service" or "cloud service". Instead, cloud services providers compete aggressively to offer innovation by way of new features or services to customers.

---

[35]   TB Working Paper, paragraph 5.14.

[36]   For further details, please see [✂].

[37]   See, for example: https://docs.aws.amazon.com/neptune/latest/userguide/feature-overview-standardscompliance.html and https://docs.aws.amazon.com/keyspaces/latest/devguide/keyspaces-vs-cassandra.html.

[38]   For further details, please see [✂].

[39]   TB Working Paper, paragraph 5.18.

[40]   See, for example, the following article which compares the features of Amazon Route 53 and Azure DNS: https://www.peerspot.com/products/comparisons/amazon-route-53_vs_azure-dns.

    c.   A few organisations said there are differences in features of managed versions of the same open-source software across different clouds, which may make it more challenging to switch or use multiple clouds.[41] However, customers have a choice between managed services and open-source solutions, so there cannot be any customer harm. Indeed, AWS offers both managed services for popular open-source software alongside open-source solutions, so customers can choose based on their own priorities. Some customers prefer to use managed open-source services, including Amazon OpenSearch Service, Amazon EMR, Amazon Elastic Kubernetes Service ("**EKS**"), and Amazon Relational Database Service ("**RDS**"), because they do not want to focus on the administrative effort involved in running and maintaining secure, up-to-date open-source software. Other customers may prefer to bear such administrative effort, because it gives them the freedom to choose where and how to run the open-source software, and allows them easily to move their existing solutions between IT services providers.[42] In any case, [✂].[43]

ii. Interface differentiation

20. The evidence presented in the TB Working Paper on the technical switching and multi-clouding costs related to interface differentiation is, at best, mixed. The TB Working Paper itself notes that "*customers largely found APIs to be differentiated across cloud providers but had mixed views on the impact of that on switching and using multiple clouds,*"[44] "*there were mixed views about the differences in interfaces in databases and storage,*"[45] and "*some customers experience difficulty when integrating multiple clouds due to the differences in APIs of core cloud services. The evidence was mixed on the degree of this impact.*"[46] Such mixed feedback cannot support a conclusion that interface differentiation results in any AEC.

21. In fact, the TB Working Paper includes several pieces of feedback which illustrate that interface differentiation does not raise any concerns. For example:

    a.   Some customers and an organisation said that APIs were differentiated for similar services across clouds, but also noted that integration using different but open APIs did not cause major challenges.[47]

    b.   Another customer said that it does not find differences in APIs to be an insurmountable challenge because it can use translation layers to convert between different clouds or have 'over-the-top'-par solutions that support multiple versions at once in a multi-cloud architecture.[48]

---

[41]   TB Working Paper, paragraph 5.25.

[42]   For further details, please see [✂].

[43]   For further details, please see [✂].

[44]   TB Working Paper, paragraph 5.35.

[45]   TB Working Paper, paragraphs 5.53 and 5.67.

[46]   TB Working Paper, paragraph 5.68.

[47]   TB Working Paper, paragraphs 5.38-5.39.

[48]   TB Working Paper, paragraph 5.40.

    c.   An organisation said that generally differentiated APIs are not a technical blocker because workarounds are available. In addition, some customers noted more generally that there are easy ways to integrate multiple clouds and that differences in interfaces were not a challenge to integration.[49]

    d.   Some customers said that there are interfaces for storage and database which are common across clouds, especially for S3, and Postgres services, a type of open-source database. Some of these parties said the similar interfaces reduce the technical cost of using or switching these services across multiple clouds.[50]

    e.   A customer said the engineering effort required to move away from Amazon S3 would be low because it would just involve moving the data across. Similarly, a Postgres database would have a low engineering effort because it is the same data format across cloud services provider.[51]

    f.   Another customer reported that there are standard APIs for accessing S3 storage, and Amazon RDS Postgres and Cloud SQL Postgres which run the same underlying database.[52]

    g.   A cloud services provider said that the Amazon S3 API had become a *de facto* standard for storage.[53]

22.   Furthermore, when put in the proper context, none of the evidence included in the TB Working Paper points towards a need for regulatory intervention with respect to interface differentiation. For example:

    a.   Some customers and organisations said that the differences in APIs of core cloud services across clouds were significant and detailed how they increased technical costs when integrating multiple clouds.[54] It is simply not possible to have identical APIs across all cloud services providers. Each cloud services provider offers its own features and has its own underlying infrastructure, control planes, internal security protocols, etc., so cloud services providers' APIs are naturally different. However, as explained in paragraph 19.a above, AWS makes many of its SDKs and APIs publicly available under open-source licences, and uses open protocols, interfaces, APIs, and data formats across services. It also publishes extensive documentation, including, where relevant, on the differences between AWS services and the underlying open source. This allows customers and competitors to build solutions that interact with the service. It does not, however, reveal how AWS implements the actions requested by those who interact with that API. "Publishing" internal systems and interfaces between

---

[49]   TB Working Paper, paragraphs 5.41-5.42.

[50]   TB Working Paper, paragraph 5.59.

[51]   TB Working Paper, paragraph 5.60.

[52]   TB Working Paper, paragraph 5.61.

[53]   TB Working Paper, paragraph 5.64.

[54]   TB Working Paper, paragraph 5.43.

those systems in this manner would be unworkable and unnecessary, as discussed in paragraphs 23-27 below.

b. Some customers and other organisations noted that there are differences in the interfaces of storage and database services across clouds and that these can create difficulties. [55] Again, it is not possible to have a single API across all storage and database services, as each cloud services provider has its own unique features and underlying infrastructure. However, AWS supports a wide range of third-party database engines and database software. For example, as part of its relational database service offering, AWS offers Amazon RDS for MySQL, PostgreSQL, MariaDB, Oracle, SQL Server, and IBM Db2. [56] In addition, the AWS Database Migration Service ("**DMS**") makes it easy for customers to move databases to and from AWS. [57]

c. A customer said that Amazon S3 is not a standard for data storage, and that Azure Blob Storage and S3 are very different with different SDKs and different APIs so that someone adapting software to both must effectively do the same work twice. [58] However, this does not mean an inability to multi-cloud. As explained above, each cloud services provider offers its own features and has its own underlying infrastructure, control planes, internal security protocols, etc., so cloud services providers' APIs are naturally different. In other words, there will never be a one-size-fits-all API. However, AWS makes the S3 API available under an open-source licence and publishes documentation on it, so customers (and competitors, including Azure) can build applications that use S3 or create software that is directly compatible with the S3 API. In addition, as S3 can store standard storage types, customers can choose the storage format. S3 has released many features over the years that enhance customers' ability to interoperate S3 with other systems. [59] Furthermore, Table 9.1 of the TB Working Paper puts forward that the evidence the CMA has seen to date indicates that "*solutions compatible with Amazon's Simple Storage Service (S3) are used extensively throughout the industry.*" [60] Thus, while Amazon S3 may not be an absolute standard for storage, the evidence presented in the TB Working Paper suggests that there is good interoperability between services built using the Amazon S3 API.

23. The TB Working Paper invites stakeholders to raise any concerns around forcing cloud services providers to publish their APIs and SDKs in the open so that they can be reviewed. [61] The TB Working Paper is unclear as to whether this refers to customer-facing APIs and SDKs, which AWS makes available to customers (as explained in paragraph 22.a above), or a requirement for cloud services providers to open up their internal systems and interfaces between those systems, which is unnecessary, unworkable, and could cause considerable harm to customers.

---

55   TB Working Paper, paragraph 5.54.

56   For further details, please see [✂]. See also paragraph 54.c below.

57   For further details, please see [✂]. See also paragraph 55.a below.

58   TB Working Paper, paragraph 5.55.

59   For examples of these feature releases, please see [✂].

60   TB Working Paper, paragraph 9.26.

61   TB Working Paper, paragraph 5.37.

As explained below, this requirement would significantly slow down AWS's ability to implement updates to its services, including security upgrades and other improvements.

24.   [✂].

25.   [✂].[62]

26.   [✂].

27.   Exposing a service's internal interfaces could also allow AWS's competitors to reverse-engineer various functions and features, and discover their underlying structure, data flow or logic. Therefore, requiring cloud services providers to open up their APIs and SDKs would harm innovation, as competitors would no longer be incentivised to innovate if they knew their successful innovations would be publicly exposed.

iii.   Asymmetry of integrations

28.   The evidence presented in the TB Working Paper on the impact of asymmetry of integrations on customers' ability to integrate and switch between multiple clouds is, at best, mixed. The TB Working Paper itself notes that "*views on [the potential impact of asymmetries of integration on customers' ability to switch clouds] were mixed amongst customers*"[63] and "*the evidence is mixed on the impact of these asymmetries.*"[64] Such mixed feedback cannot support a conclusion that asymmetry of integrations results in any AEC and therefore does not justify further investigation or intervention by the CMA.

29.   In fact, the TB Working Paper includes several pieces of feedback which illustrate that asymmetry of integrations does not raise any concerns. For example:

   a.   Customers said that asymmetry of integrations does not impact their ability to integrate multiple clouds or that asymmetry of integrations is not the main factor impacting their ability to multi-cloud and switch.[65]

   b.   A customer said that where there are asymmetries of integration this impacts its ability to integrate multiple clouds, but cloud services providers are generally actively working on the development of their integration capabilities.[66]

---

[62]   For examples of some of the vulnerabilities that unsecured APIs may create, see "*Five-step plan for securing your APIs*" available at: https://www.youtube.com/watch?v=qJOaGLXrYvc. See also: https://apisecurity.io/owasp-api-security-top-10/.

[63]   TB Working Paper, paragraph 5.87.

[64]   TB Working Paper, paragraph 5.96.

[65]   TB Working Paper, paragraphs 5.86-5.87.

[66]   TB Working Paper, paragraph 5.85.

c. One ISV suggested that cloud services providers, such as AWS, are working to address asymmetry of integrations.[67] Another ISV said that AWS provided a solution to Amazon SageMaker requiring data to be moved into an Amazon S3 bucket.[68]

30. Furthermore, when put in the proper context, none of the evidence included in the TB Working Paper points towards a need for regulatory intervention with respect to asymmetry of integrations. For example:

a. The examples identified by Ofcom of asymmetry of integrations relevant to AWS[69] are misguided and do not support the conclusion that there are concerns. Ofcom only identified interoperability limitations in a small subset of AWS's cloud services, focusing on ten services for which there are competing software solutions available that customers can run on AWS (or elsewhere). In addition, the features identified as limiting interoperability were described inaccurately, exist alongside features that ensure interoperability, or are the product of an objective technical limitation.[70] Moreover, the TB Working Paper itself acknowledges that AWS has taken steps to minimise any technical barriers resulting from differing integrations.[71]

b. IBM submitted that a player with market power could reduce the functionalities of its products when used on third-party cloud services, as opposed to when used on its own first-party cloud services.[72] AWS has collaborated with IBM so that numerous IBM services run natively on AWS.[73] [✂].

B. The technical switching and multi-clouding costs with respect to ancillary services and tools are inherent to IT services and cannot be resolved through regulatory intervention

31. The TB Working Paper's emerging view is that customers seem to experience challenges to multi-clouding and switching arising from ancillary services and tools, due to differences in both their features and interfaces.[74] As explained below, these technical challenges are inherent to IT services (*i.e.*, they cannot be resolved through regulatory intervention) and do not support the conclusion that there are any AEC requiring regulatory intervention.

i. Identity and access management

32. With respect to identity authentication, the TB Working Paper itself concludes that the costs associated with integrating identity authentication across multiple public clouds are low due

---

[67] TB Working Paper, paragraph 5.88.

[68] TB Working Paper, paragraph 5.89.

[69] TB Working Paper, paragraph 5.71.

[70] For further details, please see [✂].

[71] TB Working Paper, paragraphs 5.76-5.77.

[72] TB Working Paper, paragraph 5.90.

[73] See, for example: https://newsroom.ibm.com/2022-05-11-IBM-Signs-Strategic-Collaboration-Agreement-with-Amazon-Web-Services-to-Deliver-IBM-Software-as-a-Service-on-AWS.

[74] TB Working Paper, paragraph 6.18.

to the adoption of standard identity protocols.[75] Therefore, there is no need for regulatory intervention in this area.

33.   In fact, the TB Working Paper includes several pieces of feedback which illustrate that identity authentication does not raise any concerns. For example:

   a.   Some customers suggested that IAM may not pose technical challenges for them when adopting a multi-cloud architecture.[76]

   b.   Oracle said that there are no significant feature differences between IAM services amongst cloud services providers and that, in general, the major platforms all seem to support common industry standards.[77]

   c.   Some customers reported that they have integrated identity authentication functionality between multiple public clouds, by either 'federating' multiple cloud providers' individual IAM services or using a third-party identity management service such as Okta. Some of these customers said they were able to integrate multiple cloud services providers' IAM services across public clouds and did not experience significant technical challenges.[78]

   d.   Other customers said they were able to integrate Microsoft Entra ID with other IAM services and this does not impact their choice of cloud services provider for non-Microsoft related workloads.[79]

34.   Furthermore, when put in the proper context, none of the evidence included in the TB Working Paper points towards a need for regulatory intervention with respect to identity authentication. For example:

   a.   A customer said that it is possible to integrate Okta into multiple clouds, but the integration is complex because the integration efforts mostly sit with the customer and entail continuous work as integrations need to be applied to any new applications, services, workloads.[80] However, the fact that integration is technically complex and requires continuous work from the customer does not mean that there is customer harm. To the contrary, AWS and other cloud services providers help customers by enabling Okta integration into multiple clouds in the first place.

   b.   A cloud services provider reported that, while all major cloud services providers provide the same level of interoperability with identity providers through standard identity protocols, such as Open Authorisation ("**OAuth**"), Security Assertion Markup Language ("**SAML**"), and OpenID Connect ("**OIDC**"), making integration with third-party IAM services easier, each cloud services provider has its own nuances in the

---

[75]   TB Working Paper, paragraph 6.61.

[76]   TB Working Paper, paragraph 6.24.

[77]   TB Working Paper, paragraph 6.25.

[78]   TB Working Paper, paragraphs 6.27-6.28.

[79]   TB Working Paper, paragraph 6.29.

[80]   TB Working Paper, paragraph 6.34.

implementations of these interoperable frameworks in terms of configuration and access mapping, which can impact consistency.[81] However, as described above in relation to interfaces,[82] these nuances are inevitable because each cloud services provider has its own unique features and underlying infrastructure which will always require an element of configuration unique to that cloud services provider.

35.   While integrating authorisation across IT environments may require more technical effort, AWS has introduced a policy language that customers can use outside of AWS. In 2023, AWS launched Cedar, an open-source language and authorisation engine that can be used to express fine-grained permissions[83] as easy-to-understand policies enforced in customers' applications (*i.e.*, to determine what person or systems should have access to IT resources or functions), and an SDK with a software library to evaluate them, both open-sourced under the Apache 2.0 Licence.[84] The Amazon Verified Permissions ("**AVP**") service uses Cedar to allow customers to manage fine-grained permissions in their custom applications. Cedar can also be used with a third-party solution.[85]

36.   ISV feedback included in the TB Working Paper acknowledges AWS's authorisation solution Cedar and its impact on customers' ability to multi-cloud and switch: "*An ISV said that access management in general lacks the same level of industry standards as identity management. However, it said that, in its view, AWS is working to solve this through the open-source Cedar Policy Language and that there may be other options.*"[86] AWS's development of open-source language such as Cedar shows that cloud services providers offer ways to mitigate the costs associated with integrating access authorisation across multiple public clouds, and regulatory intervention would not benefit those efforts.

ii.   Billing

37.   AWS agrees with the emerging view in the TB Working Paper that the impact of any inherent technical challenges associated with billing services and tools on customers' ability to multi-cloud and switch is low.[87] Therefore, these ancillary tools and services do not require further investigation by the CMA.

iii.   Observability

38.   The evidence presented in the TB Working Paper on whether customers are required to change observability tools when switching or multi-clouding is too mixed to raise any competition concerns and does not support a need for regulatory intervention.

---

[81]   TB Working Paper, paragraph 6.37.

[82]   See paragraph 22 above.

[83]   Fine-grained access control is a method of controlling access to data. It evaluates individual user credentials for authentication purposes.

[84]   See: https://aws.amazon.com/about-aws/whats-new/2023/05/cedar-open-source-language-access-control/.

[85]   See also [✂].

[86]   TB Working Paper, paragraph 6.45.

[87]   TB Working Paper, paragraph 6.68.

39.   The TB Working Paper itself highlights the industry efforts to standardise observability across clouds.[88] It notes that OpenTelemetry, which consists of formal specifications and open-source software packages, improves interoperability and makes it easier for customers to switch observability tools as they switch or integrate multiple public clouds.[89] It also mentions OpenMetrics, a newly developed open standard for metrics, and the ongoing work to develop a standardised query language for observability data.[90]

40.   In addition, AWS offers observability-based cloud services that support interoperability and multi-clouding. For example:

   a.   AWS Glue,[91] a data integration service, has a function which allows data to be catalogued even when it is contained outside of AWS in another IT environment. AWS Glue supports connections from multiple data sources, including on-premises.[92]

   b.   Amazon CloudWatch[93] can be used to monitor hybrid and multi-cloud IT environments, as it can monitor the performance and health of both AWS and non-AWS resources.[94] Moreover, [✂].[95]

41.   One customer noted in relation to multi-cloud that visibility of services that have been built or migrated is probably the hardest thing to deliver.[96] The fact that one customer identifies observability as the most difficult part of its own multi-cloud architecture, while the TB Working Paper notes that "*there appears to be [only] some impact*" on customers' ability to switch or multi-cloud arising from observability tools and services, shows that technical barriers are very specific to each customer and, as discussed above, simply a product of the inherent technical barriers in the sector that AWS works hard not to exacerbate, but cannot ever fully mitigate.

   ### iv.   Provisioning and orchestration

42.   AWS agrees with the emerging view in the TB Working Paper that provisioning and orchestration services and tools do not pose significant technical challenges to customers' ability to multi-cloud and switch.[97]

---

[88]   In this regard, AWS finds it puzzling that the TB Working Paper calls out industry support for open source in observability as a mitigant to technical barriers but does not do the same for similar support in relation to core services.

[89]   TB Working Paper, paragraphs 6.70-6.73 and 6.80.

[90]   TB Working Paper, paragraph 6.75.

[91]   See: https://aws.amazon.com/glue/.

[92]   See: https://docs.aws.amazon.com/glue/latest/dg/glue-connections.html.

[93]   See: https://aws.amazon.com/cloudwatch/.

[94]   See: https://aws.amazon.com/blogs/mt/monitor-hybrid-and-multicloud-environment-using-aws-systems-manager-and-amazon-cloudwatch/.

[95]   [✂].

[96]   TB Working Paper, paragraph 6.79.

[97]   TB Working Paper, paragraph 6.87.

C. **The technical switching and multi-clouding costs with respect to data latency, skills, and transparency are inherent to IT services and cannot be resolved through regulatory intervention**

i. **Latency**

43.   The evidence presented in the TB Working Paper on whether latency has a negative impact on customers' ability to multi-cloud and switch is too mixed to raise any competition concerns and does not support a need for regulatory intervention. The TB Working Paper itself acknowledges that "*there is conflicting evidence on the importance of latency for switching between public clouds.*"[98] Moreover, the customer feedback included in the TB Working Paper confirms that the technical costs associated with latency are inherent to integrating multiple IT environments and are very customer specific.

44.   In fact, the TB Working Paper includes several pieces of feedback which illustrate that latency does not raise any concerns. For example:

      a.   Some customers indicated that latency is not a significant concern to them. For instance, a customer said that latency has a relatively low impact on its ability to integrate between clouds, because some of its workloads transfer data asynchronously, so the time taken for data to pass between public clouds is less important.[99]

      b.   Other customers noted that latency may not be a barrier to switching. For example, a customer said that it previously moved some of its latency sensitive workloads but for a time carried on using the database stored in the previous cloud services provider in the same region. The customer explained that connectivity was fast and did not cause any issues.[100]

      c.   Customers reported that having a wide geographic dispersal of data centres can be important because it may enable lower latency between their workloads and end users, such as customers and factories.[101]

45.   The Jigsaw report found that customers who need to process or transmit data in real time see latency as an inherent barrier to multi-clouding because it will likely decrease the speed of their applications and workloads.[102] However, these findings only confirm that latency is an inherent feature of cloud infrastructure, as also acknowledged in the TB Working Paper's emerging views.[103] Some customers require ultra-low latency and sometimes only one cloud services provider is geographically close enough to offer it, or the customer chooses to purchase hardware that it can place in its own premises, because no third-party infrastructure

---

[98]   TB Working Paper, paragraph 7.12.

[99]   TB Working Paper, paragraph 7.11.

[100]   TB Working Paper, paragraph 7.13.

[101]   TB Working Paper, paragraphs 7.14-7.16.

[102]   TB Working Paper, paragraph 7.9.

[103]   TB Working Paper, paragraph 7.21.

can meet its needs. In such circumstances, there is not much cloud services providers can do, other than offering infrastructure that is even closer, which may not be feasible.

46. Latency itself is not a systematic constraint on data transfer between clouds that raises competition concerns and can be resolved through regulatory intervention. Latency is fundamentally an issue of physics, which cloud services providers look to mitigate on behalf of their customers by (i) continuously evaluating the best locations for their at-scale and edge compute infrastructure and their network infrastructure, and (ii) investing in low latency networking technology. Cloud services providers are incentivised to provide their customers with a network experience that minimises latency, cost, and complexity. For example, AWS has built a first-in-class global network by investing tens of billions of dollars in proprietary networking solutions such as custom semiconductors, equipment and software, and millions of miles of terrestrial and undersea cable. These investments improve transfer speeds, reduce lag, and increase security and reliability across the entire AWS global network, which ensures 99.999% availability. In addition, AWS has tried to overcome latency challenges by offering services such as AWS Outposts, which allows customers to run their workloads using AWS infrastructure on-premises and support their hybrid infrastructure needs.[104]

47. The TB Working Paper invites views on the use of methods which allow direct connections between clouds.[105] [✂].

   ii.   Skills

48. The evidence presented in the TB Working Paper does not support the conclusion that the need for cloud-specific skills amongst employees prevents customers from multi-clouding and switching. Therefore, it does not support a need for regulatory intervention.

49. In fact, the TB Working Paper includes several pieces of feedback which illustrate that the need for cloud-specific skills does not raise any concerns. For example:

   a. Some customers said that they have found it relatively easy to retrain staff in other cloud services providers' services, or that the skills they have are transferable between cloud services providers. For example, a customer said that cloud-specific skills are not an important factor when choosing a cloud services provider because staff can be retrained and skills are transferable across clouds.[106]

   b. An ISV also said it did not find it difficult to train its staff to operate across providers.[107]

   c. A customer said that the training offerings from cloud services providers have been an active enabler to multi-cloud. It said that it is not difficult to upskill on new platforms because cloud services providers are keen to help developers learn how to use their services.[108]

---

104   See: https://aws.amazon.com/outposts/.

105   TB Working Paper, paragraph 7.7.

106   TB Working Paper, paragraph 7.34.

107   TB Working Paper, paragraph 7.35.

108   TB Working Paper, paragraph 7.41.

50. Furthermore, when put in the proper context, none of the evidence included in the TB Working Paper points towards a need for regulatory intervention with respect to skills. For example:

   a. One customer said that cloud skills are not uniform or universal across different cloud services providers, as each has its own tools, frameworks, methodologies, and best practices that require specific knowledge and training. Therefore, switching between cloud services providers would entail a continuous investment in learning and developing the relevant skills and competencies for each cloud services provider.[109] AWS fails to see how this is different from any other technical industry. The use of technology naturally means that it is necessary for customers' employees to develop and maintain technical skills, but this is true for any technical industry and not specific to cloud services. Retraining is not a barrier to companies switching IT services providers or multi-clouding, because customers invest in IT training when they are switching to a new IT services provider. Customers are willing to make this investment because they believe they will save money with their new IT services provider or will benefit from other efficiencies (*e.g.*, services of better quality) in the long term. AWS and other IT services providers also offer a variety of training courses to help customers in their transition to a new IT services provider. While it is not uncommon for customers' staff to specialise in the technologies and services of their current provider, AWS tries to make it easy for customers to gain new skills if they are considering a switch to AWS.[110]

   b. One small cloud services provider said that (i) large cloud services providers, through comprehensive training and certifications, deeply influence students and professionals, swaying their cloud preferences from an early stage; and (ii) aggressive outreach can create a tech ecosystem where new talent is predominantly trained and biased towards a single cloud vendor, limiting multi-cloud knowledge, and curbing future diversification. AWS does not agree with this characterisation. There are many free and paid options for IT users to develop their skills with different IT services providers. These include training resources from cloud services providers, as well as training from independent providers that cover multiple IT services providers. Most cloud services providers, including smaller ones, offer training courses that make it easy for IT professionals to hone their skills for their services. AWS constantly invests to assist its customers in learning new IT skills and many customers use AWS's free training offerings, including cloud-agnostic training, to increase their workforce's overall cloud-based skills.[111] AWS's education programmes, such as AWS Educate,[112] AWS Academy,[113] and AWS re/Start,[114] are not meant to "deeply influence" students and professionals, nor do they involve "aggressive outreach." To the contrary, they

---

[109]   TB Working Paper, paragraph 7.26.

[110]   See also paragraph 102 below.

[111]   See also paragraph 102 below.

[112]   See: https://aws.amazon.com/education/awseducate/.

[113]   See: https://aws.amazon.com/training/awsacademy/.

[114]   See: https://aws.amazon.com/training/restart/.

allow students and professionals to gain core skills and a familiarity with the cloud, which has an overall positive impact on the workforce active in the IT industry.

51. Skills are generally transferrable between cloud services providers. IT developers and engineers must frequently learn the nuances of new technologies, services, and projects, so their skillsets are portable and adaptable to other IT solutions. While some retraining may be necessary any time a new service or technology is adopted, it is relatively easy for professionals to apply knowledge of one cloud to another and to transfer between cloud services providers. AWS and other cloud services providers support open-source technologies and industry standard protocols, further contributing to skill portability. For example, AWS allows third parties to use AWS APIs and SDKs that are compatible with multiple commonly used programming languages outside of AWS.

iii.   Transparency

52. The evidence presented in the TB Working Paper on whether a lack of transparency has a negative impact on customers' ability to multi-cloud and switch is too mixed to raise any competition concerns and does not support a need for regulatory intervention. For example, the TB Working Paper itself puts forward the following:[115]

"*We heard differing views on the availability and discoverability of information about potential technical challenges, and whether this impacts customers' ability to switch public cloud or integrate multiple public clouds. Some said that it can impede their ability to switch;* **others said that it does not**. *Separately, a few customers said that it can impact their ability to integrate between clouds;* **others said it does not**" (emphasis added).

53. AWS already takes many active steps to inform and educate its customers about the technical aspects of its services and the integration of those services with services from other cloud services providers. For example, AWS explains the programming language behind various tools that can be used to build on AWS,[116] and documents the changes to the underlying open source of its managed open-source services. Every AWS service has extensive documentation explaining exactly how it works. Further, AWS publishes blog posts and case studies to educate customers about how to work with these services and how to migrate various workloads and application types to other clouds. AWS also works with customers to build the solutions that meet their needs. For example:

a. As illustrated by [✂], AWS publishes extensive documentation on the interoperability of its cloud services with other cloud services providers and the "reversibility" of these services (*i.e.*, the ability of a customer to retrieve and move data from one IT environment to another).

b. AWS regularly publishes blog posts explaining to customers (i) how to use AWS's cloud services in multi-cloud scenarios; (ii) how to move workloads to AWS; and (iii) how to move workloads from AWS to other cloud services providers. For instance, between 1 January 2022 and 1 October 2023, AWS published 37 blog posts showing customers how to build and manage multi-cloud applications using AWS's services, bringing the

---

115   TB Working Paper, paragraph 7.47.

116   See: https://aws.amazon.com/developer/tools/?nc1=f_dr.

total number of blog posts on multi-clouding to over 60.[117] AWS also publishes blog posts that are dedicated to the topic of migration and provide guidance on how to move workloads to or from AWS.[118]

c.   AWS has published a customer-facing guide titled "Unpicking Vendor Lock-in", which provides information on understanding and mitigating switching costs when changing cloud services providers.[119]

IV.   AWS HAS INVESTED HEAVILY TO LOWER TECHNICAL BARRIERS AND SUPPORT INTEROPERABILITY BECAUSE IT IS COMMERCIALLY INCENTIVISED TO DO SO

A.   AWS supports multi-clouding and switching in various ways

54.   Even though the technical barriers are inherent to IT services, AWS has invested heavily to ensure that their impact on the ability of customers to multi-cloud and switch remains minimal. AWS refers to AWS's Public Response for examples of how AWS supports multi-clouding and switching.[120] By way of additional examples:

a.   AWS has invested in container technology. Containers allow customers to standardise how code is deployed, making it easy to build workflows for applications that run between on-premises data centres and IT environments of other providers. Containers are built on open-source standards and customers can run their containers everywhere they run their software, including in their own data centres, on AWS, and on other cloud services providers. Containers using formats established by the Open Container Initiative[121] can be run on AWS and can be moved to other IT services providers with almost no changes required. Container services have made it easy for developers to move their applications and their data across different IT environments. Containers also often benefit from additional management and orchestration capabilities assisting customers in how to deploy and run the containers. AWS has multiple services offering such capabilities, such as Amazon Elastic Container Service ("**ECS**") and Amazon EKS. In fact, Amazon EKS Anywhere allows customers to create and operate Kubernetes clusters on their own infrastructure outside of AWS.[122] [✂].[123] [124] AWS's container services is due in large part to customers increasingly moving their application operations, data processing, and machine learning workloads to use modern technologies based on containers, Kubernetes, and compatible tools.

---

[117]   For a list of publicly available links to these blog posts, please see [✂].

[118]   For publicly available links to a number of "blog channels" which contain such blog posts, please see [✂].

[119]   See: https://docs.aws.amazon.com/pdfs/whitepapers/latest/unpicking-vendor-lock-in/unpicking-vendor-lock-in.pdf#unpicking-vendor-lock-in.

[120]   AWS's Public Response, paragraph 66. See also [✂].

[121]   See: https://opencontainers.org/.

[122]   See also paragraphs 65-66 below.

[123]   [✂].

[124]   [✂].

b. AWS services support various standard protocols to make it easier for customers to enable communications and interactions between services – whether all services are on AWS or not – in a common way. For example, AWS services support:

- REST and HTTP APIs that support OIDC and OAuth 2.0;

- Internet security protocols (*e.g.*, SSL/TLS) and ciphers (*e.g.*, RSA ciphers) both for AWS's CDN service (CloudFront) and API endpoints;

- Identity protocols including SAML 2.0 and SCIM v2.0;

- Devices and clients that use the MQTT and the MQTT over WebSocket;

- Secure communications protocols to publish and subscribe to messages, and devices and clients that use the HTTPS protocol to publish messages (AWS supports these through AWS IoT Core);

- Service configuration using JSON and YAML files through the AWS Command Line Interface;

- X.509 SSL/TL certificates that follow the CA/Browser baseline requirements;

- As discussed in more detail above, standard containers and container management solutions, such as Docker and Kubernetes;

- JDBC and ODBC interfaces;

- Standard RTMP protocols for low-latency streaming; and

- Several different programming languages, including C++, Go, Java, JavaScript, Kotlin, NET, Node.js, PHP, Python, Ruby, Rust, and Swift.

c. AWS supports the ability of customers to install and run whatever database software they want on Amazon EC2, and there are no technical barriers to AWS's ability to do so. Customers can choose to run a number of relational and non-relational database options on AWS. If they want a relational database, they can use Amazon RDS to run Amazon Aurora, third-party proprietary database engines, including Oracle Database, Microsoft SQL Server, and IBM Db2, or open-source database engines, including MySQL, PostgreSQL, and MariaDB. If they want a non-relational database, they can use Amazon DocumentDB for MongoDB-compatible workloads, ElastiCache for Redis and Memcached compatible workloads, MemoryDB for Redis-compatible workloads, and Amazon Keyspaces for Apache Cassandra compatible workloads, on AWS. In all cases, customers can choose how they want to run the database software, be it self-managed or in a managed service like Amazon RDS or Amazon Aurora.

55. While there will always be inherent technical costs to switching, customers have a range of tools to migrate to and from AWS. Indeed, applications and workloads on AWS and other clouds are very portable. For example:

a. Customers can easily move databases to and from AWS using AWS DMS.[125] AWS DMS helps customers with the task of migrating databases and data warehouses both into and out of AWS as well as between different database types. AWS DMS supports a wide variety of database sources and targets, allowing customers to move data between both homogeneous and heterogenous databases (*i.e.*, like-to-like or dissimilar databases). For example, AWS DMS allows customers to move data to Oracle Database, Microsoft SQL Server, MySQL, MariaDB, PostgreSQL, SAP Adaptive Server Enterprise, and Redis, in addition to AWS managed services. AWS DMS is very user-friendly and requires little active steps from the customer.[126]

b. Many customers choose to run their applications on EC2 virtual machines ("**VMs**") using widely available and supported operating systems such as Linux or Microsoft Windows. The nature of IT means that making any configuration changes does incur some cost, but the costs of switching a VM from one IT environment to another are relatively low. For instance, the process of migrating AWS VMs to Azure is straightforward, may only take a few hours depending on data size and bandwidth, and there are many publicly available tutorials that explain the migration process step-by-step.[127]

B.   **The existing mitigations to the technical costs to multi-clouding and switching *do* effectively support efficient multi-clouding and switching**

56.   The TB Working Paper's emerging view that existing mitigations to the technical costs to multi-clouding and switching might not effectively support efficient multi-clouding and switching for customers is wrong.[128] In particular:

i.   **Adaptors**

57.   The TB Working Paper notes that customers can procure adaptors from third parties or develop them internally.[129] Based on this understanding, it concludes that customers bear the additional cost of developing or purchasing an adaptor.[130] However, the TB Working Paper fails to grasp correctly that customers can procure these adaptors from the cloud services providers themselves, often for free or as part of the cloud services providers' other services. For example, AWS offers over 200 adaptors to customers,[131] so they can connect to data and applications that reside in other clouds like Azure or GCP, Software-as-a-Service applications like Salesforce CRM, and on-premises solutions like Oracle databases. The adaptors enable customers to work with data "in place", without having to build complex Extract-Transform-

---

[125]   See: https://aws.amazon.com/dms/.

[126]   See: https://docs.aws.amazon.com/dms/latest/userguide/Welcome.html.

[127]   See, for example: https://learn.microsoft.com/en-us/azure/migrate/tutorial-migrate-aws-virtual-machines#run-a-test-migration and https://medium.com/@root.bhargav/migration-from-aws-instance-to-azure-cloud-vm-f4c2965e1f4f.

[128]   TB Working Paper, paragraphs 8.80-8.85.

[129]   TB Working Paper, paragraph 8.6.

[130]   TB Working Paper, paragraph 8.34.

[131]   See, for example, the Azure Data Lake Storage Connector for AWS Glue, which is offered for free: https://aws.amazon.com/marketplace/pp/prodview-k5ckoqmaimk5g.

Load ("**ETL**") pipelines to transport that data to AWS or transform that data. In any case, the fact that customers bear the additional cost of developing or purchasing an adaptor does not mean that adaptors do not effectively support efficient multi-clouding and switching.

58. In relation to the built-in S3 adaptor, which can connect Amazon Kinesis Video Streams ("**KVS**") to a non-AWS service, the TB Working Paper puts forward that a customer would need to have an Amazon S3 subscription to use the adaptor, adding an additional service purchase requirement for the customer.[132] This observation is inaccurate. First, the TB Working Paper refers to the "AWS Kinesis Video Streams, Integration manual" as the source for this statement.[133] However, the webpage linked to this source[134] (i) concerns Amazon Kinesis Data Streams integrations, *i.e.*, not KVS integrations; and (ii) does not mention a built-in S3 adaptor for KVS, nor the requirement to have an Amazon S3 subscription to use the adaptor. Second, customers are not required to pay for Amazon S3 to use KVS in conjunction with a third-party cloud. If a customer wants to ingest video from a camera using KVS and send it to a non-AWS service, they can do so without storing the data in KVS since KVS allows 0-day retention for recorded video. For example, customers can call the KVS PutMedia API to ingest the video to KVS, and the KVS GetMedia API to extract the video in real-time. Third, like other AWS services, Amazon S3 is sold on a pay-as-you-go basis, so customers do not need to subscribe to Amazon S3 to use it. In other words, to the extent there was a component of Amazon S3 required to make KVS work with a non-AWS service, payment for this component would simply fall under AWS's pay-as-you-go model.[135]

59. The TB Working Paper also includes feedback from a non-profit organisation which claims that proprietary APIs reduce developers' willingness to create adaptors that work with these APIs as these may be covered by method-patents.[136] AWS disagrees with this claim. First, AWS's APIs are open sourced under the Apache 2.0 Licence, which gives implementers a patent licence to any claims covering the APIs. Second, for an invention to be patented, it cannot merely be an abstract idea.[137] Therefore, it is extremely unlikely that APIs as such would be covered by method-patents. It would also be very difficult to patent a web service API adaptor given the amount of prior art covering APIs. Third, the high number of open-source adaptors that are currently available undermines the claim that method-patents are preventing developers from creating adaptors.[138]

---

[132] TB Working Paper, paragraph 8.34.

[133] TB Working Paper, footnote 303.

[134] See: https://aws.amazon.com/kinesis/data-streams/integrations/.

[135] For the sake of completeness, if this example was meant to refer to Amazon Kinesis Data Streams, it would also be inaccurate. Amazon Kinesis Data Streams does not require the use of any S3 adaptor. There are many third-party integrations that are available (*e.g.*, Databricks, OSS Spark, OSS Flink, Kafka Connect), so that customers can build custom applications to move and consume data from Amazon Kinesis Data Streams without ever using Amazon S3 or any S3 adaptor.

[136] TB Working Paper, paragraph 8.33.

[137] See, for example: https://www.gov.uk/guidance/manual-of-patent-practice-mopp/section-1-patentability.

[138] See, for example: https://pekko.apache.org/docs/pekko-connectors/current/s3.html; https://airflow.apache.org/docs/apache-airflow-providers-amazon/stable/connections/aws.html;

ii.   Abstraction layers

60.   The TB Working Paper explains that Terraform by Hashicorp is the most cited abstraction layer used by customers and acknowledges that Terraform (i) improves the ability of customers to manage their multi-cloud architecture by abstracting the differences between cloud services providers; and (ii) simplifies the complexity of provisioning and maintenance of management services which would allow ISVs to consolidate their own efforts.[139] The TB Working Paper also mentions other IaC tools, such as Ansible and Pulumi.[140] AWS supports all three IaC tools identified in the TB Working Paper.[141] In fact, [✕].[142]

61.   However, the TB Working Paper also notes that:

a.   Deploying Terraform (and similar cloud-agnostic IaC tools) results in extra costs for customers because their staff would have to learn new skills.[143] Again, the fact that customers need to do some extra work to enjoy the benefits of abstraction layers, such as Terraform, does not mean that these tools do not effectively support efficient multi-clouding and switching. Also, customers adopt the best practice of IaC to help them reliably deploy and version their cloud infrastructure configuration changes, so even when customers choose an IaC tool provided by a cloud services provider (*e.g.*, CloudFormation by AWS or Bicep by Azure), they still need to learn an IaC tool. Learning Terraform to perform multi-cloud deployments is no different than learning one of the other IaC tools needed to perform single cloud deployments. Further, even with the time needed to learn IaC tools like Terraform, these tools can ultimately save customers time and costs with respect to deploying and versioning their cloud infrastructure configuration changes.

b.   A customer, a market research and advisory organisation, and a cloud services provider said deploying third-party abstraction layers can create vendor-specific lock-in.[144] While existing IaC tools like Terraform can become integrated into customers' development pipelines, adding steps to transition from one IaC tool to another for certain use cases, they ultimately promote and facilitate multi-clouding, making it easier for customers to use services from multiple cloud services providers.

c.   Hashicorp's announcement in 2023 to change its source code licence from Mozilla Public License v2.0 (*i.e.*, open-source licence) to the Business Source Licence has the

---

https://pekko.apache.org/docs/pekko-connectors/current/dynamodb.html; and https://developer.harness.io/docs/cloud-cost-management/use-ccm-cost-optimization/optimize-cloud-costs-with-intelligent-cloud-auto-stopping-rules/add-connectors/k8s-connector-autostopping/.

[139]   TB Working Paper, paragraphs 8.36-8.40.

[140]   TB Working Paper, paragraphs 8.37 and 8.45.

[141]   See, for example: https://aws.amazon.com/blogs/apn/terraform-beyond-the-basics-with-aws/; https://developer.hashicorp.com/terraform/tutorials/aws; https://docs.ansible.com/ansible/latest/collections/amazon/aws/index.html; and https://www.pulumi.com/aws/.

[142]   [✕].

[143]   TB Working Paper, paragraphs 8.42-8.43.

[144]   TB Working Paper, paragraph 8.44.

potential to impact all third-party use of Hashicorp products, particularly Terraform.[145] AWS understands from Hashicorp's announcing blog post that this change was not targeted at cloud services providers, as the blog post noted that "*HashiCorp APIs, SDKs, and almost all other libraries will remain MPL 2.0*" and "*[HashiCorp] will continue to work closely with the cloud service providers to ensure deep support for [its] mutual technologies. Customers of enterprise and cloud-managed HashiCorp products will see no change as well.*"[146] Further, various entities have taken action to lessen any downstream impacts of this licence change, such as creating the OpenTofu open-source project,[147] which provides an open-source version of the Terraform feature set that was available prior to the licence change.

d. The Jigsaw report found that even cloud-agnostic IaC tools can require provider-specific configurations, which means that the main technical barrier for customers to multi-cloud or switch persists and building an abstraction layer while still using provider-specific Platform-as-a-Service ("**PaaS**") required a substantial amount of effort by them.[148] In many instances, provider-specific configuration differences exist because cloud services providers enhance their services at different rates and times, resulting in provider-specific configurations for those service innovations. While this may require additional learning, many customers choose to use these service innovations and are willing to spend the additional time if it means that they are receiving the best possible service and features.

62. Furthermore, The TB Working Paper notes that customers can procure abstraction layers from a third party or develop them internally.[149] Customers can also procure abstraction layers from their cloud services provider. For example, AWS itself offers an industry-leading "abstraction layer" called the "cloud development kit" ("**CDK**"), which together with AWS CloudFormation templates uses IaC to provision resources in AWS and other clouds using familiar programming languages.[150] AWS offers a similar abstraction layer for Kubernetes clusters called "cdk8s".[151] Both the CDK and cdk8s have been open sourced, so anyone can use and/or build on these technologies.

iii. IaaS and open-source software

63. The TB Working Paper explains that the evidence suggests that PaaS is more cloud-specific with more differentiated features and interfaces than IaaS.[152] This observation is not entirely correct. While it is true that moving a VM from on-premises to the cloud or between clouds is

---

145   TB Working Paper, paragraph 8.46.

146   See: https://www.hashicorp.com/blog/hashicorp-adopts-business-source-license.

147   See: https://opentofu.org/.

148   TB Working Paper, paragraph 8.47.

149   TB Working Paper, paragraph 8.7.

150   See: https://aws.amazon.com/cdk/. See also: https://docs.aws.amazon.com/prescriptive-guidance/latest/aws-cdk-layers/introduction.html, which explains how customers can use CDK to form different "layers" of abstraction.

151   See: https://cdk8s.io/.

152   TB Working Paper, paragraph 8.8.

easier than moving PaaS applications, the TB Working Paper overestimates the difficulty of moving PaaS applications from one cloud services provider to another. For example, it is relatively straightforward to move a .NET application from Microsoft Azure App Service to AWS Elastic Beanstalk,[153] and to move a serverless app from AWS Lambda to Azure Functions using Amazon CodeCatalyst.[154]

64. Furthermore, the TB Working Paper highlights (i) the trade-offs customers face when deciding whether to use IaaS and open-source services instead of proprietary PaaS alternatives; and (ii) some of the challenges of adopting open-source technologies compared to managed versions of open-source software.[155] While these trade-offs and challenges might exist, the key point here is that customers have the choice to use IaaS and open-source services if they want to. In such circumstances, there cannot be any customer harm.

iv. Containers and Kubernetes

65. The TB Working Paper acknowledges the benefits of containers and Kubernetes, citing customer feedback which clearly confirms that containers and Kubernetes improve the portability of customers' workloads and help customers manage applications across multi-cloud environments.[156] Use of containers is broad across the IT landscape. In Research's 2022 "Voice of the Enterprise" survey, more than 58% of respondents said they were actively adopting containers, with an additional 31% in proof-of-concept stages or planning trials. According to Gartner's "The CTO's Guide to Containers and Kubernetes",[157] by 2027, more than 90% of global organisations will be running containerised applications in production.[158] Moreover, as explained in paragraph 54.a above, [✂]. As containers are designed to enable easy portability of workloads, the broad adoption of containers suggests an increased customer desire for flexibility and multi-clouding, which AWS and other IT services providers need to accommodate to compete effectively.

66. However, the TB Working Paper also includes feedback stating that containers are not completely effective at overcoming barriers to multi-clouding and switching. This feedback needs to be put in the appropriate context. In particular:

a. Some customers, a supplier of professional services, and a market research firm noted that the underlying infrastructure and supporting ancillary services are built on proprietary provider technology, making a simple lift-and-shift difficult.[159] This observation only concerns customers who choose to use cloud services providers' managed containers services, such as Amazon EKS and Amazon ECS. However, in return, customers benefit from additional management and orchestration capabilities

---

[153] See: https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/migrate-a-net-application-from-microsoft-azure-app-service-to-aws-elastic-beanstalk.html.

[154] See: https://aws.amazon.com/blogs/devops/deploy-serverless-applications-in-a-multicloud-environment-using-amazon-codecatalyst/.

[155] TB Working Paper, paragraphs 8.52-8.57.

[156] TB Working Paper, paragraphs 8.66-8.71.

[157] See: https://www.gartner.com/en/documents/3988026.

[158] See also paragraph 54.a above.

[159] TB Working Paper, paragraphs 8.72.

assisting them in how to deploy and run the containers. As explained in paragraph 19 above, some customers prefer to use managed open-source services because they do not want to focus on the administrative effort involved in running and maintaining secure, up-to-date open-source software. Indeed, AWS built a managed Kubernetes service (*i.e.*, Amazon EKS) to alleviate the operational burden connected with running Kubernetes for customers who might find that useful. However, customers can also run the open-source version of Kubernetes on AWS directly using EC2, allowing them to maintain control over the provisioning and deployment of Kubernetes if they prefer. In other words, AWS does not require customers to use its managed container services. They are free to choose. For example, Amazon EKS Distro helps customers to build self-managed Kubernetes on AWS, on-premises, or on other clouds, while Amazon EKS Anywhere allows customers to create and operate Kubernetes clusters and applications on the infrastructure of their choice, whether on-premises or with other IT services providers. While these services directly provide critical hybrid and multi-cloud capabilities, there will always be some aspects of those services that are unique to AWS due to AWS's underlying infrastructure being unique compared to other IT services providers; that makes a "simple lift-and-shift" more difficult. That said, container services have made it significantly easier for developers to move their applications and their data across different IT environments.

b. A supplier of professional services said that containers (i) do not make workloads portable because of the dependencies to the larger cloud system; and (ii) are dependent on security, monitoring, storage in the wider cloud, and have many integration points with each of these services.[160] AWS does not agree that containers do not make workloads portable. Due to inherent technical barriers, customers will always need to make certain configurations for each IT services provider they use, but they can run containers everywhere they run their software, including in their own data centres, on AWS, and on other cloud services providers. As mentioned above,[161] customers can use services like Amazon CloudWatch to monitor their data across IT services providers, so monitoring is not necessarily a "dependency" that prevents containers from promoting portability. Also, storage is of course required for data that supports any workload being ported, but this is not a "dependency" that impacts the portability of the workload via containers.

c. A customer said that the use of open-source abstraction technologies like Kubernetes increased management costs, required significant upfront investment, and can lead to greater system complexity, which can have unintended negative impacts, such as reduced platform stability and increased complexity of change.[162] If a customer chooses to run open-source Kubernetes itself, it will naturally incur management costs and face a certain level of complexity. However, as explained in paragraph 19.c above, the customer will also enjoy more freedom in running the open-source software and maintain control over the provisioning and deployment of Kubernetes. If customers do not want to bear these management costs and the burden of operating Kubernetes, they can choose to use AWS's managed open-source services,

---

[160]   TB Working Paper, paragraphs 8.73.

[161]   See paragraph 40.b above.

[162]   TB Working Paper, paragraph 8.74.

such as Amazon EKS and Amazon ECS, or those of other cloud services providers. The key point is that customers have a choice, so there cannot be any customer harm.

v. Cloud services provider and ISV workarounds

67. The TB Working Paper presents evidence from ISVs and cloud services providers that have designed workarounds to allow customers to integrate with other cloud services providers.[163] The TB Working Paper invites stakeholders to provide evidence and views on other workarounds that facilitate multi-clouding and switching.[164]

68. AWS is aware of numerous ISVs and cloud services providers that are actively promoting multi-clouding and have designed products to support this. As just a few examples, AWS cites relevant marketing materials from Snowflake, Nutanix, Veeam, and Cutover.[165] As discussed in Section C below, these are not "workarounds" but rather integrations and connections that are made possible due to the open nature of cloud services, and the strong commercial incentive to support customers using multiple solutions across IT services providers.

vi. Conclusion

69. While the inherent technical costs to multi-clouding and switching may lead to some extra work for customers, IT services providers are continuously innovating to provide solutions to mitigate these costs. The evidence presented in the TB Working Paper does not support the emerging view that these mitigations do not effectively support efficient multi-clouding and switching. Customers do have *the ability* to multi-cloud efficiently and switch when they need or want to, and these mitigations help them in doing so.

C. Cloud services providers have, and will continue to have, strong commercial incentives to lower technical barriers and support interoperability

70. As explained in AWS's Public Response [166] and [✂], [167] IT services providers are heavily commercially incentivised to ensure that customers are able to multi-cloud and switch between different IT services providers, if and when they wish to do so.

71. The TB Working Paper's emerging view that cloud services providers are not incentivised to ensure interoperability and to lower technical barriers as they would risk losing existing customers, is flawed.[168] It fails to recognise the stark reality that cloud services providers are more likely to lose customers if they do not support interoperability. This is because IT services providers compete on a workload-by-workload basis even for existing customers. If a

---

[163]   TB Working Paper, paragraphs 8.75-8.78.

[164]   TB Working Paper, paragraph 8.79.

[165]   See: https://www.snowflake.com/blog/dont-settle-for-multi-cloud-aspire-to-cross-cloud/; https://www.nutanix.com/solutions/hybrid-multicloud; https://www.veeam.com/blog/multi-cloud-resilience-tips.html; and https://www.cutover.com/content-hub/best-practices-automating-cloud-disaster-recovery.

[166]   AWS's Public Response, paragraphs 13-15.

[167]   [✂].

[168]   TB Working Paper, paragraphs 1.15 and 8.98.

customer cannot host a third-party service on AWS or cannot have an AWS service interoperate successfully with a third-party service it wishes to use, the customer will simply switch the workload away from AWS to another IT services provider or choose another IT services provider for the specific workload in the first place. Ensuring interoperability therefore is essential for cloud services providers who wish to attract new workloads and maintain existing ones.

72.   The fact that customers value interoperability is not a "simple observation" as noted in the TB Working Paper. [169] AWS's commitment to interoperability is an important reason why customers choose AWS in the first place. AWS's incentives are therefore not hypothetical – they heavily influence the commercial decisions AWS takes to ensure ease of multi-clouding and switching. While supporting interoperability makes it easier for customers to use another IT services provider, there is a strong commercial benefit to AWS offering services that interoperate with other IT services providers' services. AWS has often considered this trade-off and ultimately favoured interoperability. For example, as reflected in AWS's internal document highlighted by the TB Working Paper,[170] despite containers making it easier for customers to switch workloads away from AWS, AWS has invested heavily in containerisation technology as customers expect to use such technology in conjunction with the IT services provider of their choice, both to move applications to and from AWS's infrastructure. As evidenced by the figures in paragraph 54.a above, [✂].

73.   AWS's and other cloud services providers' commercial incentives to lower technical barriers and support interoperability will grow even stronger in the future, as new entrants and new workloads are being made available to customers. Better solutions are constantly being developed, and individual IT services providers will never reach a point where they can independently develop and offer every piece of technology that a customer wants or needs. It is a commercial imperative for IT services providers to offer interoperability with innovative solutions that customers demand, as they would otherwise be unable to attract or retain customers. Generative artificial intelligence is a recent example of how innovation is prompting customers to look for the best-in-class solutions to meet their IT needs and incentivising IT services providers to enable interoperability between these solutions and their other existing services.[171]

74.   The TB working Paper notes that "*cloud providers with large shares of supply may have a stronger incentive now than in the future to compete for customers and workloads that are new to the public cloud.*"[172] AWS finds it difficult to square this statement with the premise in the CL Working Paper that cloud services providers generate the majority of their revenue from existing customers. [173] It appears contradictory to claim that, on the one hand, competition is less likely to be driven by new customers, but, on the other hand, large cloud services providers currently have a strong incentive to compete for new customers. In any

---

[169]   TB Working Paper, paragraph 8.103.

[170]   TB Working Paper, paragraph 8.96.

[171]   [✂].

[172]   TB Working Paper, paragraph 8.101.

[173]   CL Working Paper, paragraph 5.49.

case, cloud services providers have strong commercial incentives – now and in the future – to support interoperability for both new and existing customers and workloads.

75.   In addition, the TB Working Paper puts forward that "*[w]hilst potential customers considering migrating from on-prem to cloud may value the absence of lock-in and give cloud providers a possible incentive to facilitate switching and multi-cloud, these incentives may be weaker or changed in a future where a lot of the migration to cloud has occurred.*"[174] This assertion is misguided, as it wrongly assumes that:

    a.   Switching from on-premises to the cloud is inherently easier than switching between clouds. This is not the case. In fact, customer feedback included in the TB Working Paper confirms that a switch between cloud services providers would cost a similar (if not lower) amount to the initial migration from on-premises IT infrastructure to a cloud services provider.[175] Had the CMA examined the initial migration costs from on-premises to the public cloud, the TB Working Paper would not have made this incorrect assumption.

    b.   Switching from on-premises to the cloud is permanent once it has happened. This is again not the case. As explained in AWS's Public Response,[176] on-premises IT providers continue to exert significant competitive pressure on cloud services providers. This competitive pressure extends from potential customers that are on-premises to existing customers of cloud services providers, with many customers either switching or considering a switch from the cloud to on-premises or a hybrid cloud and on-premises infrastructure. AWS frequently sees on-premises solutions as competition for customer workloads, amongst both existing on-premises customers and for established UK customers of cloud services, with some customers choosing to move away from AWS to on-premises solutions where they believe this suits their needs better. Also, a switch from on-premises to the cloud can be partial, as the switch may only concern some of the customer's workloads but not all of them. Therefore, the switch is neither complete nor permanent.

V.   REMEDIES SEEKING TO RESOLVE INHERENT TECHNICAL BARRIERS ARE UNNECESSARY AND WOULD BE DISPROPORTIONATE

76.   The TB Working Paper has not presented any evidence of AEC in relation to technical barriers, as explained in Sections II and III above. Therefore, the different potential remedies considered in the TB Working Paper are unnecessary. These potential remedies would also not achieve their desired outcome, as the technical barriers presented in the TB Working Paper are inherent to IT services and cannot be resolved through regulatory intervention. They would also severely harm innovation and customer choice, which would entail a disproportionately steep cost considering the lack of evidence of competition or customer harm.

77.   The TB Working Paper notes that some potential remedies could relate to all cloud services, others to all IaaS, all PaaS, all core services or all ancillary services and tools, while others may

---

[174]   TB Working Paper, paragraph 8.101.

[175]   TB Working Paper, paragraph 4.26.

[176]   AWS's Public Response, paragraphs 8-12.

be specific to an individual cloud service.[177] However, any regulatory intervention must be targeted at an identified concern. Wider intervention would be disproportionate and would risk significant unintended consequences, which would likely in themselves distort competition.

A. The potential remedies considered in the TB Working Paper are unnecessary and would severely harm innovation and customer choice

i. Potential remedies that require common standards for cloud services providers

78. A potential remedy proposed in the TB Working Paper is to increase standardisation of cloud services and interfaces. This potential remedy would require cloud services providers to follow voluntary or mandatory common standards for some or all cloud services and interfaces. It would require cloud services providers to contribute to the development of new standards or to adopt existing standards.[178]

79. Imposing mandatory regulator-enforced standards is incompatible with dynamic and innovative industries, such as the IT sector. IT services providers large and small are constantly developing new technologies to meet customer needs and solve novel problems, each offering a unique set of services, features, reliability, availability, scalability, and price. User requirements are also constantly changing, and technology needs to be developed and applied in different ways to address new challenges. Technological innovation does not follow a linear path, and experimentation with different solutions is required to discover their benefits and drawbacks. Enforced standards run counter to the dynamic needs of the industry and tether innovation to a static set of technologies and practices. Even if those standards encapsulate the optimal solutions at the time they are set, they will likely quickly become outdated.

80. For example, modern container technology has been evolving since the 1980s and has its roots as far back as 1960.[179] The first container concepts were shown in 1979, and the first commercial success was only seen in 2004. The popularity of containers was driven by their ability to make applications movable between different IT environments. Innovation has driven today's open standards for containers.[180] Currently, Kubernetes is the most popular container orchestration service.[181] However, ten years ago Docker was viewed as the pre-eminent containerisation service, demonstrating the pace at which innovation occurs and that in-demand technologies change. If Docker had been an enforced standard when it seemed to be the "superior" solution for containers, customers would have lost out on all the Kubernetes innovation and technology that only became available years later. Similarly, if the x86 ISA found in Intel Corporation's processor chips had been an enforced standard when it seemed

---

[177] TB Working Paper, paragraph 9.6.

[178] TB Working Paper, paragraphs 9.15-9.51.

[179] With the IBM Mainframe: https://www.ibm.com/docs/en/zosbasics/com.ibm.zos.zcourses/zcourses_MFHWenvironment.pdf.

[180] A condensed history of containers: 1960 IBM Mainframe, 1970 System/370, 1979 UNIX chroot, 1982 BSD chroot, 1993 FreeBSD Jails, 2001 Linux VServer, 2002 Linux namespaces, 2004 Solaris Zones, 2005 Open VZ, 2006 Linux cgroups, 2008 LXC, 2011 Apache Mesos, 2013 Docker, 2014 Kubernetes.

[181] More than 8,000 companies and 75,000+ individuals have contributed to Kubernetes. See: https://www.cncf.io/reports/kubernetes-project-journey-report/.

to be the "superior" solution for processors, consumers would have lost out on all the Arm-based innovation, technology, and competition that only became available years later.[182]

81.   AWS also agrees with the potential for unintended consequences of common standards highlighted in the TB Working Paper, namely (i) the risk of introducing competing standards, thereby adding complexity and ambiguity to the market; (ii) the risk of impairing innovation; (iii) the risk that the standard-setting body may not be sufficiently informed and that the standards set may not represent an optimal solution; and (iv) the risk that the standard-setting body may not be sufficiently independent.[183]

82.   As concerns the design considerations set out in the TB Working Paper:

    a.   The TB Working Paper puts forward that standardisation may be more appropriate for some ancillary services and tools, such as billing, because in general there appear to be lower levels of innovation with these services and tools when compared to other cloud services.[184] AWS does not necessarily agree with this observation, as ancillary services and tools are also characterised by a high level of innovation. The development of Cedar by AWS is a clear example of such innovation.[185] Further, the TB Working Paper highlights billing as an example of an ancillary service where standardisation may be more appropriate. However, the TB Working Paper also notes that the impact of any inherent technical challenges associated with billing services and tools on customers' ability to multi-cloud and switch is low.[186] Therefore, there is no need for standardisation with respect to billing.

    b.   The TB Working Paper puts forward that it may be appropriate to limit standardisation of interfaces to APIs used by databases and storage services, because the Jigsaw report found that challenges migrating databases and storage services are among the most commonly raised issues by customers.[187] As explained in paragraph 22 above, it is not possible to have a single API across all storage and database services, as each cloud services provider has its own unique features and underlying infrastructure. In other words, there will never be a one-size-fits-all API. However, migrating databases and storage services has never been easier due to the services and tools offered by AWS[188] and other cloud services providers, so there is no need for standardising APIs used by databases and storage services.

    c.   The TB Working Paper puts forward that standards could require cloud services providers to publish some or all of their APIs.[189] As explained in paragraphs 23-27

---

[182]   [✂].

[183]   TB Working Paper, paragraphs 9.43-9.45.

[184]   TB Working Paper, paragraph 9.29.

[185]   See paragraphs 35-36 above.

[186]   TB Working Paper, paragraph 6.68.

[187]   TB Working Paper, paragraph 9.30.

[188]   See, for example, paragraph 55.a above.

[189]   TB Working Paper, paragraph 9.33(a).

above, such a requirement is unnecessary, would be unworkable, and could cause considerable harm to customers.

d. The TB Working Paper puts forward that standards could require cloud services providers to provide equivalent access to their cloud services (so-called "equivalence of input").[190] This requirement would be practically infeasible and would harm innovation, as interoperability of novel services and features with third-party services or protocols may not always be possible. Due to the complexity of IT services and operational considerations such as security, fidelity, and latency, integration with a cloud services provider's proprietary services may be the only technically feasible way to offer a given feature given resource constraints and other customer-driven priorities. For example, Amazon Redshift developed a "bulk upload" feature whereby customers can upload data in parallel to multiple compute nodes to speed up the rate at which data is ingested into the data warehouse. Due to the significant coordination on development and customisation that is required to enable any two services to coordinate data upload across parallel nodes, this feature is purpose-built for Amazon S3. Similarly, AWS has integrations between Amazon Aurora, Amazon RDS, Amazon DynamoDB to Amazon Redshift with the recently launched "zero-ETL integrations" to simplify the user experience and performance,[191] and with Amazon Kinesis and Amazon Managed Streaming for Kafka (MSK) with Redshift streaming ingestion capability.[192] However, there are other AWS storage and database services that upload data to Amazon Redshift in the same manner as users of third-party services. Data can be loaded to Amazon Redshift from any source, such as other cloud services providers or on-premises data/file storage, via AWS Data Pipeline or SSH, or through industry-standard JDBC/ODBC driver interfaces. Amazon Redshift also provides a Data API that allows users to access data from Amazon Redshift with any application, including on-premises, cloud-native, and containerised services.

e. The TB Working Paper puts forward that standards could require cloud services providers to make the output of their cloud services comparable between cloud services providers (so-called "equivalence of output").[193] The TB Working Paper views this requirement as consistent with the "functional equivalence" requirement under the EU Data Act.[194] Such a requirement would harm innovation and is very difficult for cloud services providers to work with, as a key component of competition involves product differentiation through innovation. If cloud services providers are expected to facilitate customers achieving comparable outcomes when they move cloud services providers, then this delays or increases the cost of innovations as cloud

---

[190] TB Working Paper, paragraph 9.33(b).

[191] See: https://aws.amazon.com/blogs/big-data/announcing-zero-etl-integrations-with-aws-databases-and-amazon-redshift/. [✂].

[192] See: https://docs.aws.amazon.com/redshift/latest/dg/materialized-view-streaming-ingestion.html.

[193] TB Working Paper, paragraph 9.33(c).

[194] Article 2(37) of the EU Data Act defines the concept of "functional equivalence" as "*re-establishing on the basis of the customer's exportable data and digital assets, a minimum level of functionality in the environment of a new data processing service of the same service type after the switching process, where the destination data processing service delivers a materially comparable outcome in response to the same input for shared features supplied to the customer under the contract.*"

services providers wonder what will meet an "equivalence" standard across their broad array of customers who utilise their services in different ways. For example, a cloud services provider considering a new development to its storage service would have to factor in how that new development might impact its customers' ability to achieve a comparable outcome if they switched to another storage service.

f. The TB Working Paper puts forward that standards could require cloud services providers to ensure that their services are functionally the same (*i.e.*, the input, the way the services operate, and the output).[195] Such a requirement would be equally unworkable and harm innovation for the same reasons as described above.

83. As concerns the questions for consultation included in the TB Working Paper:[196]

a. "*Do you agree with our characterisation of common standards in cloud services and interfaces, as set out in Table 9.1 and, if not, why do you disagree?*" – AWS respectfully submits that Table 9.1's observations on common standards are too limited. For example, Table 9.1 does not cover operating systems, database types, programming languages, and other "standards" that cloud services providers can either support or not support (subject to licensing restrictions in some cases), *i.e.*, the many "standards" listed above at paragraph 54.b; all of these "standards" support interoperability. Furthermore, Table 9.1 highlights that there is limited adoption of common standards with respect to billing, which could contribute to potential difficulty in comparing billing data between cloud services providers. This observation should not raise concerns, as there is no need for standardisation with respect to billing services and tools.[197]

b. "*Do you agree that common standards and standardisation in general are more appropriate for IaaS, ancillary services and tools and interfaces (APIs) than for more abstracted types of PaaS services?*" – AWS respectfully submits that regulator-enforced common standards are not appropriate for any of these services. While common standards are likely more suitable and achievable for IaaS than for PaaS, it is not possible to break ancillary services and tools out from more abstracted PaaS in a way that does not create significant uncertainty for cloud services providers.

c. "*What are the benefits and harms of introducing common standards for IaaS, ancillary services and tools and APIs?*" – Please see paragraphs 79-82 above.

d. "*What are the benefits in having common standards for cloud services where there is more abstraction?*" – AWS respectfully submits that there are no benefits in having common standards for cloud services where there is more abstraction. Standardisation could effectively stop the development of PaaS tools if developers were forced to support all cloud services providers (*i.e.,* as they may not have the necessary resources to provide such support).

---

[195]   TB Working Paper, paragraph 9.33(d).

[196]   TB Working Paper, paragraphs 9.46-9.51.

[197]   See paragraph 82.a above.

    e. *"Which standards setting bodies have sufficient independence and could set common standards for one or more of the types of cloud service or interfaces?"* – [✂]¹⁹⁸ ¹⁹⁹ ²⁰⁰ ²⁰¹ ²⁰²

    f. *"Should the standards apply to all cloud providers that offer a relevant cloud service or should standards only apply to the largest cloud providers?"* As explained in paragraphs 104-105 below, if the CMA were to impose common standards on cloud services providers, these standards should be applied to all IT services providers rather than targeting only some cloud services providers to avoid harmful unintended consequences.

84.   [✂]:

    a. [✂] ²⁰³ ²⁰⁴ ²⁰⁵

    b. [✂] ²⁰⁶

ii. Potential remedies that use principles-based requirements

85. Another potential remedy proposed in the TB Working Paper is to require cloud services providers to comply with a set of principles to address the design of cloud services and/or the interfaces that customers, ISVs, and other parties use to access or integrate with those services. Such a principles-based approach to standard-setting would require cloud services providers to either (i) achieve a stated outcome (*e.g.*, cloud services providers must allow for comparability between bills raised by different cloud services providers); or (ii) take certain actions to achieve an outcome (*e.g.*, cloud services providers should include certain headings on their bills to ensure comparability between bills from different cloud services providers).²⁰⁷

86. This potential remedy entails a high risk of harmful unintended consequences. As noted in the TB Working Paper,²⁰⁸ imposing principles-based requirements would leave too much room for interpretation, so that there would be a high risk of misunderstanding and circumvention. Indeed, as each customer's IT architecture and use of services is different, applying broad principles that aim to achieve a certain outcome (*e.g.*, ease of switching and/or multi-clouding) would require cloud services providers to account for all individual customer circumstances,

---

198   [✂].

199   [✂].

200   [✂].

201   [✂].

202   [✂].

203   [✂].

204   [✂].

205   See also paragraphs 79-82 above.

206   [✂].

207   TB Working Paper, paragraphs 9.52-9.64.

208   TB Working Paper, paragraph 9.61.

which is impossible. In other words, it is not possible to design principles at the cloud services provider level that overcome all dependencies or other choices a customer might make when using a cloud services provider's services. Also, as this potential remedy would mainly be interpretation driven, it would result in an uneven approach and uncertainty for customers and cloud services providers on what outcome to expect. There would also be considerable uncertainty for cloud services providers on whether they would have achieved the desired outcome, as each customer would have a different opinion (*i.e.*, similar to the "functional equivalence" requirement under the EU Data Act, as discussed in paragraph 82.e above). Therefore, defining the right set of principles and subsequently monitoring and enforcing them would be an extremely difficult, if not impossible, task.

87.   As concerns the questions for consultation included in the TB Working Paper:[209]

a.   "*Is it preferable to impose broader principles-based requirements on cloud providers, or more prescriptive rules/common standards?*" – AWS respectfully submits that neither option is preferable, as both potential remedies entail a high risk of harmful unintended consequences, as explained in paragraphs 79-82 and 86 above.

b.   "*What broad principles should cloud providers be required to comply with, if we pursued a principles-based approach?*" – It is not possible to come up with a set of broad principles, as the principles would have to differ depending on the cloud services and technologies involved, as well as the specific concern the principles were seeking to address. Also, the broader the principles, the more room for interpretation, misunderstanding, and uncertainty. Conversely, the more specific the principles, the more principles the CMA would need to come up with.

c.   "*Should all cloud providers be required to comply with a principles-based approach or only the largest cloud providers?*" – As explained in paragraphs 104-105 below, if the CMA were to require cloud services providers to comply with a principles-based approach, these principles should be applied to all IT services providers rather than targeting only some cloud services providers to avoid harmful unintended consequences.

iii.   Potential remedy to improve the interoperability of cloud services through the use of abstraction layers

88.   Another potential remedy proposed in the TB Working Paper is to require cloud services providers to provide or support the provision of abstraction layers that automate or simplify the operation of the cloud technologies that sit below them in the cloud stack, allowing customers to centrally manage and use multiple clouds through a single standard interface. This potential remedy would aim to ensure that cloud services providers take steps to facilitate the development of so-called "internal developer platforms" and similar abstraction solutions.[210]

89.   Mandatory regulator-enforced abstraction layers are not desirable to help ensure interoperability and would limit the ability of IT services providers to differentiate their

---

[209]   TB Working Paper, paragraphs 9.62-9.64.

[210]   TB Working Paper, paragraphs 9.65-9.89.

services from their competitors to the detriment of innovation and customers. The management and administrative level (*i.e.*, the level where an abstraction layer operates) is one of the key areas in which IT services providers compete and offer differing functionalities for customers. Mandating an abstraction layer would require IT services providers to agree to use identical metrics and functionalities for the tools that sit below the abstraction layer. For example, Amazon S3 thinks about storage in "buckets" but Microsoft's storage services often think about storage in terms of "folders". If you were to mandate a single abstraction layer for both services, you would need to decide whether to use "buckets" or "folders" as the underlying storage metric.

90.   In other words, every time AWS would offer a new functionality or administrative tool, it would need to agree with other IT services providers how to "surface" this in the abstraction layer. This would be especially difficult in instances where other IT services providers do not offer the same tool. Mandating an abstraction layer would effectively drive IT services providers to offer the lowest common denominator in terms of services. The TB Working Paper itself acknowledges that requiring an abstraction layer could lead to reduced innovation, as it would reduce perceived differences between cloud services providers. Moreover, it notes that "*in certain circumstances, the underlying difference in functionality may be beneficial to customers, if that difference leads to better performing products/services, as providers are better able to introduce new functionality that has non-equivalent output*s."[211]

91.   Abstraction layers would also have a detrimental impact on the performance of services and would in the end harm customers. An abstraction layer is an additional layer of complexity that sits on top of everything else. Including an additional operational layer makes it more complicated to provide and operate services. An abstraction layer would therefore by its very nature have a negative impact on reliability and speed, thereby harming customers. As noted in the TB Working Paper,[212] it would also increase cloud services providers' costs if they were required to provide or support abstraction layers for free. This may result in increased prices for customers as cloud services providers would likely seek to recoup the costs of the abstraction layers elsewhere.

92.   As concerns the questions for consultation included in the TB Working Paper:[213]

   a. "*To what extent do the products already offered by the cloud providers, such as Azure Arc and Google Anthos, act as an abstraction layer and allow customers to operate across multiple public clouds?*" – Azure introduced Azure Arc to simplify the adoption of management tools provided by Azure on VMs hosted outside of Azure. Customers can register non-Azure VMs as Arc-enabled servers by installing the Azure Connected Machine agent on each VM, simplifying enablement and data collection for supported Azure services.[214] By enabling Arc on non-Azure VMs, Azure can provide customers with a single place to manage their VMs across multiple clouds in domains such as observability (with Azure Monitor), security (with Azure Sentinel), and governance (with Azure Policy).

---

[211]   TB Working Paper, paragraph 9.78.

[212]   TB Working Paper, paragraph 9.80.

[213]   TB Working Paper, paragraphs 9.81-9.89.

[214]   See: https://learn.microsoft.com/en-us/azure/azure-arc/servers/agent-overview.

b. "*To what extent do IaC products already offered by ISVs, such as Terraform (by HashiCorp) or Pulumi, act as an abstraction layer and allow customer to operate across multiple public clouds?*" – As acknowledged in the TB Working Paper, [215] IaC products already offered by ISVs, such as Terraform, improve the ability of customers to manage their multi-cloud architecture by abstracting the differences between cloud services providers, thereby allowing customers to operate across multiple public clouds. [216]

iv. Potential remedies to increase interconnectivity and reduce latency

93. Another potential remedy proposed in the TB Working Paper is to require cloud services providers to (i) connect third-party data centres by building direct fibre lines between all the data centres of different cloud services providers in a data centre hub; and/or (ii) make data centre space available for other cloud services providers. [217]

94. The physical complexity – in some cases, infeasibility – and tremendous costs associated with building direct fibre lines between all cloud services providers' data centres, or co-locating all cloud services providers' data centres, does not justify the marginal latency benefit for some specialised use cases, and would inevitably end up harming customers in the form of increased prices. Also, while latency for some specific transfer paths may improve marginally, others may deteriorate because it is simply not possible for every source location, cloud services provider location, and destination location to be geographically equidistant.

95. More specifically, requiring cloud services providers to connect data centres is unnecessary, practically unfeasible, and could have detrimental unintended consequences because:

a. There is already a network connecting all data centres, namely the internet, so it does not make sense to impose a requirement to create physical connections between all cloud services providers. Moreover, as explained in paragraph 47 above, [✂], so there is no need for this potential remedy.

b. Connecting data centres would require significant additional investment in terms of infrastructure to practically ensure interconnectivity. As a result, cloud services providers would be forced to dedicate significant time and resources to managing the increased costs brought by these changes, especially if coupled with the potential remedies considered in the CMA's Egress Fees Working Paper of 23 May 2024 on cloud services providers' ability to charge for use of their network. This significant additional investment would likely result in reduced innovation and increased prices for customers, as cloud services providers would seek to recover the cost of these investments elsewhere. Moreover, cloud services providers would likely prioritise other cloud services providers with the most traffic to/from their infrastructure, meaning that there would only be better connectivity and lower latency between already established cloud services providers.

---

[215] TB Working Paper, paragraphs 8.36-8.40.

[216] See also paragraphs 60-62 above.

[217] TB Working Paper, paragraphs 9.90-9.112.

      c.   Requiring cloud services providers to connect their data centres would mean that far fewer sites would be suitable for establishing new data centres as they would need to be reasonably close to other data centres to ensure interconnectivity. This would drive up the price for cloud services providers wishing to establish new data centres. Also, geographical concentration across cloud services providers would reduce systemic resiliency, increase correlated outage risk, concentrate power consumption in areas that may already be constrained, and risk creating new, high-risk targets with potential national security implications. Indeed, this potential remedy would lead to a range of security risks, as data centres would become accessible through numerous different entry points, some of which would be out of control of the cloud services provider tasked with ensuring the safety of the data it holds in those data centres.

      d.   AWS also agrees with the additional possible unintended consequences set out in the TB Working Paper.[218]

96.   Similarly, requiring cloud services providers to make data centre space available for other cloud services providers is unnecessary and would be practically unfeasible for the reasons set out in the TB Working Paper,[219] namely (i) it would disincentivise investment in new data centres in the United Kingdom; and (ii) there would be a range of practical difficulties around offering data centre space to competitors on fair, reasonable, and non-discriminatory terms. Furthermore, it would be impossible for cloud services providers to control security if they were forced to give up data centre space to (potentially multiple) other cloud services providers. Cloud services providers' costs would also increase, as they would need to re-design their data centres to accommodate the specific requirements of other cloud services providers (*e.g.*, in terms of cooling and energy use), as each cloud services provider may have different hardware.

97.   In view of the above, AWS welcomes the CMA's decision not to prioritise these potential remedies in its investigation.[220]

      v.   **Potential remedies requiring cloud services providers to be more transparent about the interoperability of their cloud services**

98.   Another potential remedy proposed in the TB Working Paper is to require cloud services providers to (i) publish documentation on the interoperability of each cloud service and how customers would migrate away from the cloud service; and (ii) publish information on forthcoming major changes to their underlying cloud services.[221]

99.   Requiring cloud services providers to publish documentation on the interoperability of each of their cloud services and how customers can migrate away from these cloud services is unnecessary because:

      a.   As explained in paragraph 53 above, AWS already takes many active steps to inform and educate its customers about the technical aspects of its services and the

---

[218]   TB Working Paper, paragraphs 9.101-9.104.

[219]   TB Working Paper, paragraphs 9.111-9.112.

[220]   TB Working Paper, paragraph 9.92.

[221]   TB Working Paper, paragraphs 9.113-9.140.

integration of those services with services from other cloud services providers. For example, AWS explains the programming language behind various tools that can be used to build on AWS,[222] and documents the changes to the underlying open source of its managed open-source services. Every AWS service has extensive documentation explaining exactly how it works. Further, AWS publishes blog posts and case studies to educate customers about how to work with these services and how to migrate various workloads and application types to other clouds. AWS also works with customers to build the solutions that meet their needs.

b.   That said, it is very difficult to establish a comprehensive list of all the information that is pertinent to every customer and every cloud service. Each customer, cloud service, and cloud service combination will have different requirements and sensitivities, and it would be nearly impossible to predict and provide answers for every customer question.

c.   The TB Working Paper itself acknowledges that (i) the effect of a requirement to publish details on the interoperability of cloud services could increase over time as cloud services become more numerous and complex, so that it might become too onerous and costly for cloud services providers to comply; and (ii) if too much information is published or the information published is too detailed or technically complex, some (smaller) customers may not have the resources needed to properly assess or act on the information.[223]

100.   Requiring cloud services providers to give notice and publish details of upcoming material updates to their cloud services is unnecessary, practically unfeasible, and could harm customers, competition, and innovation because:

a.   As explained in [✂], AWS already publishes details of technical updates to its cloud services in different ways. For example:

- AWS publishes details of technical updates on the "What's New with AWS" page of its website.[224]

- AWS's service teams regularly update their technical documentation with new features on the documentation page of AWS's website.[225]

- AWS announces new features and technical updates at AWS re:Invent, the annual flagship conference hosted by AWS for the global cloud computing community.[226]

---

[222]   See: https://aws.amazon.com/developer/tools/?nc1=f_dr.

[223]   TB Working Paper, paragraphs 9.123-9.124.

[224]   See: https://aws.amazon.com/new/?whats-new-content-all.sort-by=item.additionalFields.postDateTime&whats-new-content-all.sort-order=desc&awsf.whats-new-categories=*all.

[225]   See: https://docs.aws.amazon.com/.

[226]   See: https://reinvent.awsevents.com/.

b. That said, it is not possible for AWS to publish information about all its material updates to its cloud services before they are launched. Given the dynamic and competitive nature of the IT industry, publishing information about a potential service before AWS has had an opportunity to launch would allow competitors to freeride on AWS's innovative efforts. Additionally, IT services are complex and new services or features take time to develop and can lead to unexpected issues. Publishing information on services or features before AWS has been able to confirm that they will launch might risk AWS sharing incorrect information or misleading customers as to the eventual capabilities or availability of a new service or feature.

c. The TB Working Paper itself acknowledges that this potential remedy (i) may cause customers to be delayed in receiving the benefits associated with material updates because cloud services providers would be delayed in processing these updates; and (ii) could increase the regulatory burden on cloud services providers and may increase their costs.[227]

vi. Potential remedies to improve skills

101. The final potential remedy proposed in the TB Working Paper is to improve skills of cloud engineers and IT staff in the United Kingdom. This potential remedy would seek to increase the ability of technical staff to work across multiple clouds by requiring a portion of any training provided on cloud services to be cloud-agnostic.[228]

102. Any requirement to make training and education courses cloud-agnostic is unnecessary because:

a. AWS already devotes significant resources to making training available to its customers' technical staff. AWS delivers training and certification opportunities, workshops, and events to educate its customers on how to utilise AWS's services in the most cost-effective way. For example:

- As explained in AWS's Public Response,[229] many of AWS's customers use AWS's free offerings to increase their workforce's overall cloud-based skills. For example, the AWS training and certification team built a training initiative to teach foundational cloud knowledge to Principal Financial Group's workforce, helping to upskill 1,650 employees.[230]

- AWS re:Invent is a learning conference for the global cloud computing community, with keynote announcements, training and certification opportunities, and access to many technical sessions. At the last AWS re:Invent, there were [✕] sessions on the topic of multi-clouding, several of which were co-presentations with AWS partners and customers.

---

227   TB Working Paper, paragraphs 9.133-9.134.

228   TB Working Paper, paragraphs 9.141-9.151.

229   For further details, please see AWS's Public Response, paragraph 66(f). See also [✕].

230   See: https://aws.amazon.com/solutions/case-studies/principal-financial-group-case-study/.

- AWS's Migration Acceleration Program offers tailored training and content, including expertise from AWS Professional Services,[231] to help customers achieve their desired migration outcomes.[232]

    b. Technology-agnostic training would in practice not add much value for customers. While AWS offers technology-agnostic training at a basic level, the value of training often relates to teaching participants about a specific service, application, or technology, for example, serverless systems or HTTP API. The TB Working Paper itself acknowledges that this potential remedy may (i) make any training or education course less useful, as some of the training may not be specific to an individual's role or circumstances; and (ii) limit the ability of individuals to specialise in one area, as any training would likely be more general.[233]

103. In any case, the training of customers' technical staff does not prevent customers from switching or multi-clouding, because (i) customers invest in IT training when they are switching to a new IT services provider; and (ii) AWS and other IT services providers offer a variety of training courses to help customers in their transition to a new IT services provider.[234]

B. Any potential remedies should be applied to all IT services providers

104. In respect of the different potential remedies, the TB Working Paper puts forward that they could either be applied to a subset of cloud services providers (*e.g.*, large cloud services providers) or to all cloud services providers.[235] AWS respectfully submits that, if the CMA were to impose any of the potential remedies considered in the TB Working Paper, these remedies should be applied to all IT services providers rather than targeting only some cloud services providers. Applying the potential remedies selectively risks creating an artificial market fragmentation that could have harmful unintended consequences for competition and innovation.

105. As explained in AWS's Public Response,[236] given that new entrants are entering the market for the supply of IT services on a regular basis and IT services providers are constantly developing new features and services for customers, it is not clear which IT services provider will gain customers in the future and therefore it would be unfair and illogical to only impose restrictions on some IT services providers but not others.

VI. CONCLUSION

106. The evidence presented in the TB Working Paper concerning the technical costs associated with switching and multi-clouding is too mixed to raise any competition concerns and does not support the need for regulatory intervention. While customers may have to incur some

---

[231] Further information on AWS's Professional Services offers and how these help customers migrate their workloads to AWS is available at: https://aws.amazon.com/professional-services/.

[232] See: https://aws.amazon.com/migration-acceleration-program/.

[233] TB Working Paper, paragraphs 9.147-9.148.

[234] See also paragraph 50.a above.

[235] TB Working Paper, paragraphs 9.31, 9.57, 9.75, 9.108, 9.121, and 9.131.

[236] AWS's Public Response, paragraph 70.

technical costs or do some extra work to enjoy the benefits of multi-clouding and switching, the TB Working Paper has not presented any evidence of cloud services providers imposing artificial technical barriers that harm competition. To the contrary, the technical barriers identified in the TB Working Paper are inherent to IT services and often reflective of a healthy level of competition, as they indicate high levels of innovation, service differentiation, and customer choice.

107. The TB Working Paper has also not established that the existing inherent technical barriers prevent customers from switching or multi-clouding in any way that is capable of harming competition. Customers can, *and do*, switch and multi-cloud because of the effective support offered by AWS and other IT services providers, who are heavily commercially incentivised to ensure that customers have *the ability* to multi-cloud and switch, if and when they wish to do so. In fact, the introduction of cloud services by AWS and cloud services providers' continued efforts to support interoperability have made switching and multi-clouding easier than ever before.

108. In such circumstances, there cannot be any AEC. Unsurprisingly, the TB Working Paper does not present any evidence of AEC in relation to technical barriers. Therefore, the potential remedies proposed in the TB Working Paper are unnecessary and unwarranted. Indeed, the consideration of potential remedies is always contingent on an AEC finding having been reached, which is not possible based on the mixed, anecdotal, and contradictory evidence presented in the TB Working Paper. Moreover, the potential remedies considered in the TB Working Paper would come at a disproportionally high price, as they would severely harm innovation and customer choice.

* * *