# Appendix M: Egress fees – analysis of customers' cost scenarios

M.1 This appendix describes our switching and multi-cloud scenarios found in chapter 6, including background, detailed methodology and in-depth results.[1]

M.2 We have calculated these scenarios because we have not been able to rely on real-world data on egress fees for switching and multi-cloud. This is because data on egress fee spending is an imperfect proxy for the actual level of egress fees for switching and multi-cloud. The actual level of egress fees may be low if customers perceive egress fees as high and therefore avoid switching or using multiple clouds. This means that an analysis of the actual spend on egress fees for switching and multi-cloud is likely to understate what proportion of spend they could make up.

M.3 The rest of this appendix is structured as follows:

(a) motivation for this analysis;

(b) overview of the data used;

(c) switching analysis; and

(d) multi-cloud analysis.

M.4 We received submissions on this analysis from third parties, and we have incorporated their feedback into this exercise where relevant.

## Motivation

M.5 Egress fees may affect customers' incentives to switch or to run a multi-cloud architecture. To the extent egress fees for switching and multi-cloud are large, we may expect them to have an impact on customers when deciding whether to switch or to use multiple cloud providers.

M.6 Simply looking at the actual spend on egress fees for switching and multi-cloud is not sufficient to assess the likely impact of egress fees on customer behaviour, because low expenditure on egress fees may be explained by either:

[1] We previously described this as an analysis of 'hypothetical' switching and multi-cloud costs. We have removed the word 'hypothetical' to make clear that we have used actual customer data as opposed to hypothetical cost data.

(a) egress fees being low and spend therefore also being low; or

(b) egress fees being high enough that customers are deterred from incurring them by not switching or by avoiding a multi-cloud architecture.

M.7 To address the question of whether egress fees may be high enough to influence customer behaviour, we have developed some scenarios based on real-world data to explore how much a customer would have to spend, as a proportion of its total cloud spend, if it were to switch or multi-cloud.

M.8 One challenge that we faced in relation to multi-cloud is that there is no single 'use case' for multi-cloud. When considering the egress fee cost of undertaking a full switch, there is a single 'main scenario' we can look at that is informative of the overall relevance of egress fees to customers' switching decisions: the transfer of 100% of the customer's data on a single occasion.

M.9 By contrast, there is no single scenario that can characterise all multi-cloud architectures. A multi-cloud architecture, in theory, involves the transfer of any proportion of the customer's stored data at any frequency (eg hourly, daily, weekly, monthly and so on).

M.10 These scenarios are intended to be purely illustrative and therefore do not necessarily account for the boundaries of what might be practical. These scenarios are designed to illustrate whether egress fees have the potential to disincentivise switching and multi-cloud.

## Data

M.11 For both our switching and multi-cloud scenarios, we have used real-world data on the specific characteristics of UK customers of Microsoft, AWS and Google. In particular:

(a) we collected UK customer-level data on the mean volumes of data they stored for each year from 2018 to 2023;[2]

(b) we collected data on customers' annual spend; and

(c) we used the list prices on cloud providers' websites for egress fees from London via the internet.[3] We refer to such egress as 'standard tier egress'.

[2] We removed customers with zero data stored, as these scenarios would not be relevant for these customers. They represent a small percentage of all customers.
[3] Google: All networking pricing, Microsoft: Pricing – Bandwidth, AWS: Amazon S3 Simple Storage Service Pricing and EC2 On-Demand Instance Pricing, all accessed 29 Jan 2024.

We note that the list prices for Microsoft's, AWS' and Google's standard tier egress have largely not changed in the past 5 years.[4]

M.12 Additionally, for the multi-cloud scenarios, we computed customers' monthly spend by dividing their annual spend by 12.

## Switching scenarios

M.13 We calculated the 'one-off' switching cost for a customer based on a scenario where they transferred all of their data out of their current cloud and the percentage of their annual spend that cost would represent. Using the data we collected:

(a) we multiplied each customer's mean data storage volume by the list price per GB for egress that would apply were they to transfer all of this data at once;

(b) we then calculated the percentage of annual spend this would make up if spread over the duration of a typical contract.

(i) this is because switching costs can be spread over multiple years. A customer may incur switching costs once but incur the benefits of switching over multiple years. We understand that the typical contract length is three years, and therefore, we have assumed that a customer will pay one-third of the cost in a given year. We describe the estimates produced in this way as 'amortised'.

(ii) however, we consider that some customers may have to take into account the up-front cost in the initial year, rather than exclusively considering the cost of switching spread out over the course of a contract. For example, budget constraints may mean that customers are unable to spread the cost over multiple years. We have therefore also calculated the egress fees which would be faced by customers as a percentage of annual cloud spend, if they do not spread the cost over multiple years. We describe the estimates produced in this way as 'non-amortised'.

M.14 We make the following observations to our analysis of 'one-off' switching cost:

(a) our analysis is likely to understate the effect of egress fees on switching, because we have calculated these fees as a proportion of total cloud spend, but egress fees are relevant only for data-intensive workloads. By including unrelated spend in the denominator, we are likely to underestimate

[4] See our analysis of Microsoft, AWS and Google's list prices for egress, in chapter 6 of this report.

(potentially by a significant degree)[5] the effect of egress fees on switching for these data-intensive workloads.

(b) our estimates are based on a 'one-off' transfer of all the data stored on one cloud provider. Our analysis tends to understate the extent to which egress fees are paid in relation to switching, as the proportion of egress fees paid as a percentage of annual cost of cloud may be higher for customers who would need to run an integrated multi-cloud architecture for the duration of their switch.

(c) our analysis is based on list prices for the most standard or basic egress data transfer service for each cloud provider that we looked at, which may reflect egress that is primarily or partly via the public internet. Our analysis does not account for all the products or use cases for customers paying egress fees.

  (i) Microsoft and Google offer tiered routing options for their egress data transfers, which we refer to as 'standard' and 'premium',[6] whereas AWS does not offer this tiered routing. We have taken a more conservative approach by using 'standard' egress routing prices, rather than more expensive 'premium' egress routing prices which are available from Microsoft and Google.[7]

  (ii) while we also recognise that each cloud provider does offer various cost-mitigating alternatives (eg direct connect services), our understanding is that these are often for specialised use cases (eg customers with very large data transfer requirements).

(d) Our analysis does not account for discounts. We have not found a robust methodology to consistently apply discounts to egress fees (we assess [✂] proposed methodology below). Overall, our methodology does not account for discounts, including cross-service discounts like CSAs, which are common mostly among the largest customers.[8] Accounting for discounts in our analysis would lower the estimated egress fees customers who receive discounts would have to pay when switching.

M.15 Overall, we consider these factors have opposing effects on our estimates. This means the combined effect of all of these different factors is ambiguous. As such, we acknowledge our analysis does not account for some factors that may

[5] Analysis of the distribution of Azure cloud use by spend on each service (presented in chapter 7) indicates that spending on storage (the most obvious data related service) accounts for [20-30]% of total Azure spend.
[6] Microsoft's and Google's 'standard' tier involves routing primarily over the public internet, whereas 'premium' involves routing primarily over each cloud provider's own global network. Routing preference in Azure - Azure Virtual Network | Microsoft Learn; Network pricing | Google Cloud.
[7] We note also that it is the 'standard' tier routing option that is available under Microsoft's and Google's free switching programmes.
[8] For our analysis of the prevalence of committed spend discounts (CSDs), see chapter 8, paragraphs 8.41-8.47.

influence customers' actual cost of switching and that for some customers the cost of multi-cloud will likely be either higher or lower.

M.16 Oracle submitted that, in focusing on egress fees as a percentage of customer spend, we underestimate the importance of the absolute value of egress fees to the customer by failing to account for 'sticker shock' at the highest levels. Oracle explained that '[0-5%]' of spend sounds much smaller than $10 million (for example) just to switch cloud providers'.[9] Nevertheless, we have maintained our decision to focus on egress fees as a percentage of customer spend as these are more representative of the impact of egress fees on customers. For instance, the impact of $10 million to a small customer is not the same as the impact of $10 million to a large customer.

M.17 Microsoft and AWS made submissions related to the fact that our analysis of one-off switching costs is based on list prices and does not account for discounted rates:

(a) AWS submitted that [✂], and added that, since 2020, approximately [✂] of bytes transferred out of AWS are subject to [✂].[10]

(b) Microsoft submitted that customers often negotiate discounts on egress fees, and that there is a strong incentive for Microsoft to negotiate such discounts on egress fees to maintain long-term / future customer relationships.[11]

M.18 AWS submitted an analysis which contained updated estimates for the average share of one-off switching costs as a percentage of total customer spend on AWS services, including a methodology to account for customer-level discounts.[12] We consider that this analysis has a number of limitations, including:

(a) AWS uses a tiered pricing schedule for egress fees, which includes a free tier. This means that each GB of data egressed does not increase the total egress cost by the same amount. We consider that this means that the egress fees faced by customers will depend on how much data they are egressing. This is not accounted for in this analysis which limits our ability to use it to draw conclusions about the egress fees faced by customers who egress different volumes.

(b) it relies on data which only captures customers who considered egress fees low enough that they do not disincentivise data transfers. This means that it

[9] Oracle's response to the Updated Issues Statement and working papers, page 3.
[10] AWS' submission to the CMA [✂].
[11] Microsoft's response to the CMA's PDR dated 28 January 2025, footnote 107. Microsoft also noted that the Microsoft Azure website advices customers should "contact for prices over 500TB".
[12] AWS submitted that these estimates accounted for contract length, customer-level discounts (by computing average selling prices (ASP) for egress fees paid by UK customers and competition at the workload level (using a weighted average of switching costs by customer spend). [✂]. We consider accounting for contract length, and the use of averages weighted by customer spend, elsewhere in this appendix.

does not capture cases where egress fees are higher than the expected benefits of transferring data, including switching and multi-cloud, and therefore act as a barrier to this.

M.19 AWS submitted that a transfer of 100% of stored data does not necessarily reflect real-world switching scenarios.[13] Whilst we note that not every switching strategy involves transferring 100% of stored data in a single year, AWS' own free switching programme stipulates that to be eligible for the programme, customers must not leave any workloads or data on their AWS account and customers have 60 days to complete their move off of AWS,[14] which effectively requires 100% of stored data to be switched well within a year. We also note that the switching process for a customer can involve periods of multi-cloud running where customers are likely to transfer some of their data more than once.[15] We therefore consider this approach appropriate.

M.20 Microsoft submitted that cloud providers' free switching programmes eliminated all egress fees where customers transfer all of their data from one cloud provider to another,[16] and similarly, AWS submitted that it has eliminated DTO fees globally for customers switching away from AWS.[17] We have assessed these programmes in chapter 6 and in appendix N of this report, and consider that these programmes do not materially affect the conclusions of our analysis of either switching costs or multi-cloud costs.[18]

M.21 Microsoft submitted that the realities of switching and multi-cloud are more complex than indicated by this simplistic analysis.[19] In interpreting the results of this analysis, we consider egress fees in the context of switching being a complex and lengthy decision, where egress fees are one barrier among others. We do not suggest they are the only consideration made by customers when considering switching cloud providers.

**Average cost of switching, by year**

M.22 Figure M.1, Figure M.2 and Figure M.3 illustrate the average (mean) cost to a customer for 'one-off' switching as a percentage of its annual cloud spend for customers of Microsoft, AWS and Google, using data from different years to assess trends over time. We present these figures both where the cost is spread over the course of a typical contract, ie three years (amortised) and where the cost is not spread over multiple years (non-amortised).

---

[13] AWS' submission to the CMA [✂].
[14] Amazon EC2 FAQs – AWS, accessed 2 September 2024.
[15] See the description of the coverage of the free switching programmes, in Appendix N.
[16] Microsoft's response to the CMA's PDR dated 28 January 2025, footnote 107.
[17] AWS' response to the CMA's PDR dated 28 January 2025, paragraph 54.
[18] See our analysis of the relevance of free switching programmes to UK customers, in chapter 6.
[19] Microsoft's response to the CMA's PDR dated 28 January 2025, paragraph 100.

**Figure M.1: [✂]**

[✂]

*Source: CMA analysis of [✂] data [✂].*

**Figure M.2: [✂]**

[✂]

*Source: CMA analysis of [✂] data [✂].*

**Figure M.3: [✂]**

[✂]

*Source: CMA analysis of [✂] data [✂].*

M.23 Overall, average egress fees paid to move all of a UK customer's data once from one cloud provider to another would range between [0-5%] to [0-5%] of a customer's annual spend per year, if the cost is spread over the course of a typical contract (ie three years). Specifically:

(a) Figure M.1 shows that the average UK customer of AWS would incur costs of [0-5]% of their annual costs per year for three years if they were to perform a 'one-off' switch away from AWS.

(b) Figure M.2 shows that, for the average UK customer of Microsoft, 'one-off' switching would represent [0-5]% of their average annual costs per year for three years. This percentage fluctuated over the five-year period between 2018 and 2022 and was generally on an upwards trend until 2022.

(c) Figure M.3 shows that, for the average UK customer of Google, 'one-off' switching would represent [0-5]% of their average annual costs per year for three years. This has stayed consistent over the three-year period we have examined.

M.24 If customers do not spread the cost over multiple years (non-amortised), these results suggest that a customer might expect to pay egress fees of [0-5]% to [5-10]% of annual spend in a single year.

M.25 Microsoft submitted that this analysis used unweighted averages of the switching cost share of total spend, and that its calculations of a weighted average (ie shares are weighted by overall revenues) using the same methodology and data shows that one-off switching cost is less than [✂] across all years (and would be even lower where such cost is amortised across three years).[20] Similarly, AWS submitted that, to account for the fact that cloud service providers compete for

[20] [✂].

spend and workloads as opposed to individual customers, we should compute a weighted average of switching costs by customer spend.[21] We note that a weighted average by spend would be a lower value than the averages presented above, for these parties.

M.26 We consider that an unweighted average is the more appropriate methodology. This is because implementing a weighted average by customer spend would give more weight to larger customers than smaller customers. In this particular analysis, we are considering the one-off switching cost that may be faced by customers of differing sizes.

M.27 However, we acknowledge that it is relevant to consider how customers with different amounts of cloud spend are affected by egress fees. We consider that our analysis of the average costs of switching by spend band (as presented below) more directly enables us to understand how customers of different sizes are affected by egress fees for switching.

### Average cost of switching, by spend band

M.28 We have also considered the average (mean) cost to a customer for 'one-off' switching as a percentage of its annual cloud spend, for customers in different bands of annual cloud spend. Figure M.4, Figure M.5 and Figure M.6 present results split by spend band, for customers of Microsoft, AWS and Google, using data from 2022. We present these figures both where the cost is spread over the course of a typical contract, ie three years (amortised) and where the cost is not spread over multiple years (non-amortised).

**Figure M.4: [✂]**

[✂]

*Source: CMA analysis of [✂] data [✂].*

**Figure M.5: [✂]**

[✂]

*Source: CMA analysis of [✂] data [✂].*

**Figure M.6: [✂]**

[✂]

[21] [✂].

*Source: CMA analysis of [✂] data [✂].*

M.29 This analysis shows that AWS and Microsoft customers in the lowest annual spend band (ie the smallest customers) face disproportionately high egress fees as a percentage of annual cloud spend, compared to other customers. Specifically:

(a) Figure M.4 shows that the smallest AWS customers would have had to pay egress fees of [0-5]% per year over the course of a typical contract to switch cloud providers in 2022, compared to an average of [0-5%] across all AWS customers.

(b) Figure M.5 shows that the smallest Microsoft customers would have had to pay egress fees of [0-5%] per year over the course of a typical contract switch cloud providers in 2022, compared to an average of [0-5%] across all Microsoft customers.

M.30 The smallest category of Google customers also face high egress fees, as do customers in the two largest spend bands. This may be influenced by the small number of observations in these spend bands in 2022. Specifically:

(a) Figure M.6 shows that the smallest Google customers would have had to pay egress fees of [✂] per year over the course of a typical contract to switch cloud providers in 2022, above the average of [0-5]% for all Google customers. The largest customers would have had to pay egress fees of [0-5]% per year over the course of a typical contract.

M.31 Overall, this analysis shows that, across cloud providers, smallest customers face higher egress fees than other customers as a proportion of their annual cloud spend in order to switch cloud providers. The average egress fees faced by customers with annual spends of £1,000-£10,000 is [0-5]% which is, between [0-5]% percentage points higher per year than the averages of higher spend bands.

**Distribution of "one off" switching costs**

M.32 We have considered how the magnitude of "one-off" switching costs are distributed across customers, both where the cost is spread over the course of a typical contract, ie three years (amortised), and where the cost is not spread over multiple years (non-amortised), based on data from 2022.

M.33 Table M.1 presents the distribution of one-off switching costs as a proportion of annual cloud spend, if customers were to spread their costs over the course of a typical contract, ie three years (amortised).

**Table M.1: Distribution of amortised one-off switching cost percentages, 2022**

| *Distribution of amortised one-off switching cost percentages as a percentage of annual cloud spend, 2022* | *AWS* | *Microsoft* | *Google* |
|---|---|---|---|
| 0-5% | [90-100]% | [90-100]% | [90-100]% |
| 5-10% | [0-5]% | [0-5]% | [0-5]% |
| 10-15% | [0-5]% | [0-5]% | [0-5]% |
| 15-25% | [0-5]% | [0-5]% | [0-5]% |
| 25-50% | [0-5]% | [0-5]% | [0-5]% |
| 50-75% | [0-5]% | [0-5]% | [0-5]% |
| 75-100% | [0-5]% | [0-5]% | [0-5]% |
| Over 100% | [0-5]% | [0-5]% | [0-5]% |

*Source: CMA analysis of AWS, Microsoft and Google data. Responses to the CMA's information requests issued under section 174 notice to [✂].*

M.34 Overall, our analysis shows that some customers face meaningful egress fees to switch their cloud provider.

(a) [0-5]% to [5-10]% of customers would have had to pay egress fees of more than [5-10]% of their annual cloud spend per year over the course of a typical contract.

(b) These costs can also be materially higher for some customers. For example, between [0-5]% of customers would have had to pay fees over [10-20]% per year.

M.35 This means that [0-5]% to [5-10]% of customers would not have the incentive to switch to a rival cloud provider even if, over the course of a typical contract, that provider offered a price that was [5-10]% better or equivalently better quality or functionality.

M.36 Table M.2 presents the distribution of one-off switching costs as a proportion of annual cloud spend, where the customer does not spread the cost over multiple years (non-amortised).

**Table M.2: Distribution of non-amortised one-off switching cost percentages, 2022**

| *Distribution of non-amortised one-off switching cost percentages as a percentage of annual cloud spend, 2022* | *AWS* | *Microsoft* | *Google* |
|---|---|---|---|
| 0-5% | [80-90]% | [80-90]% | [70-80]% |
| 5-10% | [5-10]% | [10-20]% | [5-10]% |
| 10-15% | [0-5]% | [0-5]% | [0-5]% |
| 15-25% | [0-5]% | [0-5]% | [0-5]% |
| 25-50% | [0-5]% | [0-5]% | [0-5]% |
| 50-75% | [0-5]% | [0-5]% | [0-5]% |
| 75-100% | [0-5]% | [0-5]% | [0-5]% |
| Over 100% | [0-5]% | [0-5]% | [0-5]% |

*Source: CMA analysis of AWS, Microsoft and Google data. Responses to the CMA's information requests issued under section 174 notice to [✂].*

M.37 This shows that if customers did not spread the cost over multiple years, around [10-20]% to [20-30]% of customers would have had to pay egress fees of more than [5-10]% of their annual cloud spend in a year. These costs can also be materially higher for some customers. For example, around [0-5]% to [5-10]% of

customers would have had to pay egress fees of more than [15-30]% of their annual cloud spend in a year.

## Multi-cloud scenarios

M.38 In order to illustrate the potential cost of a range of possible multi-cloud scenarios, we have calculated the monthly egress fees a customer would pay (as a proportion of their overall cloud spend) depending on:

(a) the customer's size (proxied by their annual spend on cloud services – this is based on real-world data); and

(b) the proportion of their stored data that they may transfer.

M.39 As egress fees are billed monthly, we have analysed these costs on a monthly basis. All data transfer volumes, total egress fee costs and the percentages of cost that egress fees may make up therefore reflect one individual month.

M.40 Using the data outlined in paragraph M.11, we calculated the percentage of total monthly spend that egress fees could make up in various data transfer scenarios across six different annual spend bands. To illustrate this, we created a table, with spend bands on the rows and data transfer volumes on the columns.

M.41 We chose to conduct this analysis by customer size rather than just volume of data stored in order to capture the non-linear pricing of egress fees and its subsequent effects. In addition, using spend bands already used elsewhere in our investigation allowed us to map the results of this analysis across to other sections, where relevant.[22]

### Methodology to calculate multi-cloud scenarios

M.42 Our methodology is described step-by-step below:

(a) first, we took the average (using both the median and mean) of each spend band's mean volume of data stored in the cloud. In doing so, we excluded the top and bottom 2% of data storage values to exclude potential outliers.[23]

(b) second, data transfer volumes, expressed as a percentage of data stored in the cloud, were applied to the average data stored for each spend band. This gave us values in gigabytes (GB) for amounts of data transferred per month in our multi-cloud architectures.

[22] These spend bands were originally used by Ofcom. See Ofcom's Final Report, Table 5.7
[23] Where spend bands have fewer than 50 observations, we remove the maximum and minimum value for data storage only.

(c) third, we multiplied these data transfer volumes by the applicable per-GB price for data egress (accounting for any free tiers) based on list prices taken from cloud providers' websites to give a monthly cost (USD) of egress fees. These prices are applied as tiered (eg the first 100 or 200GB is free, then the following 100TB is split into 50TB at $0.05, 50TB at $0.04) rather than charged at a uniform rate (100TB all at $0.04). test

(d) fourth, we took the average (using both the median and mean) of total annual cloud spend for customers within each spend band. This was divided by 12 to give average monthly spend for each spend band.

(e) finally, we divided the costs computed in step (c) by the average monthly cloud spend computed in step (d) to obtain the percentage of monthly spend that egress fees would make up in a multi-cloud scenario.

M.43 As set out in (a) and (d) above, we have calculated averages within each spend band, for volume of data stored and total cloud spend, using both the mean and median, which we present separately below.

**Caveats to our analysis of multi-cloud scenarios**

M.44 When interpreting the results of our analysis, the same caveats as our switching analysis apply, plus:

(a) some spend bands have very few observations because not many customers fell within that spend bracket in 2022. Therefore, it may be the case that these numbers are not necessarily representative of other customers who would be in those spend bands in other years.

(b) the purpose of this analysis is to give an indication of the broad order of influence of fees that would be incurred for different levels of transfer. Where this analysis suggests that egress fees would be expensive for a given level of (ongoing) data transfer, whether that level of data transfer would be possible due to barriers unrelated to egress fees is a separate question.

(c) our scenarios are simplifications of what a multi-cloud architecture might look like, focusing only on data transfer volumes per month. Other considerations (eg data sovereignty requirements, different methods of data transfer, different types of data storage) are not captured in our analysis.

(d) our analysis only assesses egress fees for data transfers between clouds, as this is the only type of egress relevant to whether egress fees are a barrier to multi-cloud. Hence, we do not consider any other type of egress fees in our analysis (eg egress to end users, or 'serving' egress).

M.45 We received a number of submissions that our analysis is based on list prices for data egress and does not take account of cost-mitigating strategies and discounts.

(a) Microsoft submitted that the egress products considered in this analysis does not capture cheaper (per GB) alternatives that may be available to their customers.[24]

(b) Microsoft submitted that customers often negotiate discounts on egress fees, and that there is a strong incentive for Microsoft to negotiate such discounts on egress fees to maintain long-term / future customer relationships.[25]

(c) a cloud provider submitted that our analysis fails to account for the different types of egress (and their related costs) customers can adopt.[26]

(d) AWS submitted that these tables may overestimate the cost of data transfer, and in particular, the tables do not account for volume discounts, private pricing, or cost-mitigating strategies available to customers when calculating the cost of data transfers. Further, the total cloud spend corresponds to net spend and will account for these factors, leading to an upward bias.[27]

(e) AWS submitted an analysis which accounted for discounts by computing customer-specific average egress prices paid by UK customers (ie the total egress revenues received by AWS from a given customer, divided by that customer's egress usage for each customer).[28]

M.46 We note that these submissions overlap with those we received on our switching analysis and therefore our assessment of them above also applies here.[29] Overall, we consider these factors have opposing effects on our estimates. This means the combined effect of all of these different factors is ambiguous. As such, we acknowledge our analysis does not account for some factors that may influence customers' actual cost of multi-cloud, and that for some customers, the cost of multi-cloud will likely be either higher or lower.

M.47 AWS, Microsoft and Google all submitted that our data transfer volumes are not reflective of customers' current architectures:

(a) AWS and Google submitted that our assumption that customers may need to transfer up to 100% of their data stored to multi-cloud is not justified by any consideration of how customers actually multi-cloud.[30]

---

[24] Microsoft's response to the CMA's PDR dated 28 January 2025, footnote 107.
[25] Microsoft's response to the CMA's PDR dated 28 January 2025, footnote 107. Microsoft also noted that the Microsoft Azure website advices customers should "contact for prices over 500TB".
[26] [✂] submission to the CMA [✂].
[27] AWS' submission to the CMA [✂].
[28] [✂].
[29] See under 'Switching scenarios' earlier in this appendix.
[30] AWS submission to the CMA [✂] Googles submission to the CMA [✂].

(b) Microsoft submitted that the realities of switching and multi-cloud are more complex than indicated by this simplistic analysis.[31] It submitted that the hypothetical multi-cloud use case set out in the Egress fees working paper, which Microsoft characterised as one that requires significant data to be moving between clouds on a constant basis, does not exist in practice. Microsoft added that the examples of potential frictions created by multi-cloud provided are not realistic and mainstream.[32]

(c) Microsoft further submitted that this analysis is unrealistic and overestimates the volume of monthly data egressed by UK Azure customers. It submitted that the average monthly data egressed by UK Azure companies was [✂] in 2022, which is c. [✂]% of the average monthly data started summed across all UK customers in 2022 and therefore our analysis implies an egress spend share of less than [✂]% across all UK Azure customers in 2022. It added that even if the total volume of data egress doubled (ie customers transfer [✂]% of the monthly volume of data stored), the analysis would still show that egress costs were still less than [✂]% of total spend across all UK customers.[33]

(d) a cloud provider submitted that, in its own experience, customers who multi-cloud typically transfer less than 10% of their total stored data on a monthly basis.[34] It further submitted that customers who multi-cloud have strong incentives – unrelated to egress fees, such as compliance, security and data localisation requirements – to keep the volume of data transferred for the purposes of multi-cloud to the minimum amount required.[35]

M.48 We consider that the current level of integration (and therefore the level of data transferred across clouds) in multi-cloud architectures is likely to be influenced by the current level of egress fees. This is because the higher egress fees are, the more expensive it is to multi-cloud. Therefore, the current level of data transferred across cloud does not provide a reliable indication of how much data would be transferred in the absence of such fees.

**Results and interpretation**

M.49 In the tables below, the percentages along the top row represent different volumes of data transfer as a percentage of a customer's total data stored in the cloud. This is the easiest way to capture a wide range on multi-cloud architectures (ie different degrees of integration). The spend bands in the column on the left are categories

[31] Microsoft's response to the CMA's PDR dated 28 January 2025, paragraph 100.
[32] Microsoft's response to the Competitive landscape, Committed spend agreements and Egress fees working papers, page 25.
[33] Microsoft's submission to the CMA [✂].
[34] [✂] submission to the CMA [✂].
[35] [✂] submission to the CMA [✂].

based on how much customers spend on cloud services per year (USD) and are a proxy for different customer sizes. Each cell of the table represents how much it would cost a customer of a certain size to transfer a given share of its stored data to another cloud, as a percentage of their total monthly cloud spend.

M.50 The results give us an indication of the egress fee costs associated with different levels of data transfer in a multi-cloud architecture and whether some sizes of customers face especially high costs to transfer data. By extension, these tables can help to indicate whether egress fees may disincentivise moving to multi-cloud architectures that require a certain volume of data transfer (ie increasing the level of integration of clouds).

M.51 We produced tables for customers whose primary cloud is Microsoft, AWS and Google respectively who are transferring data to another cloud via the internet. It shows how much different customers might expect to pay as a percentage of their total cloud spend to transfer data from their incumbent cloud to another cloud (or indeed to any other destination over the internet), depending on the customer's size – proxied by annual spend on their cloud – and how much data (as a share of their total stored data) they intend to transfer.

M.52 This analysis considers the costs of different multi-cloud strategies for a representative customer of each spend band, where this customer has a given volume of cloud storage and annual cloud spend. As set out above, we have calculated these costs based on the average volume of cloud storage and average annual cloud spend, using both the mean and median within each spend band.

**Median customer analysis**

M.53 Figure M.7, Figure M.8 and Figure M.9 below show the cost of egress as a percentage of monthly cloud spend, depending on the proportion their total stored data transferred, for customers of AWS, Microsoft and Google respectively. We calculated these cost percentages for each spend band, based on the median of customers' volumes of data stored and total annual cloud spend within each band.

**Figure M.7 [✂]**

*Source: CMA analysis of [✂] data [✂].*

**Figure M.8: [✂]**

*Source: CMA analysis of [✂] data [✂].*

**Figure M.9: [✂]**

**[✂]**

*Source: CMA analysis of [✂] data [✂].*

M.54 Looking at an average customer (in terms of volumes of data stored and total annual cloud spend) within each spend band, and across the three main cloud providers, we found that:

(a) for an average customer in the lowest spend band (£1,000 and 10,000), the cost of egressing [20-30]% of their data ranges between [0-5]% of their monthly cloud spend, ranged across cloud providers. This increases to a range of [0-5]% to [5-10]% when egressing 50% of data stored.

(b) for an average customer in spend bands spending between £10,000 and and £20million, the cost of egressing [20-30]% of their data ranges from [0-5]% to [10-20]% of their monthly cloud spend, ranged across spend bands and cloud providers. [36] This increases to a range of [5-10]% to [20-30]% when egressing [50-60]% of data stored. [37]

(c) for an average customer spending more than £20million in 2022, the cost of egressing [20-30]% of data results in a cost of between [0-5]% and [10-20]% of their monthly cloud spend, ranged across cloud providers. [38] This increases to a range of [✂]% to [✂]% when egressing 50% of data stored. [39]

M.55 Some customers are covered by their cloud provider's free tier and therefore pay zero egress fees, up to a certain volume of data transfer. We found that:

(a) an average customer in the lowest spend band (£1,000-10,000) is covered by cloud providers' free tiers when transferring up to [10-20]% to [30-40]% of their data, depending on their cloud provider; and

(b) the average customer in the next-lowest spend band (£10,000-1million) is also covered when transferring up to [✂] of their data, depending on their cloud provider.

---

[36] We note that some spend bands have very few observations because not many customers fell within that spend bracket in 2022. Therefore, as a sensitivity we considered the egress fees faced by customers in spend bands with at least [✂] customers and found that this range was [0-5]% to [5-10]%.
[37] Considering the egress fees faced by customers in spend bands with at least [✂] customers, that this range was [5-10]% to [10-20]%
[38] Considering the egress fees faced by customers in spend bands with at least [✂] customers, that this range was [0-5]% to [5-10]%.
[39] Considering the egress fees faced by customers in spend bands with at least [✂] customers, that this range was [5-10]% to [10-20]%.

### Mean customer analysis

M.56 Figure M.10, Figure M.11 and Figure M.12 below show the cost of egress as a percentage of monthly cloud spend for customers of AWS, Microsoft and Google respectively, depending on the customer's annual cloud spend and the proportion their total stored data they intend to transfer, using the mean of customer's cloud spend and total annual cloud spend in each spend band.

**Figure M.10: [✂]**


[✂]

*Source: CMA analysis of [✂] data [✂].*

**Figure M.11: [✂]**


[✂]

*Source: CMA analysis of [✂] data [✂].*

**Figure M.12: [✂]**


[✂]

*Source: CMA analysis of [✂] data [✂].*

M.57 Looking at an average customer (in terms of volumes of data stored and total annual cloud spend) within each spend band, and across the three main cloud providers, we found that:

(a) for an average customer in the lowest spend band (£1,000-10,000), the cost of egressing [20-30]% of their data ranges between [0-5]% and [5-10]% of their monthly cloud spend, ranged across cloud providers. This increases to a range of [✂]% to [✂]% when egressing 50% of data stored.

(b) for an average customer in spend bands spending between £10,000 and £20million, the cost of egressing 20% of their data ranges from [✂]% to [✂]% of their monthly cloud spend, ranged across spend bands and cloud providers.[40] This increases to a range of [✂]% to [✂]% when egressing 50% of data stored.[41]

(c) for an average customer spending more than £20 million in 2022, the cost of egressing 20% of data results in a cost of between [✂]% and [✂]% of their monthly cloud spend, ranged across cloud providers.[42] This increases to a range of [✂]% to [✂]% when egressing 50% of data stored.[43]

M.58 Some customers are covered by their cloud provider's free tier and therefore pay zero egress fees, up to a certain volume of data transfer. We found that:

(a) an average customer in the lowest spend band (£1,000-10,000) is covered by cloud providers' free tiers when transferring up to [5-10%] to [10-20]% of their data, depending on their cloud provider;

(b) the average customer in the next-lowest spend band (£10,000 and £1million) is not covered when transferring [0-5]% of their data, for any of the three main cloud providers.

**Interpretation**

M.59 The results of our analysis are broadly consistent whether we use the mean or median to estimate the average volume of data and cloud spend. However, we observe that:

(a) the results using the mean tend to be higher than the median. This difference is most significant amongst the smallest customers (those who spend

[40] We note that some spend bands have very few observations because not many customers fell within that spend bracket in 2022. Therefore, as a sensitivity we considered the egress fees faced by customers in spend bands with at least [✂] customers and found that this range was [0-5]% to [5-10]%.
[41] Considering the egress fees faced by customers in spend bands with at least [✂] customers, that this range was [5-10]% to [20-30]%.
[42] Considering the egress fees faced by customers in spend bands with at least [✂] customers, that this range was [0-5]% to [0-5]%.
[43] Considering the egress fees faced by customers in spend bands with at least [✂] customers, that this range was [5-10]% to [10-20]%.

between £1,000-£10,000) when egressing at least 20% [20-30]% of their data per month; and

(b) compared to larger customers, the difference between the median and the mean egress fees faced by the smallest customers increases at a faster rate as the volume of data being transferred increases.

M.60 This suggests that some of the smallest customers face considerable egress fees when transferring at least 20% [20-30]% of their data per month. These customers may have a particularly high disincentive to increase the level of integration between clouds or multi-cloud. This is because the costs they would face to transfer data between clouds at higher data transfer levels, in terms of percentage of their annual spend, are higher than larger customers.

## Google's submission on the costs of multi-cloud

M.61 Google submitted an analysis which sought to quantify the difference between a customer's internal data transfer cost between multiple regions in a single-cloud setup and a customer's cross-cloud cost in a multi-cloud setup.[44] Google submitted that its results show that the cost difference between the two data transfer types is 'de minimis' in the context of overall cloud spend.[45]

M.62 For this analysis, Google created a hypothetical customer, based on a real-world multi-cloud customer who uses an integrated multi-cloud architecture. The hypothetical customer is launching a new digital bank.[46] In this hypothetical customer's multi-cloud setup, the hypothetical customer transfers some amount of data between clouds in order to make use of products elsewhere.[47]

M.63 In this hypothetical customer's single-cloud setup, Google's model assumes that this customer will replicate certain parts of its setup in at least two cloud regions for resiliency and availability purposes as a result of 'the nature and geographic footprint of [the customer's] banking operations'. A large amount of data is replicated and transferred across regions, incurring egress fees across regions. 50% more data is egressed in the single-cloud setup than in the multi-cloud setup.[48]

M.64 After comparing the costs of data egress as a percentage of total workload-related cloud spend across the two different architectures, Google concludes that the

[44] Google's submission to the CMA [✂].
[45] Google's submission to the CMA [✂].
[46] Google's submission to the CMA [✂].
[47] Google's submission to the CMA [✂].
[48] Google's submission to the CMA [✂].

difference in data transit costs is too small to suggest that egress fees are a barrier to multi-cloud.[49]

M.65 We consider that Google's submission sets an artificially high benchmark for finding a concern due to its choice of hypothetical customer and the modelling choices it made in a single-cloud scenario:

(a) first, we do not consider the customer chosen to be representative of the wider market. The customer chosen, a large bank, is likely to be subject to various requirements (eg resiliency/availability thresholds, cross-region data replication, data sovereignty) that other customers will not need to consider. This inflates single-cloud egress costs, as this customer transfers more data between regions than we may otherwise expect.

(b) second, we consider that the assumptions made when modelling the customer's single-cloud architecture artificially inflate egress fees in the single-cloud case, therefore reducing the difference in egress fees between the single-cloud and multi-cloud cases.

M.66 We further consider that internal transfer fees are not a relevant benchmark against which to assess the potential impact of egress fees, as not all customers will incur them, eg if they do not operate across multiple regions.

M.67 As such, we consider Google's analysis does not show that egress fees are not a barrier to multi-cloud.

[49] Google's submission to the CMA [✂].