

Cloud services market investigation

Egress fees working paper

23 May 2024

This is one of a series of consultative working papers which will be published during the course of the investigation. This paper should be read alongside the [Issues Statement](#) published on 17 October 2023 and other working papers published.

These papers do not form the inquiry group's provisional decision report. The group is carrying forward its information-gathering and analysis work and will proceed to prepare its provisional decision report, which is currently scheduled for publication in September/October, taking into consideration responses to the consultation on the Issues Statement and responses to the working papers as well as other submissions made to us. Parties wishing to comment on this paper should send their comments to CloudMI@cma.gov.uk by **13 June 2024**.

© Crown copyright 2024

You may reuse this information (not including logos) free of charge in any format or medium, under the terms of the Open Government Licence.

To view this licence, visit www.nationalarchives.gov.uk/doc/open-government-licence/ or write to the Information Policy Team, The National Archives, Kew, London TW94DU, or email: psi@nationalarchives.gsi.gov.uk.

The Competition and Markets Authority has excluded from this published version of the working paper information which it considers should be excluded having regard to the three considerations set out in section 244 of the Enterprise Act 2002 (specified information: considerations relevant to disclosure). The omissions are indicated by [X]. [Some numbers have been replaced by a range. These are shown in square brackets.] [Non-sensitive wording is also indicated in square brackets]

Contents

1.	Introduction and framework	3
	Conceptual framework and analysis structure	3
	Background on egress fees	5
	Types of interconnections	6
	Pricing	8
	Recent developments	11
2.	Analysis of egress fees	13
	Prevalence of egress fees	13
	Cloud providers' views	13
	Data on egress fee spending	13
	Emerging views on prevalence of egress fees	14
	Magnitude of egress fees and role of egress fees on customers' choice	14
	Customer data on spending on egress fees	15
	Customer views	19
	Internal documents	23
	Emerging views on magnitude of egress fees and role of egress fees on customers' choice	24
3.	Cloud providers' cost of egressing data	26
	Relevant costs for egress data transfers	26
	Common data transfer costs vs specific costs for egress	27
	Incremental costs	28
	Analysis of egress costs	30
	Analytical approach	30
	Our analysis	31
	Emerging views on cost of egress fees to providers	35
4.	Potential remedies	36
	Overview of potential egress fee remedies	36
	Interactions with regulations outside the UK	37
	Our analysis of potential egress fee remedies	37
	Restricting the level of egress fees: price control remedies	38
	Banning egress fees	40
	Capping egress fees by reference to other fees charged	44
	Capping egress fees by comparison to costs incurred	46
	Other considerations	48
	Information transparency remedies	49
	Summary of potential remedies and invitation to comment	50
	Appendix A: Cloud providers' cost of egress	51
	Assets cloud providers use to provide data transfer services	51
	Costs cloud providers incur in providing data transfer services	52
	Cloud provider submissions on egress-specific costs	52
	Cloud provider submissions on incremental costs	53
	Appendix B: Cloud providers' connectivity arrangements	55
	Connectivity costs	56
	Appendix C: Data available for analysis of egress costs	58

1. Introduction and framework

- 1.1 This working paper presents our initial analysis of the potential impact of egress fees on competition in connection with the supply of public cloud infrastructure services in the UK, in particular, the potential impact on switching and multi-clouding.
- 1.2 We set out in turn:
- (a) our conceptual framework for assessing the potential impact of egress fees on competition,
 - (b) some background information on egress fees; and
 - (c) our initial analysis, based on evidence we have seen to date, of:
 - (i) the prevalence of egress fees;
 - (ii) the magnitude of egress fees and role of egress fees on customers' choice; and
 - (iii) the cloud providers' cost of egressing data.
- 1.3 We include our early thinking on potential remedies, together with some specific questions on which we would welcome views. We also welcome views on any other potential remedies.
- 1.4 Parties wishing to comment on this paper should send their comments to CloudMI@cma.gov.uk by 13 June 2024.

Conceptual framework and analysis structure

- 1.5 Egress fees are charges to customers when they transfer data out of their provider's cloud, either to an end user or application, when moving data between the cloud and its on-premises data centres, when moving data between different cloud providers, either as part of a multi-cloud architecture, or as part of switching between two cloud providers.
- 1.6 In our [Issues Statement](#), we set out a theory of harm relating to egress fees. In particular, we indicated that we would assess the extent to which egress fees may prevent, restrict or distort competition by creating barriers to switching and the use of multiple clouds, by contributing to the unpredictability of costs for customers, or by leading cloud providers to entrench their position.¹
- 1.7 Under this theory of harm, we are considering whether, and to what extent, egress fees may make it more costly and difficult for customers to choose the best value offers for them and may be deterred from switching or using multiple clouds.
- 1.8 Customers considering whether to switch or multi-cloud face a trade-off between the expected benefits and expected costs of doing so.
- (a) The expected benefits of switching or using multiple clouds may include lower spend on cloud services, higher flexibility, and the ability to access better-quality and better-fitted services or new innovations.

¹ CMA, [Issues Statement](#), paragraphs 6 and 28.

- (b) Expected costs may include the time and cost of moving workloads, the cost of reconfiguring the cloud architecture, the cost of retraining staff and the increased management complexity for customers using multiple clouds.
- 1.9 Where the expected benefits of switching (or using multiple clouds) exceed the costs, we would expect customers to switch (or multi-cloud). Where the costs exceed the benefits, customers are likely to choose not to use an alternative provider, even though the offer may be better for that customer. The greater the switching cost, the more likely it is that customers will refrain from switching or using multiple clouds despite the availability of products that would otherwise represent a better overall offer in terms of price, quality, range of features/capabilities, etc...
- 1.10 Where switching costs² limit switching (or the use of multiple clouds), the competitive process may be distorted. Customers may be harmed where they are deterred from availing themselves of attractive competitive offerings.
- 1.11 Switching costs may further harm competition in the long run where they make it more difficult for smaller rivals to expand, benefit from economies of scale, and compete with larger rivals on a stronger footing.
- 1.12 Switching costs may therefore reinforce or increase the level of concentration in a market. We are also more likely to be concerned where switching costs are present in markets that are more concentrated, especially when they are imposed by the larger providers in those markets. We consider the extent of economies of scale and the current level of concentration in the supply of cloud infrastructure services in the Competitive landscape working paper.
- 1.13 Egress fees represent a cost both when switching or when using multiple clouds:
- (a) When switching public cloud provider, a customer will need to transfer data between public clouds and will incur a fee if the volume of data transferred is large enough.
 - (b) When using an integrated multi-cloud architecture, a customer will transfer data back and forth between public clouds and will incur ongoing egress fees over time if the volume of data transferred reaches a sufficient level.
- 1.14 The extent of the barrier from egress fees to both switching and using multiple clouds depends on the amount of data a customer seeks to transfer and the frequency with which the data must be transferred.
- (a) In relation to switching, customers could move all the data they seek to transfer as a 'one-off' transfer.³ When transferring as a 'one-off', the impact of egress fees are incurred only during the switching process rather than on a recurring basis.
 - (b) In relation to the use of multiple clouds, the extent to which egress fees may be a barrier to multi-cloud depends on the degree of integration of customers' multi-cloud architecture. Under a siloed multi-cloud architecture where workloads have low levels of interdependence, smaller and less frequent data transfer between them may imply lower egress fees. Under a more integrated multi-cloud architecture, with larger and more frequent data transfers across clouds, egress fees may be larger.
- 1.15 In this paper, we set out analysis relating to the following issues:

² At times we use 'switching costs' as shorthand for switching and using multiple cloud costs.

³ This might involve a period of running two services in parallel with each other, which would contribute to the cost of switching, particularly if data must be egressed back and forth between those clouds.


- (a) First, we consider the prevalence of egress fees, ie the extent to which egress fees are payable as well as the extent to which they have been paid. Where egress fees are applied to only a small subset of customers – by number and by value – this may be indicative, but not probative, of how significant they are or whether they are acting as a barrier to switching and using multiple clouds;
 - (b) Second, we consider a range of evidence on the role of egress fees in customers' choice to switch or multi-cloud including evidence on the likely magnitude of egress fees that customers would experience when switching or using multiple clouds. Switching costs are more likely to reduce customers' propensity to switch when they are large in magnitude.
- 1.16 Large cloud providers have said that egress fees are reflective of costs,⁴ although we note that providers did not explain how egress fees being cost-reflective or not impacts on the extent to which they are a switching cost that could distort competition.
- 1.17 Despite this we have sought to understand the relationship between egress fees and underlying costs. We set out the type and nature of costs identified by cloud providers as incurred in providing egress fees, and analysis of egress costs compared to revenues and prices to consider the extent to which egress fees may be reflective of costs.
- 1.18 The assessments set out in this working paper reflect the evidence we have seen to date but our evidence gathering is still ongoing. Any eventual findings on whether egress fees constitute a feature that prevents, restricts or distorts competition will be made in the round and on the basis of all the evidence; we have not drawn any provisional findings in this paper. In that context, where possible, we have set out some emerging views but our thinking will continue to evolve as this market investigation progresses.

Background on egress fees

- 1.19 Cloud providers offer networking services to allow cloud customers to transfer data:
- (a) into a cloud provider's infrastructure ('ingress');
 - (b) within a cloud provider's infrastructure ('internal transfers'); and
 - (c) out of a cloud provider's infrastructure ('egress').
- 1.20 Cloud providers charge for the use of most types of their data transfer services. In this paper we will refer to these charges as data transfer fees and these are set out in Figure 1.1.

⁴ [AWS response to the Issues Statement](#) dated 17 October 2023, paragraph 24; [Microsoft response to the Issues Statement](#) dated 17 October 2023, paragraph 36; [Google response to the Issues Statement](#) dated 17 October 2023, paragraph 26.

Figure 1.1: Types of data transfers



*A content distribution network (CDN) is a geographically distributed network that delivers content to end users or applications.

Source: CMA analysis of a Google submission [34]

- 1.21 As mentioned in Figure 1.1, cloud providers generally do not charge for data ingress, but most providers charge customers when transferring data within ('internal transfer fees') and out of a cloud provider's infrastructure. We refer to egress fees as the charges customers incur when they transfer their data out of their cloud.⁵ This can occur when:
- (a) customers move their data from their cloud to their on-premises data centres;
 - (b) customers move data between different cloud providers as part of a multi-cloud architecture or as part of switching; and
 - (c) customers deliver content to end users or applications, either directly or via content distribution networks (CDNs) such as Cloudflare.
- 1.22 The focus of our assessment in this working paper is on egress fees incurred when transferring data to other cloud providers as customers are likely to do this when switching cloud provider or using multiple clouds. However, we have been told by more than one cloud provider that they cannot:
- (a) identify the purpose of a customer's data transfer out of their cloud infrastructure; or
 - (b) consistently identify the peer (ie the company and/or the relevant business unit/subsidiary within that company) to which the data is transferred and whether the peer is the end destination.⁶

Types of interconnections

- 1.23 There are different ways through which customers can move data into, out of and within a cloud provider's infrastructure. The fee that a customer incurs depends on the type of

⁵ The terminology and the definition of these fees are not consistently used by cloud providers.

⁶ See for instance: Responses to the CMA's information requests [34].

route used.⁷ Table 1.1 sets out the different routes through which customers can transfer data in, out and within a cloud provider's infrastructure.

Table 1.1: Data transfer routes

Transfer route	Internal/External	Fees
Intra-virtual network transfer	Internal	X
Inter-virtual network transfer	Internal	✓
Internet transfer	External	✓
Direct connection	External	✓
Physical transfer	External	✓
Content delivery networks	External	✓

Source: CMA analysis of publicly available information (Microsoft website, [Pricing – Bandwidth | Microsoft Azure](#); AWS website, [EC2 On-Demand Instance Pricing – Amazon Web Services](#); Google website, [All networking pricing | Virtual Private Cloud | Google Cloud](#) [accessed 14 November 2023]).

1.24 As summarised in Table 1.1, there are two types of internal cloud transfers:

- (a) Intra-virtual network transfers are transfers of data within a virtual network, eg when customers transfer data from a database to a virtual machine hosted within the same virtual network to run analytics workloads. Such transfers are free of charge within an availability zone, otherwise a fee applies.
- (b) Inter-virtual network transfers are transfers of data within a cloud that is not within a virtual network. For instance, transfer of a customer's data between storage in US and UK regions, or transfer of data between two companies' virtual networks within the same cloud. Internal transfer fees apply to such transfers.

1.25 Table 1.1 also sets out the four main routes we have identified that customers can use to transfer data out of a cloud provider's infrastructure:

- (a) Internet transfer is the transfer of data from a public cloud over the internet. This may be to another resource using a public IP address in the same cloud, to the customer's on-premises network, or to end users. Egress fees are charged on such transfers. Internet transfers can be standard or premium internet transfers.
 - (i) Standard internet transfers typically occur when the data exits the cloud provider's network close to the source and is then carried over the public internet by an internet service provider ('ISP') to the destination.
 - (ii) Premium internet transfers occur when the data is carried over the cloud provider's network backbone until close to the destination. The ISP will still carry the data for the final stretch over the public internet.
- (b) Direct connection⁸ is a way to transfer data between multiple public clouds or on-premises networks without using the internet. Data is routed over the network backbone of the cloud provider to co-location centres, typically internet exchange

⁷ For completeness, we note that there are also services that can be added onto transfer routes and generally don't change the underlying egress fees. For example, Network Address Translation ('NAT') gateways add security features to network connections over the internet. Customers buying NAT gateways incur a running cost that is listed separately, but the level of egress fee paid for the transfer remains the same with or without the service (see Google website, [Pricing | Cloud NAT | Google Cloud](#); Microsoft website, [Azure NAT Gateway - Pricing | Microsoft](#); AWS website, [NAT gateways - Amazon Virtual Private Cloud](#), accessed 14 November 2023).

⁸ We note that the terminology differs across providers. For instance, AWS refer to direct connections as 'Direct Connect', Microsoft as 'ExpressRoute' and Google as 'Dedicated Interconnect'.

points ('IXPs') run by third parties. We note that there is a distinction between direct connections to on-premises and to other clouds:

- (i) to on-premises, direct connections can be dedicated (on a private physical cable) or via a connectivity partner (on shared private cables);
 - (ii) to other clouds, direct connections can be via on-premises, via IXP or via cloud provider (ie Google's recent Cross-Cloud Interconnect).
- (c) Direct connections are generally used in cases of high volume or latency-sensitive traffic. Egress fees are charged on such transfers, but at a lower rate than for internet transfer. Hourly running fees are also charged.
- (d) Physical device transfer occurs when data is uploaded and downloaded from a physical device which is shipped back and forth between the customer's on-premises and public cloud data centre. Such transfers are generally used for the migration of data between on-premises and the cloud where the network connection is poor, or the volume of data is very large.
- (e) Content delivery networks are a specialised type of transfer where data is transferred first to a geographically distributed network of data centres and then to end users over the internet. The goal is to provide high availability and performance by distributing the service close to the end user. Cloudflare is an example of a content delivery network. These networks are typically used to distribute a relatively small number of files to a large number of users where latency is important (eg streaming media). In niche cases they might be used to switch or facilitate hybrid clouds.

1.26 Based on the characteristics and common use cases for each of these types of transfers, as described above, we currently consider that (i) internet transfer and (ii) direct connections are the most relevant external transfer routes for the purpose of assessing our theory of harm that focuses on switching and/or multi-cloud (see paragraph 1.5).

Pricing

1.27 Data transfer fees are generally volume-based and region-specific. The total charge incurred when a customer egresses data also depends on a number of factors:

- (a) Where the data is stored. Cloud providers set data transfer fees differently across their regions.
- (b) Where the data is going. The level of data transfer fees depends on whether data is going to the public internet or somewhere else (eg within cloud providers' network).
- (c) Whether the cloud provider offers free tiers and/or rates that reduce with increased data transfer volumes (eg Microsoft, AWS, Google, Oracle and IBM all offer a free tier for certain services⁹);


⁹ Microsoft website, [Pricing – Bandwidth | Microsoft Azure](#); AWS website, [Amazon S3 Simple Storage Service Pricing - Amazon Web Services](#) and [EC2 On-Demand Instance Pricing – Amazon Web Services](#); Google website, [All networking pricing | Virtual Private Cloud | Google Cloud](#); Oracle website, [Cloud Networking Pricing | Oracle United Kingdom](#), all accessed 12 March 2024; [X] response to the CMA's information request [Y].

- (d) Whether a CDN is used. For transfers to end users via the public internet, the level of the data transfer fees also depends on whether the provider's own or a third party CDN is used.
- (e) Whether the cloud provider has different routing offerings – Google and Microsoft offer two egress tiers – one where data is transferred through their own network and one where data is routed via transit (internet service provider, 'ISP') networks. Microsoft describes ISP network routing as a "cost-optimized" option – it avoids some of the internal transfer fees that would be incurred by routing through their backbone network.¹⁰ AWS also offers a similar 'premium' internet transfer using its accelerator services. This is structured differently to the Google and Microsoft offerings but serves a similar purpose.¹¹

Egress fees pricing

1.28 Figure 1.2 below shows different providers' egress fee list prices¹² for data transfer out via the public internet¹³ in March 2024.

Figure 1.2: Egress fee list prices for data transfer out via the public internet (March 2024, GBP per GB)



¹⁰ Google website, [Network Service Tiers overview | Google Cloud](#); Microsoft website, [Routing preference in Azure - Azure Virtual Network | Microsoft Learn](#), accessed 14 March 2024.

¹¹ See AWS website, [S3 Transfer Acceleration \(amazon.com\)](#); AWS website, [AWS Global Accelerator features - Amazon Web Services](#), accessed 19 March 2024.

¹² Note: All prices have been converted to GBP from USD where applicable, based on the GBP/USD exchange rate as of 19 March 2024. Google's prices have been converted from per-GiB to per-GB.

¹³ 'Standard tier' egress via the internet, or equivalent.

Sources: Microsoft website, [Pricing – Bandwidth | Microsoft Azure](#); AWS website, [Amazon S3 Simple Storage Service Pricing - Amazon Web Services](#) and [EC2 On-Demand Instance Pricing – Amazon Web Services](#); Google website, [All networking pricing | Virtual Private Cloud | Google Cloud](#); Oracle website, [Cloud Networking Pricing | Oracle United Kingdom](#), accessed 19 March 2024.; OVHcloud's website, [Price list: a comparison of our Public Cloud offers | OVHcloud UK](#), accessed 10 May 2024.

- 1.29 Figure 1.2 illustrates that list prices for AWS, Microsoft, and Google may include:
- (a) a higher marginal price for low volumes of data transfer outside of the free tier; and
 - (b) declining marginal prices for higher volumes of data transfer.
- 1.30 AWS, Microsoft, Google, and Oracle, for certain services, provide a 'free tier' whereby customers are able to egress a certain volume of data without paying any fees.¹⁴
- 1.31 Finally, Figure 1.2 shows that Oracle, and OVHcloud also have a flatter fee structure – outside of the Oracle free tier, they charge one price irrespective of the amount of data transferred. Oracle and OVHcloud also charge materially lower fees than other providers included in the comparison.
- 1.32 Figure 1.2 is based on list prices, though some customers may not be subject to list prices for egress fees via the public internet as depending on the cloud provider they may be able to negotiate:
- (a) private discounts on egress fees specifically; or
 - (b) cross-product discounts (ie committed spend agreements) that could be applied to egress fees.¹⁵
- 1.33 We have reviewed internal documents from the largest cloud providers relating to the negotiation of recent contracts with UK customers for the provision of public cloud infrastructure services where the providers agreed to offer discounts on egress fees. In the sample of documents reviewed so far, the customers had all negotiated discounts across products rather than discounts specifically related to egress fees (we consider discounts across products in a separate working paper).¹⁶
- 1.34 Microsoft submitted that the effective price paid for egress fees over the internet has declined. Microsoft also submitted that, on average, real prices for both 'metered and unmetered' egress plans have fallen since 2018.¹⁷ We are still considering this submission and as part of this working paper we have not assessed how egress fees have changed over time. We note however that, even if egress fees were shown to be declining over time, the current level of egress fees may still constitute a barrier to switching and multi-cloud, which is the focus of our theory of harm (see paragraph 1.5).

¹⁴ Microsoft website, [Pricing – Bandwidth | Microsoft Azure](#); AWS website, [Amazon S3 Simple Storage Service Pricing - Amazon Web Services](#) and [EC2 On-Demand Instance Pricing – Amazon Web Services](#); Google website, [All networking pricing | Virtual Private Cloud | Google Cloud](#); Oracle website, [Cloud Networking Pricing | Oracle United Kingdom](#), all accessed 12 March 2024.

¹⁵ Responses to the CMA's information requests [3<].

¹⁶ [3<].

¹⁷ Microsoft submission [3<].

Recent developments

- 1.35 We also describe some recent developments following the European Union's (EU) Data Act (the 'EU Data Act') coming into force.¹⁸
- 1.36 The EU Data Act entered into force on 11 January 2024.¹⁹ The EU Data Act governs the use and exchange of data within the EU and imposes obligations on cloud providers via provisions that refer to 'data processing services'.²⁰
- 1.37 The EU Data Act contains provisions relevant to egress fees for EU customers:
- (a) Article 29 requires that any 'switching charges' (which includes data egress charges relating to switching) 'cannot exceed costs incurred by the provider of data processing services that are directly linked to the switching process concerned' and applies from 11 January 2024 until 12 January 2027.²¹ From 12 January 2027 a full removal of switching charges is required.²²
 - (b) Article 34(2) requires that egress data charges for customer use of data processing services in parallel with another data processing service (which we consider to cover multi-cloud use) cannot exceed costs incurred by the provider. Our understanding is that this will take effect along with other main provision of the EU Data Act on 12 September 2025.
- 1.38 In response to the start of the application of Article 29(2) of the EU Data Act, Google announced a global programme of providing free egress for customers switching to other cloud providers (or to on-premises) on 11 January 2024.²³ The programme applies to customers who want to do a complete exit from Google Cloud. Google has indicated that it will also review cases where customers want to migrate only a portion of their data off Google Cloud on a case-by-case basis.²⁴ The free switching egress only applies to data residing in Google Cloud data storage and data management products. The approach to identifying relevant egress data transfers is for customers to submit a Google Cloud exit form, which the Google Cloud support team will then assess for eligibility and approval.

¹⁸ [Regulation \(EU\) 2023/2854 of the European Parliament and of the Council of 13 December 2023 on harmonised rules on fair access to and use of data and amending Regulation \(EU\) 2017/2394 and Directive \(EU\) 2020/1828 \(Data Act\)](#) ('EU Data Act'). The EU Data Act came into force on 11 January 2024.

¹⁹ [EU Data Act](#), Article 50.

²⁰ A data processing service is defined in the Data Act as 'a digital service that is provided to a customer and that enables ubiquitous and on-demand network access to a shared pool of configurable, scalable and elastic computing resources of a centralised, distributed or highly distributed nature that can be rapidly provisioned and released with minimal management effort or service provider interaction' (Article 2(8)). See, also, recital 81 of the Data Act: 'The generic concept 'data processing services' covers a substantial number of services with a very broad range of different purposes, functionalities, and technical set-ups. As commonly understood by providers and users and in line with broadly used standards, data processing services fall into one or more of the following three data processing service delivery models, namely Infrastructure as a Service (IaaS), Platform as a service (PaaS) and Software as a Service (SaaS). Those service delivery models represent a specific, pre-packaged combination of ICT resources offered by a provider of data processing service. Those three fundamental data processing delivery models are further complemented by emerging variations, each comprised of a distinct combination of ICT resources, such as Storage as a Service and Database as a Service.'

²¹ [EU Data Act](#), Articles 29(2) and 29(3).

²² [EU Data Act](#), Article 29(1).

²³ [Eliminating data transfer fees when migrating off Google Cloud | Google Cloud Blog](#), accessed 22 February 2024.

²⁴ Google's FAQs state 'For cases where customers are migrating only a portion of their data and not completely migrating off Google Cloud, customers are not required to terminate their Google Cloud agreement. Please contact us by completing the form with the specifics of your migration and the Google Cloud Support team will review the case and get back to you.' [Google Cloud Exit free data transfer request](#)

Customers have 60 days following approval notification to complete their migration out of Google Cloud.²⁵

- 1.39 On 5 March 2024 AWS announced a similar programme of providing free data transfer out to the internet to customers globally who want to move to another IT provider (effective as of the date of announcement).²⁶ The programme is described as being for data transfer fees when a customer is moving all data off of AWS (and won't be covered by AWS' existing free tier). However, the programme FAQs also note that customers who want to move their total usage of a single service, but not everything, off AWS can contact AWS customer support. Customers are required to contact AWS customer support to initiate the process of getting approval for this programme and once approved, have 60 days to complete their switch. Customers with less than 100GB of data stored on AWS and egress data transfers from specialised services such as AWS Direct Connect are not included in the programme.²⁷
- 1.40 On 13 March 2024, Microsoft announced that it also offers, globally, free egress for customers leaving Azure when taking their data out of the Azure infrastructure via the internet to switch to another cloud provider or an on-premises.²⁸
- 1.41 [REDACTED].²⁹
- 1.42 [REDACTED].³⁰
- 1.43 As these free switching programmes have just been announced, we will continue to gather further information to assess in detail the implications of these voluntary commitments in our consideration of the role of egress fees and any potential remedies and we welcome any views on this.
- 1.44 We will also seek to gather further information on how cloud providers will comply with the requirements of Article 34(2) of the EU Data Act for 'in parallel use' data transfer fees.

²⁵ [Google Cloud Exit free data transfer request](#), accessed 22 February 2024.

²⁶ [Promoting customer choice: AWS takes another step to lower costs for customers changing IT providers | Networking & Content Delivery \(amazon.com\)](#), accessed 22 February 2024.

²⁷ [Amazon EC2 FAQs – AWS](#), accessed 22 February 2024.

²⁸ Microsoft website, [Now available: Free data transfer out to internet when leaving Azure | Aggiornamenti di Azure | Microsoft Azure](#), accessed 14 March 2024.

²⁹ [REDACTED] response to the CMA's information request [REDACTED].

³⁰ [Google Cloud Exit free data transfer request](#), accessed 22 February 2024; Responses to the CMA's information requests [REDACTED]; [Cancel and delete your Azure subscription - Microsoft Cost Management | Microsoft Learn](#), accessed 22 February 2024.

2. Analysis of egress fees

Prevalence of egress fees

- 2.1 All UK public cloud customers have been subject to egress fees if they transfer data out of their provider's cloud, as long as the amount of data transferred does not fall within a cloud provider's free tier. These egress fees can vary based on whether a customer is eligible for additional discounts on egress fees but are structured similarly otherwise.³¹
- 2.2 In order to examine the prevalence of egress fees being paid by customers, we consider:
- (a) cloud providers' views on the prevalence of egress fees; and
 - (b) data on egress fees actually paid by customers.

Cloud providers' views

- 2.3 Cloud providers told us the proportion of their customers who pay egress fees. These varied widely:
- (a) AWS said that over 90% of customers globally with data transfer out ('DTO') usage do not pay egress fees.³² We note that this calculation is based on number of accounts rather than individual customers and is global rather than UK specific.³³
 - (b) Microsoft said that just under half of its UK customers, with less than 100 USD spend on first party public cloud infrastructure services in calendar year 2022, paid egress fees in 2022, unweighted by yearly cloud spend.³⁴ Microsoft also submitted that smaller customers - that represent a large proportion of the overall customers in 2022 - are more likely to face no egress fees.³⁵
 - (c) Google said that [REDACTED] of its UK customers paid egress fees in 2022, unweighted by yearly cloud spend. Google said that [REDACTED].³⁶

Data on egress fee spending

- 2.4 We analysed data on the prevalence of egress fees being paid by UK customers of the three largest cloud providers.
- 2.5 We gathered data from AWS, Microsoft, and Google on UK customers' cloud spend in 2022, including data on any amounts spent on egress fees. We make the following observations relating to the data, and the implications for our analysis:
- (a) The data provided by the cloud providers related to customers that spent at least 1,000 USD on cloud services each year. This threshold excludes customers with much lower spend that may be more likely to be only trialling a cloud provider's services rather than being a committed customer of cloud services. One implication of this threshold is that the proportion of customers in the dataset that paid egress fees in 2022 may be higher than in the total population of customers to the extent

³¹ Responses to the CMA's information requests [REDACTED].

³² AWS response to the CMA's information request [REDACTED].

³³ Based on [REDACTED] response to Ofcom's information request [REDACTED].

³⁴ Microsoft's response to the CMA's information request [REDACTED].

³⁵ Microsoft submission [REDACTED].

³⁶ Google response to the CMA's information request [REDACTED].

that smaller customers are more likely to fall within a cloud provider’s free tier for egress fees.

- (b) Cloud providers told us that they were unable to identify whether a customer that egresses data is sending it to a different cloud infrastructure or elsewhere (eg to end users). Therefore, our customer-level data on egress fees encompasses all the use cases set out in paragraph 1.21. These include some use cases that appear to be less directly related to customers’ ability to switch and/or multi-cloud, such as streaming content to end users. Therefore, our analysis – and more generally cloud providers’ data – will tend to overstate the extent to which egress fees were paid in relation to switching and multi-clouding.

2.6 Table 2.1 shows our estimates of both the unweighted and weighted percentages of AWS, Microsoft, and Google’ customers that paid egress fees in 2022.

Table 2.1: Unweighted and weighted proportions of customers, with at least 1,000 USD in annual spend, paying egress fees (2022)

Party	Unweighted proportion of customers paying egress fees	Weighted proportion of customers paying egress fees
AWS	[50-60]%	[90-100]%
Microsoft	[60-70]%	[90-100]%
Google	[80-90]%	[90-100]%

Source: CMA analysis of AWS, Microsoft, and Google data [Responses to the CMA’s information requests [36]].³⁷

2.7 Table 2.1 shows that most customers of AWS, Microsoft, and Google paid egress fees in 2022. Table 2.1 also shows that a higher proportion of Google’s customers paid egress fees than either of the other large cloud providers. This is consistent with the fact that in 2022, unlike Microsoft and AWS, Google did not offer a ‘free tier’ for standard tier internet data transfer fees which was introduced in 2023.³⁸ Finally, in Table 2.1 the proportion of customers paying egress fees weighted by cloud spend indicates that egress fees are more prevalent among those with higher spends on cloud.

Emerging views on prevalence of egress fees

2.8 Our emerging view on the prevalence of egress fees is that, to date, they are common and widespread across all use cases. All UK public cloud customers are subject to egress fees if they transfer data out of their cloud provider infrastructure, unless the amount of data transferred falls within a cloud provider’s free tier.

Magnitude of egress fees and role of egress fees on customers’ choice

2.9 As mentioned in paragraph 1.10 switching costs limit switching (or the use of multiple clouds), the competitive process may be distorted, because they can deter customers from switching to rivals (or using multiple clouds) that have a better overall competitive offering. The higher the level of egress fees associated with transferring data between clouds, the higher the cost of switching and/or using a more integrated multi-cloud

³⁷ The analysis is conducted at the Salesforce ID level, which AWS uses as a close proxy for unique customer. However, some of AWS’ [36] customers do not have a corresponding Salesforce ID and are excluded from the analysis. Furthermore, the analysis also excludes around [36] of AWS customers where AWS was not able to determine those customers’ egress spend.

³⁸ Google’s free tier was introduced 1 October 2023 (Google website, [Announcing 200 GB free Standard Tier internet data transfer per month | Google Cloud Blog](#), accessed 21 March 2024).

architecture becomes. Higher levels of switching cost will tend to make customers less responsive to the efforts of rivals to compete by improving their offer.

2.10 In this section, we consider a range of evidence relevant to the role of egress fees on customers' choices. In particular, we consider evidence on:

- (a) customer data on spending on egress fees;
- (b) the cost of egress under hypothetical scenarios;
- (c) customer views on the scale of egress fees; and
- (d) internal documents relevant to egress fees.

Customer data on spending on egress fees

2.11 We analysed data on the amount of egress fees being paid by customers of cloud infrastructure services. The considerations set out in paragraph 2.5 relevant to the analysis of prevalence of egress fees are also relevant here.

2.12 In addition, we observe that:

- (a) Data on egress fee spending is an imperfect proxy for the actual importance of egress fees. Where egress fees are high, customers may seek to avoid them by deciding not to switch or multi-cloud (or doing so to a lesser extent). This means that an analysis of actual egress fee spend is likely to understate the importance of egress fees.
- (b) Cloud providers told us that they were unable to identify whether a customer that egresses data is sending it to a different cloud infrastructure or elsewhere (eg to end users). Therefore, our customer-level data on egress fees encompasses all the use cases set out in paragraph 1.21. These include some use cases that appear to be less directly related to customers' ability to switch and/or multi-cloud. Therefore, our analysis – and more generally cloud providers' data – will tend to overstate the extent to which egress fees were paid in relation to switching and multi-cloud.
- (c) When switching, the cost of egress fees can be defrayed over a number of years. Therefore, the proportion of egress fees spend over the total spend in a given year may overstate the impact on customers' decision making.

2.13 In light of the observations above, and in particular given the effects in terms of understating or overstating the significance of egress fees, data on actual spend on egress fees must be assessed with caution:

- (a) High spend on egress fees may be explained by the rate of the fee being set at a high level, but also potentially by large spend on egressing fees to end users (which is less relevant for our purposes).
- (b) Low spend on egress fees may be explained by the rate of the fee being set at a low level, but also potentially by customers' responses to high fee rates.

2.14 Therefore, while it is useful to know the current level of egress fee spending as a matter of context, other evidence on the significance of egress fees and their impact on customer behaviour is likely to be more informative.

2.15 Using the data described in paragraph 2.5, we analysed the current magnitude of egress fees as a share of cloud spend by the customers in the dataset – ie customers spending at least 1,000 USD on cloud services in 2022. Table 2.2 sets out our estimates of the

percentage of the revenues generated by egress fees in 2022 for AWS, Microsoft and Google.

Table 2.2: Cloud providers – percentage of revenue generated by egress fees in the UK, for customers with at least 1,000USD in annual spend (2022)

Party	Percentage of revenues generated by egress fees, 2022
AWS	[0-5]%
Microsoft	[0-5]%
Google	[0-5]%

Source: CMA analysis of AWS, Microsoft and Google data Responses to the CMA's information requests [3<].³⁹

2.16 This data indicates that egress fees account for a small proportion of the providers total UK public cloud revenues, ie between [0-5]% and [0-5]% (see Table 2.2).

2.17 Table 2.3 illustrates, for AWS, Microsoft, and Google respectively, the distribution of customers based on the percentage of their total spend in 2022 that was on egress fees.

Table 2.3: Egress fees as a percentage of total spend for customers with at least 1,000USD in annual spend, 2022 (by customer, cumulative percentage)

Egress fees as a percentage of total spend, 2022 (by customer, cumulative percentage)	AWS	Microsoft	Google
Less than 1%	[70-80]%	[80-90]%	[50-60]%
Over 1%	[20-30]%	[10-20]%	[40-50]%
Over 2%	[10-20]%	[10-20]%	[20-30]%
Over 5%	[5-10]%	[5-10]%	[10-20]%
Over 10%	[5-10]%	[0-5]%	[10-20]%
Over 25%	[0-5]%	[0-5]%	[0-5]%
Over 50%	[0-5]%	[0-5]%	[0-5]%
Over 75%	[0-5]%	[0-5]%	[0-5]%

Source: CMA analysis of AWS, Microsoft, and Google data Responses to the CMA's information requests [3<].

2.18 These figures show that a significant majority of customers in 2022 paid egress fees making up less than [0-5]% of their total spend on cloud.

2.19 Around [10-20]% of AWS and Microsoft's UK customers paid [0-5]% of their annual cloud spend on egress fees. This proportion is higher for Google's UK customers, where [20-30]% paid over [0-5%] of their annual cloud spend on egress fees. Roughly [10-20]% of Google customers paid egress fees representing [0-5%] of their total cloud spend, whereas about [5-10]% of AWS customers and [5-10]% of Microsoft customers did.

The cost of egress under hypothetical scenarios.

2.20 An alternative way of considering the magnitude of egress fees is to consider how much it would cost to switch or to pursue a multi-cloud strategy based on the levels of fees charged by providers and data on customer characteristics that are relevant to the cost of switching or using a multi-cloud architecture.

2.21 In the sections below we will discuss evidence on the hypothetical scenarios on switching and multi-cloud separately. This includes the views of cloud providers and estimated egress costs based on hypothetical scenarios.

³⁹ The analysis is conducted at the Salesforce ID level, which AWS uses as a close proxy for unique customer. However, some of AWS' [3<] customers do not have a corresponding Salesforce ID and are excluded from the analysis. Furthermore, the analysis also excludes around [3<] of AWS customers where AWS was not able to determine those customers' egress spend.

Hypothetical switching scenario

2.22 Below we analyse how costly it would be to egress data for customers willing to migrate to another cloud provider. To assess whether potential egress fees would represent a significant proportion of customers' annual cloud spend, we will consider evidence from:

- (a) cloud providers' views, and
- (b) data on the egress fees that would be incurred based on a 'one-off' switch of provider.

Cloud providers' submissions relating to the hypothetical 'one-off' switching costs

2.23 AWS, Microsoft and Google said that egress fees were a small percentage of customers' total costs of switching. For example:

- (a) AWS, using its own pricing calculator,⁴⁰ estimated that the egress fees incurred by switching customers represent less than [X] of their annual spend on AWS. AWS said this is a level unlikely to deter customers from switching and is consistent with only [X] of respondents to Ofcom's market research perceiving 'data charges' as the main challenge of switching provider completely. AWS also submitted a secondary analysis, based on its own data, of data transfer out charges incurred by customers in the process of switching which estimated that these charges represent less than [X] of these customers' annual spend.⁴¹
- (b) Microsoft referenced an IDC Report produced for the European Commission which estimated that the one-off egress fee cost of a customer switching all data was approximately 0.3% – 0.35% of the expected annual operating cost using either Azure or AWS for cloud services.⁴² Microsoft also said that it provides customers with an option to use direct connections to Azure that allow customers transferring significant amounts of data to reduce their overall egress fee charges and to transfer at faster latency (eg Azure ExpressRoute).⁴³
- (c) Google estimated that egress fees as a direct result of switching ('True Exit fees') exceed no more than [X] of a customer's contract spend over a typical three-year period. They add that egress fees covering all types of egress traffic are still a small proportion of an average customer's spend.⁴⁴

2.24 OVHcloud submitted that egress fees are unduly used by AWS, Microsoft and Google to lock customers in their ecosystem and prevent them from switching.⁴⁵

Hypothetical 'one-off' switching cost

2.25 We calculated the hypothetical 'one-off' switching cost for a customer based on a scenario where they transfer all of their data out of their current cloud, and the percentage of their annual revenue that cost would represent. This was based on the specific characteristics of UK customers. In particular:

⁴⁰ [AWS Pricing Calculator](#)

⁴¹ AWS submission to Ofcom [X].

⁴² [Microsoft's response to Ofcom's Interim Report](#), 31 May 2023, paragraph 237.

⁴³ [Microsoft response to the Issues Statement](#) dated 17 October 2023, paragraph 41.

⁴⁴ Google submission to Ofcom [X].

⁴⁵ OVHcloud submission [X].

- (a) we collected UK customer-level data on the minimum, maximum and mean volumes of data they stored for each year from 2018 to 2023;
- (b) we used the list prices on providers' websites for data transfer from London via the internet;⁴⁶
- (c) we multiplied each customer's mean data storage volume by the price per GB for egress that would apply were they to transfer all of this data at once; and
- (d) we then calculated the percentage of annual spend this would make up for each customer in each applicable year.

2.26 We would make the following observations to our analysis of 'one-off' switching cost:

- (a) Our estimates are based on a hypothetical 'one-off' transfer of all the data stored on one cloud provider. Our analysis will tend to understate the extent to which egress fees are paid in relation to switching. 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.
- (b) Transfer costs may differ based on the type of interconnection used, or discounts agreed with a provider. We currently understand that egress fees specific discounts are rare, but that discounts on egress fees may arise from wider cross-product discounts which we consider separately in the Committed spend agreements working paper.
- (c) To the extent that switching involves a one-off transfer of data (which may depend on if some form of integrated multi-cloud is needed during the switching process) then egress fees are incurred only during the switching process rather than on a recurring basis as would be the case if a customer adopted an integrated multi-cloud architecture.

2.27 Figure 2.1, Figure 2.2 and Figure 2.3 illustrate, for each of AWS, Microsoft, and Google, the cost to a customer of 'one-off' switching as a percentage of its total spend for each available year.

Figure 2.1 [X]

Figure 2.2: [X]

Figure 2.3: [X]

2.28 Overall, egress fees paid to move all of a UK customer's data once from one cloud provider to another range between [0-5]% to [5-10] % of a customer's annual spend. Specifically:

- (a) Figure 2.1 shows that the average UK customer of AWS would incur costs of around [0-5]% of their annual costs if they were to perform a 'one-off' switch away from AWS. This in 2023 is explained by annual revenues for that year covering only January to September.

⁴⁶ Google: [All networking pricing | Virtual Private Cloud | Google Cloud](#), Microsoft: [Pricing – Bandwidth | Microsoft Azure](#), AWS: [Amazon S3 Simple Storage Service Pricing - Amazon Web Services](#) and [EC2 On-Demand Instance Pricing – Amazon Web Services](#), all accessed 29 Jan 2024.

- (b) Figure 2.2 shows that, for the average UK customer of Microsoft, ‘one-off’ switching would represent around [0-5]% to [0-5]% of their average annual costs. This percentage has fluctuated over the six-year period between 2018 and 2023, generally rising over time.
- (c) Figure 2.3 shows that, for the average UK customer of Google, ‘one-off’ switching would represent [0-5]% of their average annual costs. This has stayed consistent over the three-year period we have examined.

2.29 We also considered the distribution of customers’ ‘one-off’ switching cost as a percentage of annual cost in 2022. Table 2.4 shows our results.

Table 2.4: Distribution of one-off switching cost percentages, 2022 (cumulative percentage)

<i>Distribution of one-off switching cost percentages, 2022 (cumulative percentages)</i>	<i>AWS</i>	<i>Microsoft</i>	<i>Google</i>
Less than 1%	[50-60]%	[30-40]%	[30-40]%
Over 1%	[40-50]%	[60-70]%	[60-70]%
Over 2%	[30-40]%	[40-50]%	[40-50]%
Over 5%	[10-20]%	[10-20]%	[20-30]%
Over 10%	[5-10]%	[5-10]%	[10-20]%
Over 25%	[0-5]%	[0-5]%	[0-5]%
Over 50%	[0-5]%	[0-5]%	[0-5]%
Over 75%	[0-5]%	[0-5]%	[0-5]%

Source: CMA analysis of [redacted] Responses to the CMA’s information requests [redacted].

2.30 Overall, our analysis shows that the egress fee charge for a one-off switch is subject to a broad distribution:

- (a) Most of AWS’, Microsoft’s and Google’s UK customers would have had to pay less than 5% of their total annual spend in 2022 if they were to transfer all of their data out of their current cloud. Between [10-20]% and [20-30]% of customers would have had to pay egress fees of more than 5% of their total annual spend.

Hypothetical multi-cloud scenario

2.31 In its market study, Ofcom estimated how costly it would be for customers using multiple clouds to egress data between cloud providers. To do this Ofcom used some stylised scenarios of how much data customers might transfer between cloud providers in a hypothetical multi-cloud architecture. Ofcom’s model found that a hypothetical multi-clouding customer could face egress fees costing between 5% and 15% of its overall annual cloud spending.⁴⁷

2.32 [redacted].⁴⁸

2.33 We are still considering the scenarios used by Ofcom [redacted]⁴⁹ as well as gathering additional evidence to consider whether and how to do any further work on analysing hypothetical scenarios that estimate the costs of egressing data when using a multi-cloud architecture.

Customer views

2.34 In order to assess whether egress fees are having an impact on customer behaviour, or are likely to do so, we asked customers a range of questions. In this section, we set out

⁴⁷ Cloud services market study final report (ofcom.org.uk), paragraphs 5.147 to 5.151.

⁴⁸ [redacted] submission [redacted].

⁴⁹ [redacted].

the evidence gathered from large customers on the significance of egress fees in the context of switching and operating a multi-cloud architecture.

- 2.35 In line with the qualitative nature of the evidence we gathered, we have given a narrative summary of the key points that we consider emerge from the evidence.
- 2.36 We also commissioned primary market research from Jigsaw Research. This research was intended to capture a wider range and a different set of customers from those we spoke to through direct channels. We are still considering the full details of this primary research and have only incorporated key takeaways where relevant in this working paper. The evidence is set out in full in a separate paper (the Jigsaw report) and we will consider the evidence from this research alongside the evidence outlined in this paper in our ongoing work.

Large customers' views on switching

- 2.37 In interpreting evidence from customers at this stage, we consider that:
- (a) some customers may not have been in a position to give an informed view on switching. For example, if a customer has considered switching or has switched they are likely to have a more informed view of whether egress fees are a challenge to switching. In contrast, if a customer is relatively new to public cloud and has not considered switching then they are likely to be relatively less informed. In that context, we may place more weight on customers who have switched or had considered switching.
 - (b) egress fees may be a problem but may be overshadowed by other more issues⁵⁰ or customers are not complaining about egress fees, although such fees may nonetheless have affected their switching behaviours.⁵¹
- 2.38 In this context, we asked large UK customers whether they had previously switched or considered switching between public clouds and, if so, to explain any challenges they encountered (or anticipate they would encounter), distinguishing between commercial (including egress fees, committed spend discounts, software licensing costs) and technical aspects. We asked large UK customers who had never considered switching why this was the case.

Customers that have switched or considered switching

- 2.39 Only a few customers that we spoke to spontaneously identified egress fees as a challenge. This included because they were a cost to switching and could be significant, even if they could be mitigated to some extent.
- 2.40 Only a few other customers spontaneously identified that egress fees were not a challenge. This included because egress fees could be mitigated in some instances.

⁵⁰ For example, a customer may have considered switching and done an initial assessment of the technical effort involved in switching and decided against switching without considering other commercial aspects. In this situation it is plausible that egress fees could still have affected switching behaviour in the absence of that technical effort.

⁵¹ Customers may have become accustomed to egress fees and consider it an unavoidable cost of the service, given all the largest providers charge them.

Customers that have never considered switching

- 2.41 Many customers we spoke to had never considered switching because they were satisfied with or had only recently moved to their current provider.
- 2.42 A few of these customers explicitly identified egress fees as a factor that could disincentivise switching. This is because egress fees can be significant given the volume of data stored.

Other customer evidence

- 2.43 Not all customers received the same set of questions, and some customers we spoke to who were not asked the questions set out above also discussed egress fees in the context of switching. We will consider this evidence in the round and alongside other evidence received and discussed above.
- 2.44 A few of these customers indicated egress fees would impact their decision making when considering a change of cloud provider as these fees could be substantial if switching applications from one cloud provider to another or, more generally were cited as a barrier to switching.⁵²
- 2.45 A few of these customers said that they had not found egress fees an issue or they had not explicitly considered them and that there were other costs to switching that were far more significant.⁵³

Evidence from the Jigsaw report on switching

- 2.46 The Jigsaw report sets out how some participants said that egress fees play a part in disincentivising them from potentially considering a migration to another cloud provider. However, the Jigsaw report highlights that in almost no cases were egress fees considered as the main or even one of the main barriers to switching.
- 2.47 For most participants, the technical challenges combined with the lack of a clear business case were the main barriers to switching, with egress fees in some cases contributing to the reluctance to consider a switch. When egress fees did contribute to this reluctance this typically involved a company with relatively high data volumes, or mature companies who have accumulated a large amount of data over years.
- 2.48 Finally, the Jigsaw report shows that those participants who did switch, consider the egress fees they incurred a price worth paying to deliver the cloud strategy that makes most sense for their business.⁵⁴

Large customers' views on multi-cloud

- 2.49 We asked large customers to explain any challenges they encountered in setting up or operationalising the multiple public clouds they were using and in setting up or operationalising any integrations between them. We asked them to distinguish between commercial aspects (explicitly referring to egress fees, committed spend discounts, and

⁵² Notes of meetings with [redacted]; [redacted] response to the CMA's information request [redacted].

⁵³ Notes of meetings with [redacted].

⁵⁴ Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024).

software licensing costs as examples) and technical aspects. We consider the points set out at paragraph 2.37 apply equally to this evidence.⁵⁵

- 2.50 As described in the Competitive landscape working paper, multi-cloud architectures lie along a spectrum of different degrees of integration. The degree of integration may influence the responses of customers as well as the interpretation of those responses.
- 2.51 As set out in paragraph 1.14(b), the extent to which egress fees are incurred depends on the volume and the frequency of the data transfers between public clouds (ie the degree of integration). We therefore expect customers with more integrated forms of multi-cloud to have more insights on whether egress fees have been a challenge when setting up their multi-cloud architecture, whereas customers with limited integration across clouds may not have formed a view on whether egress has been a barrier to set up their multi-cloud architecture.
- 2.52 We consider responses below based on the self-reported type of multi-cloud architecture the customer we spoke to currently uses.

Customers with integrated cloud architectures

- 2.53 Many customers we spoke to identified themselves as using an integrated form of multi-cloud architecture, although their multi-cloud architecture is on the less integrated end of the spectrum.⁵⁶
- 2.54 A majority of these customers said that egress fees have been a challenge to multi-cloud architectures and/or took them into consideration when taking their decision to set up a multi-cloud architecture. Reasons for this included that they make it difficult to use the lowest cost service across cloud provider and they incentivise the use of one cloud provider.
- 2.55 A few of these customers said that they did not consider egress fees to be a material challenge in setting up a multi-cloud infrastructure.⁵⁷

Customers with siloed cloud architecture

- 2.56 A handful of customers said that they used a siloed multi-cloud architecture.⁵⁸
- 2.57 Only a few of these customers suggested that egress fees had a significant impact on their multi-cloud architecture. Reasons included because egress fee costs were significant enough to be a consideration in the design phase of a multi-cloud architecture.
- 2.58 One of these customers said that egress fees were not a material challenge because they were not a material disincentive to multi-cloud for its use cases.⁵⁹

Customers with single-cloud architectures

- 2.59 A few of the customers we spoke said they used a single cloud provider.⁶⁰

⁵⁵ For example, some customers may not have perceived egress fees to be a challenge because they may have only a limited level of integration in their multi-cloud infrastructure for other reasons.

⁵⁶ Responses to the CMA's information requests [§].

⁵⁷ [§] response to the CMA's information request [§].

⁵⁸ Responses to the CMA's information requests [§].

⁵⁹ [§] response to the CMA's information request [§].

⁶⁰ Responses to the CMA's information requests [§].

- 2.60 We asked these customers the reason for using a single public cloud and any potential challenges that they may encounter if moving to a multi-cloud architecture. None of the customers commented on whether egress fees were a reason to adopt a single cloud architecture or whether egress fees had been a challenge to start using multi-cloud.

Other customer evidence

- 2.61 Not all customers who we spoke to received the same set of questions and some customers who were not asked the questions set out above also discussed egress fees in the context of using multiple clouds. We will consider this evidence in the round and alongside other evidence received and discussed above.
- 2.62 A handful of these customers indicated egress fees would impact their decision making when considering a switch to multi-cloud or changing cloud providers, as egress fees are expensive given the volume of data that they would need to transfer or more generally were identified as a barrier to switching.⁶¹
- 2.63 One of these customers said that they had mitigated egress fees through the architecture they had employed.⁶²

Evidence from the Jigsaw report on multi-cloud

- 2.64 In relation to egress fees, the Jigsaw report is fairly consistent with respect to switching and the use of multiple clouds. As with switching, some participants said that egress fees play a part in disincentivising them from potentially considering a multi-cloud approach.
- 2.65 However, the Jigsaw report highlights that in almost no cases were egress fees considered as the main or even one of the main barriers to a multi-cloud approach. For most participants, the technical challenges combined with the lack of a clear business case were the main barriers to switching, with egress fees in some cases contributing to the reluctance to consider the use of multiple clouds.
- 2.66 When egress fees did contribute to this reluctance this typically involved a company with relatively high data volumes, or mature companies who have accumulated a large amount of data over years.
- 2.67 Finally, the Jigsaw report shows that those participants who did multi-cloud, consider the egress fees they incurred a price worth paying to deliver the cloud strategy that makes most sense for their business⁶³

Internal documents

- 2.68 In this section, we set out the evidence from cloud providers' internal documents on their commercial strategies in relation to egress fees, including the rationale for charging these fees.
- 2.69 Overall the cloud providers' internal documents show that:
- (a) Decisions regarding egress fees are usually discussed together with internal transfer fees, and that internal transfer fees appear to be lower than egress fees.⁶⁴

⁶¹ Responses to the CMA's information requests [3<], Notes of meetings with [3<].

⁶² [3<] response to the CMA's information request [3<].

⁶³ Cloud Services Market Investigation Qualitative Customer Research conducted by Jigsaw (2024).

⁶⁴ See for instance: Responses to the CMA's information requests [3<].

- (b) When setting their networking fees, the cloud providers usually consider both their own costs and/or margins and their competitors' pricing amongst other factors.⁶⁵
- (c) [X] and [X] internal documents [X] show that direct connections are seen by cloud providers as being used to target multi-cloud customers. [X] documents also suggest that direct connections are useful for customers to reduce the cost of egress when migrating content to another cloud provider.⁶⁶
- (d) [X] internal document shows that it discussed setting egress fees to zero to gain a strategic advantage over its competitors. In doing so, [X] internal document framed the question as whether to be the 'no lock-in cloud' which indicates that [X] sees egress fees contribute to lock-in. The proposal to set egress fees to zero, was explored, but this strategy would be 'easily matched by competitors and it was too 'costly [X] and the industry as a whole'. We note that [X] was discussing egress fees in general and not just egress fees paid specifically for the purpose of switching and/or multi-cloud.⁶⁷
- (e) [X] internal documents show that it has considered using internal transfer fees and egress fees to influence the regions that customers use (ie move demand away from capacity-constrained regions). However, ultimately [X] decided that this would not be effective unless it also used the pricing of other products such as its compute and storage products to influence the regions that customers use. This is consistent with data transfer fees (internal and egress) being one factor that can influence customer demand.⁶⁸

Emerging views on magnitude of egress fees and role of egress fees on customers' choice

- 2.70 Based on the evidence we have seen to date, our emerging view on the magnitude and role of egress fees on customers' choice is that egress fees can be relevant to customers' decision making when switching or using multiple clouds. Specifically:
- (a) Our estimates of hypothetical 'one-off' switching costs indicate that customers would have to pay [0-5]% to [5-10]% of their annual spend to switch all the data stored in with their current cloud providers. These costs can also be materially higher for some customers. This additional cost could make it more difficult for customers to switch cloud provider. In addition, to the extent customers would need to run an integrated multi-cloud architecture for the duration of their switch, egress fees may represent an even higher proportion of their annual cost of cloud.
 - (b) Some respondents explicitly identified egress fees as a challenge when considering or a factor disincentivising switching. In addition, some respondents said that egress fees had been a challenge to multi-cloud architectures and/or took them into consideration when taking their decision to set up a multi-cloud architecture; and
 - (c) One cloud provider's internal documents show how data transfer fees could be used, alongside other cloud services, to move customer demand away from capacity-constrained regions. Also, another cloud provider's internal document suggests making egress fees free as a way to reduce customer lock-in.

⁶⁵ See for instance: Responses to the CMA's information requests [X].

⁶⁶ Responses to the CMA's information requests [X].

⁶⁷ [X] response to the CMA's information request [X].

⁶⁸ [X] response to the CMA's information request [X].

- 2.71 Our analysis on current spend on egress fees must be interpreted cautiously, given the scope for customer responses to high egress fees to cause the estimates to understate the magnitude of egress fees. Issues with the data, including the inability to distinguish egress to end users means the estimates include costs that are unrelated to switching or the use of multiple clouds.
- 2.72 The data therefore serve only to set out some contextual information about current levels of spending on egress fees. In that context, a majority of the cloud providers' UK customers in the relevant data set paid egress fees of less than 1% of their annual spend on cloud. However, a minority paid more substantial proportions.
- 2.73 We are considering, and will continue to consider, this evidence in its appropriate context alongside other evidence we received in response to this working paper.

3. Cloud providers' cost of egressing data

- 3.1 This section considers the costs cloud providers incur in providing egress data transfers and the extent to which these costs are egress-specific and incremental in nature. We also analyse egress costs relative to the egress fees charged by cloud providers, to understand the extent to which egress fees charged to customers reflect cloud providers' underlying costs of transferring data.
- 3.2 We have undertaken this analysis as large cloud providers have said that egress fees are important for recovering costs and investment. We are also aware that large cloud providers generally charge internal (within-cloud) data transfer rates which are substantially lower than their egress fee list prices and that they do not charge for ingress data transfers, and we are interested in understanding the extent to which there are differences in the cost of providing egress that these pricing differences reflect.
- 3.3 At this stage our analysis is focussed on understanding the costs of providing egress data transfers overall, and relative to egress fees charged. This is to assess the rationale put forward by large cloud providers and to consider whether in the current market egress fees are reflective of costs, or may be at a level above what we consider to be relevant benchmarks.

Relevant costs for egress data transfers

- 3.4 As part of our analysis, we set out to understand what costs are incurred by cloud providers in providing data transfer services, what costs are specifically relevant to egress data transfers, and the extent to which these apply incrementally per instance of egress for switching or multi-clouding.
- 3.5 We have summarised the assets cloud providers have identified as relevant to data transfers, what cloud providers have said are the costs of data transfer, and how they are incurred in Appendix A. Broadly, cloud providers identified network infrastructure costs and connectivity costs (to connect network locations within a cloud and to external parties, and to transfer data out of a cloud) as incurred in providing data transfers for customers. Some cloud providers also cite shared and/or overhead business costs such as sales and marketing as costs for providing data transfer services.⁶⁹
- 3.6 As set out above, cloud providers generally do not charge for data ingress, but most providers charge customers for egress. Large cloud providers also typically charge a much higher rate for egress data transfers compared to internal (within cloud) data transfers.⁷⁰ Given this, we have sought to understand the extent to which there are costs which are specific to egress which may explain the difference in egress pricing relative to ingress and internal data transfer pricing. Additionally, as egress data transfers are charged based on volumes (ie at a \$ per GB rate) we consider the extent to which the costs for providing egress are incrementally incurred, to understand the extent to which volume-based pricing may reflect incrementally incurred cost recovery.

⁶⁹ Responses to the CMA's information requests [8].

⁷⁰ See for example: [EC2 On-Demand Instance Pricing – Amazon Web Services](#) prices for data transfer out to the Internet compared to data transfer within the same AWS region and to other AWS regions; [Pricing | Cloud Storage | Google Cloud](#) prices for general network usage (eg data transfers to the Internet) compared with data transfer within the same Google region; [Pricing - Bandwidth | Microsoft Azure](#) prices for Internet egress compared to data transfers between Availability Zones or within the same Availability Zone. Webpages accessed on 20 May 2024.

- 3.7 Our analysis of egress-specific costs highlights connectivity as an element that cloud providers commonly identify. Given this, we also consider how, and the extent to which, connectivity costs are incurred by cloud providers.

Common data transfer costs vs specific costs for egress

- 3.8 Cloud providers all said that generally the same fixed assets are or may be used to ingress and egress customer data.⁷¹
- 3.9 One cloud provider noted that the same network supports inbound and outbound transfers at the same time, but that their capacity is built to support the requirements of the dominant direction of traffic (which for the cloud provider is outbound dominant) rather than the sum of both directions.⁷² The cloud provider has said that, as at September 2023, over [X] of its regions are outbound dominant (including [X]) but it does not account for costs for inbound transfers in any regions.⁷³
- 3.10 Based on this shared use of assets, it appears that the capex incurred for network fixed assets and any ongoing annual expenses associated with these assets (eg depreciation, maintenance) are common to both the ingress and egress of data, and hence the recovery of this fixed cost expenditure is applicable to provision of both ingress and egress services. Whilst egress may be a larger driver than ingress for additional network investment for some cloud providers, it appears that any investment will be used by both ingress and egress services, and we view enabling ingress to be an essential function for cloud providers to be able to provide other cloud services such as data storage.
- 3.11 We note that cloud providers generally provide ingress data transfers for free, which does not appear to be purely reflective of a difference in costs between providing ingress and egress, given cloud providers have identified shared costs for egress and ingress (ie ingress is not considered to be nil cost to provide).
- 3.12 With regard to comparisons between egress and internal data transfers, cloud providers indicated that much of the same network infrastructure is or can be used to support both internal and egress data transfers.⁷⁴ Cloud providers expressed somewhat varying views on what were common vs specific costs for egress compared to internal data transfers. However, cloud providers identified some costs which were specific to egress and not relevant to internal data transfers. We summarise cloud providers' submissions on egress-specific costs in Appendix A.
- 3.13 Internet transit (ie bandwidth) and peering costs appear to be a common egress-specific cost across a number of cloud providers.⁷⁵ Additionally, a number of the cloud providers appear to have some overlap in their views of public internet connection infrastructure and edge networking costs (including transit fees, peering charges and the costs of edge network assets) as being specific to egress. We note that IP transit and peering charges are typically variable (eg \$/mbps) whereas asset costs are fixed.
- 3.14 This is consistent with OVHcloud's view that the only technical justification for egress fees is the cost for using internet bandwidth from ISPs or from other operators, such as suppliers of connectivity solutions.⁷⁶

⁷¹ Responses to the CMA's information requests [X].

⁷² [X] response to the CMA's information request [X].

⁷³ [X] response to the CMA's information request [X].

⁷⁴ Responses to the CMA's information requests [X].

⁷⁵ See Appendix A

⁷⁶ OVHcloud submission [X].

- 3.15 If both egress and internal data transfer fees charged by cloud providers are based on cost recovery, then the difference between egress and internal data transfer fees would be expected to reflect these identified egress-specific costs.

Incremental costs

- 3.16 As noted above, we are interested in the extent to which the costs for providing egress are incrementally incurred to understand how egress prices may reflect incrementally incurred cost recovery, and as egress data transfers are charged to customers based on incremental volumes. For fixed costs or additional investment (eg in network expansion) specifically for data transfer services, cloud providers may have some degree of control over the timing of those costs and how quickly costs are recovered through fees, which may also mean that they are less directly attributable per instance of egress and less comparable between cloud providers. Additionally, for infrastructure investment and business costs that may support the provision of multiple cloud services (eg data centres, HR costs), the extent to which costs are recovered from data transfer services and other cloud services will be subject to cloud providers' business strategy and pricing decisions.
- 3.17 Cloud providers have provided us with limited detail on their incremental costs of transferring or egressing customer data. They also expressed differing views on what should be considered incremental costs for egress.⁷⁷
- 3.18 As set out in Appendix A, cloud providers have highlighted the costs of connectivity to the public internet as an egress-specific cost. We note that some of the incremental costs identified by some cloud providers listed above relate to connectivity costs. We note that some arrangements (paid peering and IP transit⁷⁸) are identified as being charged per mbps based on the amount of traffic exchanged between the networks.⁷⁹ We consider the extent to which costs of connectivity for egress data transfers are incrementally incurred in the section below.

Extent to which identified costs are incrementally incurred

- 3.19 Based on the descriptions of costs provided and set out in the previous section, bandwidth excess fees (or 'burst fees') may be the only type of cost that can be incremental per individual data egress instance (when a cloud providers' data traffic has exceeded monthly fixed bandwidth levels). For some of the costs identified, eg hardware maintenance and energy consumption, it is not clear that these are incurred per instance of egress, but instead appear more likely to be variable costs based on general level of activity.
- 3.20 We recognise that there will likely be limits to the total egress volume that a cloud provider can deliver for a given amount of network infrastructure,⁸⁰ and therefore accommodating future capacity may incur additional investment. However, we also recognise that increasing capacity to accommodate additional egress would be step-up investment rather

⁷⁷ Cloud providers' submissions in relation to incremental costs are summarised in Appendix A.

⁷⁸ Peering (which for a cloud provider may be paid or settlement-free, and done through interconnection or an Internet exchange (IX) port) is the exchange of traffic between networks. Peering uses direct connections between networks to send traffic to these networks, whereas for IP transit traffic is exchanged between the networks and routed by the transit provider and could terminate anywhere in the world.

⁷⁹ See Appendix B where we set out cloud providers' connectivity arrangements and identified costs in more detail.

⁸⁰ Oracle said total egress volume is typically limited by 4 primary factors: 1) total bandwidth to egress networks, 2) total circuits (ie network infrastructure ports) connected to egress networks, 3) traffic engineering, and 4) other costs. Oracle response to the CMA's information request [3<].

than incrementally incurred costs per egress transfer, for which cloud providers would have some degree of control over how and when such investment is recovered (particularly if investment supports not only egress data transfer services, but other cloud services).

- 3.21 It is also not clear how there is a difference in how costs are incurred for ingress compared to egress. Ingress would seem to be justifiable as a one-time cost only if customers only ingressed data once; this does not seem applicable for customers that ingress multiple times due to multi-cloud or hybrid cloud (cloud and on-premises) use.
- 3.22 Additionally, there may be circumstances where there is no incremental cost to a cloud provider for an egress data transfer due to having settlement free peering in place. Our analysis indicates that a substantial portion of some cloud providers' peering for egress data transfers is settlement-free. One cloud provider said most of its internet peering⁸¹ is settlement-free.⁸² Another cloud provider said internet peering arrangements with ISPs are mostly on a settlement-free basis.⁸³ As noted in Appendix B, Microsoft said it has settlement-free contracts for all its UK ISPs, so it does not pay any fixed or variable costs to ISPs in the UK.⁸⁴
- 3.23 Given this, it appears that there may be some peering or IP transit costs that are unavoidable in order for a cloud provider to be able to provide egress data transfer services to any end user location, but that as cloud providers, CDNs, NSPs and ISPs increasingly interconnect and establish free peering and data transfer arrangements, the proportion of network spend on transit is likely to reduce. Our initial analysis also shows that transit costs comprise a relatively small proportion of networking costs for large cloud providers.⁸⁵
- 3.24 Moreover, incremental egress costs for cloud-to-cloud data transfers, which applies to transfers for switching or multi-cloud use, appear to be avoidable where a cloud provider can negotiate settlement free peering with the destination cloud network, at the relevant location.
- 3.25 This is consistent with OVHcloud's view that some cloud providers have made private peering or private network interconnect (PNI) agreements to interconnect their services to reduce costs. OVHcloud said that by using private network to transfer data, they do not require the use of internet bandwidth.⁸⁶
- 3.26 A CDN provider said that the marginal costs of data transfer for AWS, Microsoft, and Google are often near-zero for large customers, because they use third party CDNs who have interconnected infrastructure to these cloud providers.⁸⁷
- 3.27 We note that some cloud providers also offer private connection services, which may avoid the use of public internet egress, which may offer customers benefits such as greater reliability, security, and/or lower latency for transferring data but generally have different pricing structures to standard data transfers. They are not necessarily an overall cost-minimising option for customer egress (in particular for switching and/or multi-cloud).

⁸¹ We take internet peering to collectively refer to peering directly with ISP networks and other networks (eg other cloud providers) and via IX, and IP transit.

⁸² [3<] response to the CMA's information request [3<].

⁸³ [3<] response to the CMA's information request [3<].

⁸⁴ Microsoft response to the CMA's information request [3<].

⁸⁵ CMA analysis of transit costs: responses to the CMA's information requests [3<].

⁸⁶ OVHcloud submission [3<].

⁸⁷ Cloudflare response to Ofcom's Interim Report, 5 April 2023, page 2 [Cloudflare \(ofcom.org.uk\)](https://www.cloudflare.com/uk/interim-report)

Analysis of egress costs

- 3.28 We have analysed cost data from cloud providers in order to understand the extent to which egress fees are reflective of costs, and the extent to which pricing differences between egress and other data transfer services (ingress and internal data transfers) are reflective of the difference in the cost of providing egress.
- 3.29 In this section we set out the approach we took to our analysis of the extent to which egress fees are reflective of the costs of providing egress.⁸⁸

Analytical approach

- 3.30 In considering how to assess the extent to which egress fees are reflective of egress costs overall and differences in costs to other data transfer services, we considered that the following could be useful indicators:
- (a) Margin analysis: comparison of total relevant costs incurred in providing egress data transfers to total egress revenues, to determine margins achieved. Whilst we would not necessarily expect margins to be zero in order for egress fees to be broadly cost reflective, we are interested in whether margins may be considered high. As a starting point, we use our estimate of cloud providers' cost of capital for their wider service (ie provision of cloud services)⁸⁹ as an indicative benchmark to compare margins against.⁹⁰
 - (b) Unit cost analysis: comparison of average relevant costs per GB egressed to egress fee list prices (per GB) and average egress revenues per GB. The average revenues and average relevant costs are calculated using the same revenue and cost data as used in (a) above (so the difference between average revenues and costs will be consistent with margins from our margin analysis), but by converting to average \$ per GB rates we are able to compare these to egress fee list prices. List prices and/or average revenues that are higher than average costs may be consistent with egress fees not being reflective of costs.
 - (c) Service comparison: comparison of relevant costs for internal data transfers to egress data transfers, given the common assets used, and costs incurred, in providing these services as set out earlier in this section. Revenue and/or list price differentials between these services being higher than the cost differential may be consistent with egress fees not being reflective of costs.
- 3.31 We consider that use of total and average revenues for comparison to costs allows the application of free tiers and any discounts for egress to be taken into account. When comparing to list prices for the large cloud providers, we have used the prices applicable for data transfers via the internet from London/UK and the lowest non-zero egress fee price in the pricing structure to be conservative. However, we note that comparisons to list prices do not reflect any free tier or customer-specific discounts that apply.
- 3.32 In determining what cost data to use for our analysis, we also consider whether there are adjustments which should be applied to the data as submitted. In the analysis below, we

⁸⁸ A summary of the data used in our analysis and its challenges is set out in Appendix C.

⁸⁹ See Competitive landscape working paper [Cloud services market investigation - GOV.UK \(www.gov.uk\)](https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/674441/Cloud_services_market_investigation_-_GOV.UK.pdf)

⁹⁰ We acknowledge that margins (which have revenue as denominator) are calculated on a somewhat different basis to the returns on capital employed (which have capital employed as denominator) which we benchmark against our estimated cost of capital in our profitability analysis. Revenues and capital employed can be of a different scale, although from our profitability analysis of [3<].

present both costs as submitted and with adjustments to exclude costs which we do not consider to be relevant to assessing the cost justification for egress fees; namely, exclusion of indirect overhead costs (such as sales and marketing and income tax).

- 3.33 We consider these indirect overhead costs to be general business costs which cloud providers have discretion over how to recover.
- 3.34 [REDACTED].⁹¹
- 3.35 We considered applying further adjustments to assess egress margins based on transit costs only, given cloud providers' submissions indicate that this may be the only egress-specific cost. However, given the results of our data analysis on a submitted cost and direct cost basis were already indicating high margins, we have not performed analysis on this basis.
- 3.36 Given the challenges in collecting comparable data from cloud providers, we have also compared transit capacity rates across providers to assess the extent to which incremental costs for egress data transfers are similar or differ between cloud providers. We compare the extent to which transit rates vary between providers to the difference in egress fee pricing structures between the larger cloud providers and others.

Our analysis

- 3.37 This section sets out the analysis we have undertaken to understand the extent to which egress fees are reflective of costs, and the extent to which pricing differences between egress and other data transfer services (ingress and internal data transfers) are reflective of the difference in the cost of providing egress.
- 3.38 We have analysed margins and unit costs of egress data transfer services from one provider; a comparison of revenues and costs for egress data transfer services to internal data transfer services for two others and a comparison of ISP transit rates and list pricing for a number of cloud providers.
- 3.39 We have also supplemented these analyses with data from cloud provider internal documents, where relevant. We are also considering whether there are other practicable benchmarks which may be relevant to consider for our analysis, for example margins for other comparable services provided by cloud providers.
- 3.40 One of the key limitations of the analysis is that we have been unable to complete the same analysis for all the cloud providers. We are also unable to determine if the cost data provided by cloud providers is exactly like-for-like (including in relation to how they treat internal cost recharges).
- 3.41 Despite this, we consider this analysis to be informative given:
- (a) It covers AWS (one of the two largest cloud providers) in detail using each of the identified indicators, and considers both AWS' fully-allocated costs for a conservative approach, and direct costs only for a closer approximation of what we currently consider to be relevant costs.
 - (b) It includes some analysis of Microsoft's costs. As such, both of the two largest cloud providers are included to some extent in the analysis.

⁹¹ [REDACTED] response to the CMA's information request [REDACTED].

- 3.42 We would expect the cost base for cloud providers' networks to be similar.⁹² Given these cloud providers also charge very similar egress list prices, it appears likely that any findings on the cost-reflectivity of egress fees for one provider will also have some applicability to the others.

Margin and unit cost analysis

AWS

- 3.43 AWS does a full cost allocation for its data transfer services. We have used the revenue and cost data obtained from AWS to calculate profit margins and unit costs for its egress data transfer services for 2019 to 2023. We also consider the trend from AWS's forecast revenue and costs for 2024, but give this less weight given forecast data is subject to more assumptions.
- 3.44 We have focused our analysis on egress data transfer services excluding AWS's CDN service, CloudFront. Our understanding is that CloudFront is primarily used for data transfers to end users (eg streaming video content) and is not generally used for customer switching or multi-cloud. As such, we consider the data for egress excluding CloudFront more closely reflects the margins and unit costs that are relevant for switching and multi-cloud egress data transfers and we use egress list prices for data transfers out to the internet in our analysis (not CloudFront list prices).⁹³
- 3.45 AWS's cost allocations for its data transfer services include allocations for fixed and variable AWS and Amazon group indirect costs. It also includes an allocation of income tax. We do not consider it relevant to include income tax in an assessment of profitability, given that income tax will be based on current or prior period profitability for the business overall. We exclude income tax from all of our analysis.
- 3.46 As discussed above, we do not consider it necessarily relevant to include indirect business costs such as sales and marketing, [REDACTED]. As such, we have also calculated AWS egress profit margins using costs it has identified as direct costs only.⁹⁴
- 3.47 AWS egress margins (excluding CloudFront) have been [REDACTED] for 2019 to 2023 under all scenarios.⁹⁵ We consider egress margins excluding indirect costs to be most relevant to our assessment and these have been [REDACTED] for 2019 to 2023.
- 3.48 We have analysed AWS margins on egress data transfer services (excluding CloudFront). This indicates that [REDACTED].
- 3.49 [REDACTED].⁹⁶ [REDACTED].
- 3.50 [REDACTED].⁹⁷

⁹² Google Cloud has 187 Points of Presence (POPs), over 7,500 edge node locations, and is available in over 200 countries and territories. [Global Locations - Regions & Zones | Google Cloud](#), [Expanding our global infrastructure with new regions and subsea cables \(blog.google\)](#) In comparison, AWS has over 600 CloudFront POPs and serves 245 countries and territories [Global Infrastructure - AWS \(amazon.com\)](#) and Microsoft Azure has 192 POPs [Azure CDN POP locations by region | Microsoft Learn](#) and is available in 140 countries and territories [Azure pricing FAQ | Microsoft Azure](#), all accessed 20 May 2024.

⁹³ Note: data transfers out from CloudFront edge locations are subject to different pricing. See [Amazon CloudFront CDN - Plans & Pricing](#), accessed 20 May 2024.

⁹⁴ [REDACTED] [AWS response to the CMA's information request [REDACTED]].

⁹⁵ [REDACTED].

⁹⁶ [REDACTED] AWS response to the CMA's information request [REDACTED].

⁹⁷ [REDACTED] AWS response to the CMA's information request [REDACTED].

- 3.51 We have also analysed AWS average unit costs for egress, to compare these to list prices and average revenues. Whilst using total and average revenue data takes into account free tier egress and discounts applied, the list prices do also reflect the incremental costs that some customers will be paying.
- 3.52 We have analysed AWS average revenues compared to average costs on a \$ per GB basis for egress data transfers, using the same AWS direct cost data (exclusive of indirect costs) as above, and AWS current list prices for data transfers out from Amazon EC2 to the internet, from the London region, after the 100GB/month free tier.
- 3.53 AWS' average revenues per GB from egress have been [redacted] than its average costs per GB for egress. It also shows that AWS's current list prices are [redacted] AWS's average egress costs per GB over the period 2019 to 2023.⁹⁸
- 3.54 As a sensitivity, we also analysed average egress costs including indirect costs. Including indirect costs, the average cost figures (excluding income tax) range from [redacted] per GB between 2019 and 2023 [redacted].
- 3.55 [redacted].
- 3.56 We consider the difference between AWS' unit costs for egress (even when including AWS' indirect cost allocation) and its average revenues and list prices for egress to be consistent with egress fees not being reflective of costs.

Microsoft

- 3.57 Microsoft told us that it does not specifically monitor the cost of egressing data. However, we have identified the following evidence from Microsoft internal documents which we consider to be relevant:
- 3.58 Microsoft's internal transfer prices charged by Azure to other parts of Microsoft's business indicate that for FY24, Microsoft's internal price for egress is [redacted] its egress list prices charged to customers (depending on region and whether routed via Microsoft's network or the ISP network). The internal price for egress from Europe made by other parts of Microsoft's business is [redacted] Microsoft's current first tier of egress list prices charged to customers.⁹⁹ We would generally expect internal transfer prices to follow the arm's length principle, so that prices are comparable to a transaction between two unrelated parties.
- 3.59 Microsoft's 'workload' gross margins, representing units of accountability for the Microsoft engineering teams (but which do not align with particular sets of Azure services), show Networking Services gross margins were [redacted] in FY23, [redacted] in FY22 and 93% in FY21.¹⁰⁰ Whilst Networking Services includes ingress, internal and egress data transfer services as well as other Microsoft networking services (such as CDN, VPN Gateway, Virtual Network, and ExpressRoute private connections), in its response to our clarification request on this data Microsoft recommended viewing the gross margin of Networking Services as a proxy for Bandwidth (ie ingress, internal and egress data transfer services) and Virtual Network gross margins.¹⁰¹

⁹⁸ [redacted].

⁹⁹ The internal transfer pricing appears to be calculated using cost data for long-haul, metro and transit costs plus [redacted] uplift for overheads. Microsoft responses to the CMA's information requests [redacted].

¹⁰⁰ Microsoft includes an allocation for internal use cost recovery [redacted] Microsoft's responses to the CMA's information requests [redacted].

¹⁰¹ Microsoft response to the CMA's information request [redacted].

- 3.60 Additionally, Microsoft has settlement free peering in place with all its UK ISPs. Microsoft said it does not pay any fixed or variable costs for its UK ISP suppliers, and that all costs associated with ISP supply are indirect costs like cabling and hardware.¹⁰² This may indicate that Microsoft's incremental costs for UK egress data transfers are nil, or close to nil, unless there are other interconnection costs incurred which have not been identified.

Google

- 3.61 Google does not track costs at an egress data transfer level. However, we have identified the following data points from Google internal documents which we consider to be relevant:
- 3.62 In a recent internal business case for Google's Cloud networking services, Google showed modelled gross margins on its Cloud Interconnect egress fees of [redacted] based on its 2023 listed pricing, and gross margins of [redacted] based on Google's new pricing taking effect from 1 February 2024.¹⁰³ The document does not state what costs are included in the margin modelling, and they may not be directly comparable to egress fee costs and margins. However, we note that Google's egress fees for Cloud Interconnect data transfers are generally lower than general egress fees to the Internet (Standard and Premium tiers), although additional connection charges apply for Cloud Interconnect usage.¹⁰⁴

Service comparison to transit costs

AWS

- 3.63 AWS submitted that transit costs were only related to egress and not related to transfers of data within the AWS network.¹⁰⁵
- 3.64 Some cloud providers have said that transit costs are also the only incremental cost of providing egress. Therefore, we consider it relevant to compare the difference between egress and internal transfer prices and revenues to the cost of transit as an indicator of whether the relative difference between egress and internal data transfer fees is reflective of costs.
- 3.65 Our analysis indicates that, on a total revenue and average revenue per GB of customer egress basis [redacted]. We consider the difference between AWS's internal data transfer and egress fees revenue to be consistent with egress fee pricing relative to internal data transfer pricing not being reflective of costs.

Microsoft

- 3.66 We also received global transit costs from Microsoft for its cloud infrastructure and have data transfer revenue data for Microsoft for 2021 and 2022 to compare this to.
- 3.67 Our analysis indicates that, on a total revenue basis, [redacted]. We consider the difference between Microsoft's internal data transfer and egress fees revenue to be consistent with egress fee pricing relative to internal data transfer pricing not being reflective of costs.

¹⁰² Microsoft response to the CMA's information request [redacted].

¹⁰³ Google response to the CMA's information request [redacted].

¹⁰⁴ [Pricing | Cloud Storage | Google Cloud](#), accessed 22 February 2024; [Pricing | Cloud Interconnect | Google Cloud](#), accessed 22 February 2024.

¹⁰⁵ AWS response to the CMA's information request [redacted].

- 3.68 We note also that the costs submitted by Microsoft may not include [§], so these transit costs could potentially be overstated, in which case the difference between Microsoft's egress and internal data transfer revenues would be even less cost-reflective.

Transit rates

- 3.69 We also asked cloud providers to provide their transit rates for their largest ISP suppliers, as transit has been identified to be an egress-specific cost and can be incremental per instance of egress data transfer, depending on the nature of the contracts cloud providers have in place.¹⁰⁶
- 3.70 Cloud providers appear to each pay transit rates in the UK which are broadly in the same range.¹⁰⁷ Additionally, some cloud providers appear to pay global transit rate which are broadly around the same range.¹⁰⁸ We are also considering further any differences between global and UK transit rates.
- 3.71 Evidence we have seen to date suggests that, despite having some similarity in transit rates, cloud providers have quite different pricing structures for egress data transfers.

Emerging views on cost of egress fees to providers

- 3.72 This is our initial analysis, and we are continuing to refine our evidence and understanding. We also recognise that AWS is the only cloud provider that has provided us with a detailed cost breakdown by data transfer service, and we are investigating ways to further include Microsoft, Google and/or other cloud providers in our analysis.
- 3.73 Internet transit and peering costs appear to be the main egress-specific cost across cloud providers. It also indicates that the 'burst fees' charged for exceeding committed bandwidth usage are potentially the only cost incrementally incurred per instance of egress.
- 3.74 Our analysis also indicates that egress fees may not be reflective of the costs of providing egress for some providers while others (particularly smaller providers) may have fees that are more cost-reflective.

¹⁰⁶ A number of cloud providers have minimum monthly commitments for bandwidth, with capacity rates that apply for usage in excess of the minimum commitment.

¹⁰⁷ Responses to the CMA's information requests [§].

¹⁰⁸ Responses to the CMA's information requests [§].

4. Potential remedies

- 4.1 In the event that we find that egress fees are a feature that gives rise to an AEC, we are required to decide whether, and if so what, remedial action should be taken to address that AEC.¹⁰⁹ In this section we outline our emerging views on potential remedies relating to egress fees.
- 1.2 We described the CMA's approach to remedies in our issues statement.¹¹⁰ We are at an early stage of considering potential remedies and as our understanding of the market(s) and the potential issues develops, we expect our consideration of potential remedies to evolve. As set out in the CMA's guidance,¹¹¹ we will consider and discuss potential remedies alongside working on understanding what features of the market may give rise to adverse effects. Consistent with this, we set out in this section our early views on potential remedies to any potential AEC(s) relating to egress fees and invite submissions from parties on these to help inform our emerging thinking.
- 4.2 Accordingly, in the rest of this section, we set out the potential remedies we are considering and the views that stakeholders have submitted on potential remedy options. We then provide further detail on the aims and overall approach to the potential remedy options; and set out some key design considerations and questions that we invite submissions on.

Overview of potential egress fee remedies

- 4.3 In our issues statement we set out and sought comment on four categories of remedies which we had identified as potentially relevant to our assessment in relation to egress fees, given the nature of the theories of harm under investigation:
- (a) Preventing cloud providers from charging egress fees, (ie a ban);
 - (b) Capping egress fees by reference to other fees charged by the cloud provider (eg ingress fees or other data transfer fees);
 - (c) Capping egress fees by comparison to the costs incurred by the cloud provider; and/or
 - (d) Increasing the visibility and understanding of egress fees for potential customers, potentially as part of wider requirements to improve the predictability and control spend on cloud.¹¹²
- 4.4 Our issues statement also noted that we are considering the potential for cross-cutting remedies or a package of remedies which would combine to remedy, mitigate or prevent any AECs or their detrimental effects on customers.¹¹³ Our assessment will involve us reviewing potential remedies as standalone egress fee remedies or as components of a package of remedies, for which we would also need to consider any interactions between egress fees and the other components of the package.

¹⁰⁹ Section 134(4) of the Enterprise Act 2002.

¹¹⁰ [Issues statement \(publishing.service.gov.uk\)](https://publishing.service.gov.uk)

¹¹¹ [CMA3 \(Revised\), Market Studies and Market Investigations: Supplemental guidance on the CMA's approach \(publishing.service.gov.uk\)](#), paragraph 3.50.

¹¹² [Issues statement \(publishing.service.gov.uk\)](https://publishing.service.gov.uk)

¹¹³ [Issues statement \(publishing.service.gov.uk\)](https://publishing.service.gov.uk), paragraphs 57 and 58.

- 4.5 There are also cross-cutting design elements which may apply across a number of the potential remedies we are considering, including for egress fees (eg, geography in scope for the purposes of any potential remedies). We will be considering such cross-cutting design elements in a later working paper on potential remedies.

Interactions with regulations outside the UK

- 4.6 Earlier in this paper, we noted that, in response to the EU Data Act, some cloud providers have recently introduced global programmes to provide egress data transfers for switching for free. To the extent that these could impact any egress fee features arising in the UK, we might expect these free switching programmes to impact the relative cost of switching for affected customers, and cloud providers' margins on egress data transfers for switching.
- 4.7 Our consideration of the relevance of these recent developments to our assessment of any AEC is ongoing. However, when finalised, we will consider how this impacts any egress fee features and the need for remedies in relation to this. In considering the role of these programmes in the design of any remedies, we will take account of the scope of the programmes and their voluntary nature including the ability of cloud providers to reverse them.
- 4.8 In that context, we make the following preliminary observations::
- (a) The applicability of these programmes to UK customers could be changed at any point in time: both AWS and Microsoft have noted that they may make changes with respect to the free data transfer policies at any time.¹¹⁴
 - (b) The scope of the programmes indicates that they relate to switching only and not multi-cloud use.
- 4.9 We will consider to what extent, if any, these voluntary commitments, if suitably formalised would be sufficient to remedy any AEC. We welcome any additional submissions on the role of these programmes in designing appropriate remedies, to the extent necessary.

Our analysis of potential egress fee remedies

- 4.10 In this section we set out the rationale for the potential remedies for egress fees that we are considering and summarise the views that cloud providers and other parties have submitted, as well as some of the design and assessment issues that arise.
- 4.11 We structure this section as follows:
- (a) We provide an overview of our overall approach to the three potential remedies which would restrict the level of egress fees (which we refer to as price control remedies), and general views from parties relevant to price control remedies.
 - (b) We then set out parties' views and our current analysis of design considerations for each of the three potential price control remedies:
 - (i) Banning egress fees;

¹¹⁴ [Amazon EC2 FAQs – AWS; Cancel and delete your Azure subscription - Microsoft Cost Management | Microsoft Learn](#), accessed 20 May 2024.

- (ii) Capping egress fees by reference to other fees charged by cloud providers; and
 - (iii) Capping egress fees by reference to the costs incurred by cloud providers.
- (c) We then set out our current thinking on information transparency remedies.

Restricting the level of egress fees: price control remedies

Aims and approach

- 4.12 The primary aim of a price control remedy would be to provide an ongoing restriction on prices so that the cost to customers of data egress does not make it more costly and difficult for customers to choose the best value offers for them and deter them from switching or using multiple clouds.
- 4.13 We would expect price control remedies to reduce the cost to customers of switching or using multi-cloud, under current market pricing structures. We note that large cloud providers provide a free tier for many egress services, provide ingress data transfers for free, and generally provide internal data transfers at lower prices than egress.
- 4.14 The three price control options are mutually exclusive. However, each option could work on its own or in combination with the other remedy measures to help reduce customer detriment by reducing costs (absolute and/or relative to within-cloud usage) as a barrier to switching and multi-cloud.

Design considerations

- 4.15 Remedies which control outcomes, such as price controls, need to specify the products or services that are subject to control and the basis for the level of the price that is set.¹¹⁵
- 4.16 There are a variety of ways in which a price control that constrains the price of egress data transfers can be specified. We are continuing to consider the following key parameters related to the design of any remedy option that controls the price of egress fees:
- (a) How the level of any price cap should be determined;
 - (b) Which cloud providers should be subject to any price cap;
 - (c) Which of the egress data transfer services provided by cloud providers should be included within the scope of any price cap;
 - (d) What should be the duration for any price cap and the basis for any sunset provision; and
 - (e) How any price cap could be designed to mitigate any unintended distortions to market outcomes that may arise from the introduction of price regulation.
- 4.17 In the following section we set out our initial thinking on these parameters for each of the three price control options being considered, taking into account the views submitted by parties in response to our issues statement, submissions from cloud providers and other

¹¹⁵ CC3 (Revised), [Guidelines for market investigations: Their role, procedures, assessment and remedies](https://publishing.service.gov.uk/guidelines-for-market-investigations) (publishing.service.gov.uk), Annex B paragraph 87.

market participants in response to information requests, and the initial analysis set out earlier in this working paper.

General comments from parties relevant to egress fee price control remedies

- 4.18 Some cloud providers opposed setting a fixed price or cap for egress fees, on the basis that it would be unfair to customers (with smaller usage customers subsidising larger customers)¹¹⁶ and could reduce investment or innovation in network infrastructure.¹¹⁷
- 4.19 AWS said that compensating cloud providers for the substantial cost of transferring data and building and operating a complex and expensive network is the only way to facilitate large-scale data transfer of the kind that customers demand.¹¹⁸
- 4.20 Google said that it does not see data transfer charges as a barrier to switching or multi-cloud. Google said data transfer fees allow cloud providers to recover costs of network infrastructure usage and support its ongoing efforts to invest in the further expansion of its global network infrastructure and to develop innovative networking solutions to benefit customers.¹¹⁹
- 4.21 Google said setting an egress fee price or cap at the same level across cloud providers would not recognise that cloud providers compete on parameters including quality and innovation.¹²⁰
- 4.22 Microsoft said that an element of charging for data transfers is important for cost recovery and efficient usage of cloud infrastructure.¹²¹
- 4.23 IBM said that if the costs incurred by a cloud provider are not recovered in one form, they will likely be recovered elsewhere, in order to allow the provider to achieve an economic return.¹²² IBM also said that its experience is that customers are more focused on technical barriers than on the cost of egress fees.¹²³
- 4.24 The UK Competitive Telecommunications Association (UKCTA), a trade association for UK fixed line telecommunications companies, said that excessive egress charges can act to deter consumers of cloud services from exercising optimal choices in their solution design and raise the costs associated with switching suppliers. However, UKCTA was of the view that cloud providers should be given the opportunity to remedy concerns around egress fees themselves in the first instance, volunteering a solution that seeks to preserve consumer choice and improve overall market flexibility with the minimum level of disruption and cost.¹²⁴

¹¹⁶ Google suggests that requiring all data transfers to be priced at the same level removes the ability for cloud providers to offer differently priced data transfer services (eg premium and CDN services) to cater to different customer needs. AWS suggests that large customers would be subsidised both as a result of the likely need to increase other cloud fees to cover costs if egress fees were banned (impacting smaller customers' broader cloud spend), and if a flat fee was set for unlimited data transfers, regardless of size of transfers.

¹¹⁷ [AWS response to the Issues Statement](#) dated 17 October 2023, page 14; [Google response to the Issues Statement](#) dated 17 October 2023, pages 8-9.

¹¹⁸ [AWS response to the Issues Statement](#) dated 17 October 2023, page 14.

¹¹⁹ [Google response to the Issues Statement](#) dated 17 October 2023, page 7.

¹²⁰ [Google response to the Issues Statement](#) dated 17 October 2023, page 8.

¹²¹ [Microsoft response to the Issues Statement](#) dated 17 October 2023, page 9.

¹²² [IBM response to the Issues Statement](#) dated 17 October 2023, page 3.

¹²³ [IBM response to the Issues Statement](#) dated 17 October 2023, page 4.

¹²⁴ [UKCTA response to the Issues Statement](#) dated 17 October 2023, page 1.

- 4.25 Another cloud provider said that cloud providers should increase interconnectivity by connecting their data centres to other cloud providers' data centres. It said that this was the most efficient way to exchange traffic and avoided alternatives with a higher cost and lower assurance over reliability/performance, but that there should be choice over interconnection mechanisms and there should be a mechanism to ensure that costs are reasonable.¹²⁵
- 4.26 In contrast, OVHcloud said that interconnecting data centres (eg through PNI agreements) can reduce cloud provider costs for data transfers between providers by avoiding bandwidth costs but is not an adequate option to foster multi-cloud strategies due to the current lack of interoperability between cloud provider services. It noted that connecting infrastructures of different cloud providers requires a stack of software solutions on top of the services provided.¹²⁶
- 4.27 No additional remedies were proposed in relation to egress fees by respondents to our issues statement.

Banning egress fees

Cloud providers' views

- 4.28 Some cloud providers support the complete removal of egress fees.¹²⁷ They provided the following rationales for this view:
- (a) Oracle said that cloud providers should not be competing on data transfer fees as it considers that the value it provides to customers is based on its service offerings. It also noted that as data remains the property of the customers, the customer should be readily able to move their data among various cloud services and providers.¹²⁸
 - (b) OVHcloud said as it considers egress fees to represent artificial costs and to be unduly used by the largest cloud providers to lock customers in their ecosystems and prevent them from switching.¹²⁹
- 4.29 The largest cloud providers said removing egress fees could result in a free-rider problem, where customers have inefficient excess egress usage because there is no cost to them for egress.¹³⁰
- 4.30 Large cloud providers said the charges for other cloud services were likely to increase if egress fees were removed.¹³¹
- 4.31 AWS said that requiring egress to be free was uneconomic because cloud providers could not provide customers with the reliable, fast, and secure global network that they require

¹²⁵ [redacted] response to the CMA's information request [redacted].

¹²⁶ OVHcloud submission [redacted].

¹²⁷ OVHcloud submission [redacted]; Oracle says data mobility fees should be zero. [Oracle response to the Issues Statement](#) dated 17 October 2023, page 5.

¹²⁸ [Oracle response to the Issues Statement](#) dated 17 October 2023, page 5.

¹²⁹ OVHcloud submission [redacted].

¹³⁰ [Microsoft response to the Issues Statement](#) dated 17 October 2023, page 14; [AWS response to the Issues Statement](#) dated 17 October 2023, page 14.

¹³¹ [Microsoft response to the Issues Statement](#) dated 17 October 2023, page 14; [AWS response to the Issues Statement](#) dated 17 October 2023, page 14; [Google response to the Issues Statement](#) dated 17 October 2023, page 8.

without being able to recoup the significant expenses associated with building and maintaining such a network.¹³²

- 4.32 Google said that banning egress fees would lead to unintended consequences, including reduced transparency of costs (as they might be replaced by higher charges for other services), unfair outcomes for customers and reduced innovation in the development of network infrastructure.¹³³
- 4.33 A cloud provider also disagreed with banning egress fees, but on the basis that smaller providers and new entrants to the cost-intensive IaaS market do not have the requisite capital or scale to cross-subsidise and offer egress for free as it suggested the large cloud providers do. It said that 'An egress ban therefore risks entrenching hyperscalers' oligopoly leading to less competition.'¹³⁴

Our emerging views

How the level of this form of price cap should be determined

- 4.34 This potential remedy would result in a ban on egress fees. The scope of who and what this ban could apply to is discussed below. This potential remedy is analytically straightforward from the perspective of determining a methodology for setting the price.
- 4.35 This remedy would completely remove egress fees costs as a barrier to switching and multi-cloud. As part of our evaluation of this remedy option we will consider the materiality of relevant egress costs to cloud providers which may need to be absorbed if they can no longer recover these costs through egress fees.

Which cloud providers should be subject to this form of price control

- 4.36 The scope of a ban on egress fees could be market-wide. However, the scope of the remedy in terms of the providers captured will be dependent on any AEC finding.
- 4.37 We note that in our issues statement we indicated that we were considering whether any barriers created by egress fees could lead to cloud providers entrenching their positions. To the extent this is reflected in a feature of any AEC, we would consider whether the scope of any remedy should apply to all providers or be based on a threshold criteria, for example including a cloud provider with public cloud infrastructure services revenues above a certain threshold.
- 4.38 Additionally, a market-wide remedy might be detrimental to smaller cloud providers' ability to compete if they were to lose access to a profitable revenue stream and were unable to recover any costs (egress-related or otherwise) through egress fees. This could potentially make investment to build scale more challenging.

Which services should be in scope

- 4.39 As set out above, we are considering whether, and to what extent, egress fees may deter customers from switching and using multiple clouds, and are therefore focussed on

¹³² [AWS response to the Issues Statement](#) dated 17 October 2023, page 14.

¹³³ [Google response to the Issues Statement](#) dated 17 October 2023, pages 8-9.

¹³⁴ [36] response to the CMA's information request [36].

egress data transfers for switching and multi-cloud use.¹³⁵ Egress data transfers out of a cloud provider's infrastructure can take place via the internet, direct connection, CDN or physical device transfer. We are primarily interested in transfers via the internet as we understand that these are more likely to be the route for switching or multi-cloud than CDNs or physical device transfers.

- 4.40 Our understanding is that direct connections are currently used more often for customer transfers to on-premises infrastructure, however they can also be used for data transfers between clouds. Direct connections also generally have different pricing structures to data transfers via the internet, with lower egress fees but other charges in addition such as connection charges. We are considering whether egress fees for direct connections should be included in the scope of any potential remedies. We will also consider whether there may be potential unintended consequences if direct connections are not included (for example, potential price increases to recoup lost revenues for internet egress).
- 4.41 We are considering ways that cloud providers can identify egress data transfers specifically for switching or multi-cloud, given cloud providers have said that they cannot identify the purpose of customers' data transfers.
- 4.42 Customer self-nomination/attestation appears to be the most accurate means of doing so, however it is not automatic and relies on customers being proactive. It may also place some administrative burden on cloud providers if customer checks are considered necessary (and friction if these checks create new challenges for customers).
- 4.43 A more automated but potentially less accurate proxy may be the identification data associated with data transfers. Several cloud providers said they may be able to identify the peer to which data is transferred using the Border Gateway Protocol (BGP) peer Autonomous System Number (ASN) for the transfer.¹³⁶ However, they noted issues including:
- (a) The ASN for a data transfer may not reflect the ultimate use or end destination for the data, for example it may be sent to another cloud network for on-premises hosting services or to be transferred to end users and/or 'eye-ball networks' (those used primarily for browsing the internet or consuming content). In addition, not all networks have an assigned ASN, in which case a cloud provider may be unable to identify the end destination network.
 - (b) The ASN may be for a network/company that provides cloud and non-cloud services and it is not possible to identify which business unit or subsidiary the data is transferred to.
- 4.44 We consider it might be possible to create technical workarounds to some of these issues for ASNs, such as asking the customer whether they are transferring data to another cloud, but we also recognise that such workarounds may be complex and might themselves give rise to unintended consequences, for example potentially preventing or

¹³⁵ Note, as set out in paragraphs 1.35 to 1.44 above, the EU Data Act contains provisions that set requirements on egress data charges for switching and multi-cloud use for EU customers (with different requirements and timing for taking effect between these two use cases). We note that only requirements in relation to switching charges have already taken effect under the EU Data Act, and we understand that the EU Data Act requirements in relation to multi-cloud use only take effect on 12 September 2025. We also understand that the global free switching programmes that AWS, Microsoft and Google have introduced apply to switching egress but not multi-cloud egress.

¹³⁶ Responses to the CMA's information requests [38].

limiting customers' ability to do automated egress usage, which could potentially deter multi-cloud usage.

- 4.45 We note that Microsoft said it can identify traffic to another cloud or CDN network only if it has direct peering (as AWS and Google Cloud do). Such traffic data shows usage of the peering service at an aggregate level, but Microsoft cannot identify which customer is transferring and what amount of data is being transferred by a specific customer. Microsoft said that customers transfer data to various destinations routing through a third-party ISP. Identification to an identifiable location, company or method of transfer is only possible if a customer uses private peering over ExpressRoute or VPN for a given customer.¹³⁷
- 4.46 One option would be to include all egress data transfers in the scope of any potential remedy, until such a point that providers can accurately identify egress related to switching and multi-cloud. This would mean that egress that is not related to switching or multi-cloud would initially be included within the scope.

What should be the duration for any price cap

- 4.47 We may consider to specify a limited duration (for example, by means of a long-stop date in a 'sunset clause') for individual remedy measures, for example where these are designed to have a transitional effect. This might occur if we expect any AEC to be time-limited, or if a particular element of a remedy package is intended to be a temporary arrangement to deliver improvements in the short term, while other longer-term measures take effect.¹³⁸
- 4.48 Should we find an AEC in relation to egress fees and we considered it likely to endure, then we may require a price control remedy such as a ban for a longer duration. However, we also note that price controls may give rise to distortion risks, if retained over a long period, and will not generally be preferred as a long-term solution.¹³⁹ As such it may be appropriate to include a long-stop date – either a sunset clause or a review - for a price control remedy such as a ban.

How price cap design could mitigate any unintended distortions

- 4.49 We will consider the potential for unintended consequences. For a price control remedy such as a ban. At this early stage of our remedies thinking, we have focused on identifying these potential risks and unintended consequences.
- 4.50 The large cloud providers have said that preventing them from charging egress fees may lead to increased prices for other cloud services.¹⁴⁰ This could impact a broader set of customers than the customers that currently incur egress fees.
- 4.51 We set out earlier in this paper the proportion of UK public cloud revenue accounted for by egress fees.¹⁴¹ Given our initial analysis of the extent to which egress revenues are higher than the associated costs, this would indicate that any necessity to recover costs through other services would be below this level, reducing the risk of any distortions.

¹³⁷ Microsoft response to the CMA's information request [38].

¹³⁸ CC3 (Revised) paragraph 339.

¹³⁹ CC3 (Revised) paragraph 389.

¹⁴⁰ Microsoft response to the Issues Statement dated 17 October 2023, page 14; AWS response to the Issues Statement dated 17 October 2023, page 14; Google response to the Issues Statement dated 17 October 2023, page 8.

¹⁴¹ These revenues include egress data transfers via CDN, which would not necessarily be included within the scope of any price control remedy.

- 4.52 Large cloud providers also said that removing egress fees could result in inefficient excess egress usage by customers because there is no incentive to them to minimise their usage.¹⁴² However, the cloud providers have not provided any evidence to substantiate this. We therefore invite evidence from the providers that addresses the following:
- (a) How likely it may be that customers would continue to have an incentive to design their IT architectures efficiently and avoid significant volumes of unnecessary egress, for example for latency, availability, security and/or data governance reasons.
 - (b) How easy or difficult it may be for a customer to inadvertently accumulate a high volume of egress via the internet.
 - (c) Whether there is any existing evidence of free egress leading to inefficient egress usage, given that the large cloud providers currently provide a free tier of egress per month.

Capping egress fees by reference to other fees charged

Parties' views

- 4.53 [redacted].¹⁴³ AWS said it incurs less cost for internal data transfers than for data transfers out to the Internet.¹⁴⁴
- 4.54 In response to our issues statement, AWS also said that a flat fee for unlimited data transfers, regardless of size, would force smaller customers to subsidise larger customers' use of network resources through their broader cloud spend, and that other hypothetical alternatives that would ultimately force cloud providers to charge for transferring data into their clouds would make switching more costly, and would impose new data transfer costs on customers who currently don't pay any data transfer fees.¹⁴⁵
- 4.55 OVHcloud stated that it does not support capping fees by reference to other fees charged. It said while this might remedy the issue of large cloud providers allocating common costs to egress in order to justify egress fees, it would not address switching costs, could have indirect adverse effects such as overall fee increases, and would require a very complex compliance and monitoring process.¹⁴⁶
- 4.56 A CDN provider supported capping egress fees at the costs incurred with porting data into the cloud, ie ingress fees, on the basis that there is no justification for cloud providers charging more for data transfers out than in.¹⁴⁷

Our emerging views

How the level of this form of price cap should be determined

- 4.57 This remedy would be intended to reduce egress fees' potential significance and impact on decision-making for customers when considering switching or multi-cloud use. It would

¹⁴² [Microsoft response to the Issues Statement](#) dated 17 October 2023, page 14; [AWS response to the Issues Statement](#) dated 17 October 2023, page 14.

¹⁴³ AWS response to Ofcom's Interim Report, paragraph 18 [Amazon Web Services \(ofcom.org.uk\)](#)

¹⁴⁴ AWS response to the CMA's information request [redacted].

¹⁴⁵ [AWS response to the Issues Statement](#) dated 17 October 2023, page 14.

¹⁴⁶ OVHcloud submission [redacted].

¹⁴⁷ [redacted] submission [redacted].

also be intended to allow cloud providers to continue to determine the extent to which they recover costs through data transfer fees, but remove potential deterrence to switching or using multi-cloud arising from relative cost differences to customers between transferring data within the same cloud or between different clouds.

- 4.58 The intended impact of this remedy would be to set egress fees at a level equivalent to another comparable data transfer service, so that there is no cost difference to the customer between either (i) transferring their data within their existing cloud provider's service as compared to another cloud provider's service, or (ii) having their data enter their cloud provider's infrastructure as compared to having data exit it.
- 4.59 Given the commonality in assets and costs set out in Section 3 above, for a remedy capping egress fees by reference to other fees charged, we consider the following to be the most relevant reference point fees to consider:
- (a) Ingress fees, on the basis that ingress data transfers follow the same available routes into cloud provider networks as for egress data transfers out of the cloud via the internet, and cloud providers have said ingress and egress generally use the same fixed assets. We note that the large cloud providers all currently provide ingress for free.
 - (b) Inter-region or inter-continent internal data transfer fees, on the basis that the journey of such data transfers is comparable to egress data transfers. Our understanding is that egress data transfers routed via cloud providers' own network (eg Microsoft Global Network, Google Cloud Premium Tier) will follow the same path as inter-region or inter-continent transfers, with just a final mile exit of the cloud network to the transfer destination. As set out in the previous section, cloud providers have said that internal data transfers have common fixed assets and costs with egress but have said there are some egress-specific costs such as transit and other public internet connectivity costs.
- 4.60 We consider intra-zone, inter-zone and/or intra-region internal data transfer fees to be less relevant reference products given that these relate to transfers over shorter distances, whereas our understanding is that egress data transfers often cover longer distances. We are considering whether the egress fees charged for direct connection solutions by cloud providers is another relevant reference point.

Which cloud providers should be subject to this form of price control

- 4.61 The same considerations as for a ban on egress fees are applicable here.
- 4.62 We also note that designing a market-wide remedy for capping egress fees by reference to other charges is also likely to be more complex compared to a remedy for specified cloud providers, given the range of business models and data transfer charges the market will encompass.

Which services should be in scope

- 4.63 The same considerations as for a ban on egress fees are applicable here.

What should be the duration for any price cap

- 4.64 The same considerations for a ban on egress fees are applicable here.

How price cap design could mitigate any unintended distortions

- 4.65 Much of the same considerations for a ban on egress fees are applicable here. However, we also note a potential remedy that would cap egress fees by reference to another non-zero fee would still allow cloud providers to earn some revenues (and hence, recover some costs) for egress data transfers.
- 4.66 For an egress fee cap by reference to other cloud fees, there is also a risk of unintended consequences as egress fees become set relative to another price, rather than specified based on a certain value or methodology. For example, to the extent that the reference fee is meant to represent a reasonable cost measure, there is a risk the price cap results in a distortion of the reference fee. This risk likely increases if the value of egress increases relative to the value of the reference product. There may also be additional complexity in choosing a reference product that is differently defined and/or priced across cloud providers (although we note that ingress fees do not currently have this issue).
- 4.67 There is also potentially a greater circumvention risk (compared to a ban), given cloud providers may have the ability to change their pricing structure for the reference fee in a way that means there is, in effect, limited change in the cost of egress data transfers to customers (overall, and/or relative to ingress or within cloud data transfers).

Capping egress fees by comparison to costs incurred

Parties' views

- 4.68 AWS said that banning the possibility to earn profit on data transfers out to the Internet will act as a disincentive for cloud providers to innovate and invest to improve this service for customers. It also said, if in the form of price controls, it will be costly and complex to implement due to the variation in transfer costs depending on time, volume and location.¹⁴⁸
- 4.69 One cloud provider said that setting egress fees to costs means fees will fluctuate and offer customers less certainty.¹⁴⁹ The cloud provider also said that capping egress fees at cost may also promote inefficient use of providers' networks by some customers, which could put excess strain on networks and increase cloud providers' costs, which will be passed on via increased fees for all customers.¹⁵⁰
- 4.70 OVHcloud said capping egress fees at cost could be a temporary and transitional remedy towards complete removal but should only be transitional as it does not consider egress fees to be justified by economic or technical reasons. OVHcloud also suggested that a remedy capping egress fees at cost would only ensure workable competition if there was a strict scope of application to costs that are directly and solely related to the switching process.¹⁵¹
- 4.71 An ISV was supportive of egress fees better reflecting actual network costs incurred by cloud providers.¹⁵²

¹⁴⁸ AWS response to Ofcom's Interim Report, para 18 [Amazon Web Services \(ofcom.org.uk\)](https://www.ofcom.gov.uk/consult/condocs/amazon/amazonweb/amazonweb.pdf)

¹⁴⁹ [redacted], submissions to the CMA [redacted].

¹⁵⁰ [redacted], submissions to CMA [redacted].

¹⁵¹ OVHcloud submission [redacted].

¹⁵² [Company A response to the Issues Statement](#) dated 17 October 2023, page 3.

How the level of this form of price cap should be determined

- 4.72 For a cap on egress fees linked to the costs incurred by the cloud provider, the previous section sets out the evidence we have collected on the costs that cloud providers have said are relevant to providing data transfers. The intended impact of this remedy would be to set egress fees at a level equivalent to recovery of relevant costs for cloud providers' provision of egress data transfers.
- 4.73 We consider that there are a range of options for how a cap could be designed in principle. These include:
- (a) issuing cost guidelines, which set out the types of costs which are suitable to include in determining egress fees and relies on cloud providers to implement;
 - (b) a negotiation/arbitration model, whereby parties would make submissions on egress fee costs to an oversight body to reach an agreed view on the egress fee cap, and a third-party arbitrator may be used if agreement cannot be reached; and
 - (c) setting egress fees to a cost-based price level that we would determine. There are a range of methodology options for this approach, such as cost benchmarking across companies or calculating company long run incremental cost.
- 4.74 We are in the early stages of our analysis and remedies assessment and will consider effectiveness and proportionality further as our work progresses. However, we note that for an egress fee cap at cost, there may be trade-offs between the design complexity required for implementation and the ongoing monitoring and enforcement required.
- 4.75 For example, using cost guidelines (as described above) to set egress fees at cost may be easier to implement as less detail and data is needed, however it may create challenges for monitoring and/or enforcement due to it being less specific on an egress fees value.
- 4.76 Conversely, capping egress fee rates at cost via a CMA-determined price level may require less ongoing monitoring and enforcement but would involve substantial complexity to implement in a way which is effective and minimises distortion risk, and the availability and comparability of cost data between cloud providers is likely to be a challenge.
- 4.77 We note that capping egress fees at cost would allow cloud providers to recover relevant costs included in the cap. Evidence we have seen to date indicates that there is a significant amount of commonality in the assets required and costs incurred in providing ingress, internal data transfer and egress. As set out in the previous section, transit costs are generally the only egress-specific cost commonly identified by cloud providers and other stakeholders, and our understanding is that transit costs comprise a relatively small proportion of networking costs for large cloud providers.¹⁵³ However, cloud providers did have differing submissions on the extent to which there are other types of costs which are egress-specific, and what costs incurred in providing data transfer services are incremental vs fixed.
- 4.78 We will consider further the extent to which costs are common (and therefore more subject to business decisions as to how these costs are recovered) as opposed to egress-

¹⁵³ CMA analysis of transit costs: responses to the CMA's information requests [38].

specific, and fixed as opposed to incremental, and how to address differences in submissions, when considering relevant costs for any of the design options above.

Which cloud providers should be subject to this form of price control

- 4.79 The same considerations for a ban on egress fees are applicable here.
- 4.80 We also note that designing a market-wide remedy for capping egress fees at cost is also likely to be more complex compared to a remedy for specified cloud providers, given the range of business models and data transfer charges the market will encompass.

Which services should be in scope

- 4.81 The same considerations for a ban on egress fees are applicable here.

What should be the duration for any price cap

- 4.82 The same considerations for a ban on egress fees are applicable here.

How price cap design could mitigate any unintended distortions

- 4.83 Much of the same considerations for a ban on egress fees are applicable here. However, we also note a potential remedy that would cap egress fees at cost would still allow cloud providers to earn some revenues (and hence, recover some costs) for egress data transfers.
- 4.84 There is also a risk for any of the cost cap options for egress fees that it risks becoming distortive or redundant as networking practices and cost efficiencies change over time.

Other considerations

- 4.85 In relation to the potential for cross-cutting remedies or a package of remedies, we recognise that there may be some interactions between egress fees and other features we are investigating. For example, overall customer discounts (such as committed spend agreements) are applicable to egress fees for a number of cloud providers,¹⁵⁴ so any potential remedies for committed spend agreements, if an AEC arising from these was found, could have interactions with the significance of egress fees to customers.
- 4.86 We will consider interactions between potential remedies for different potential features further as our thinking on cross-cutting remedies or a package of remedies develops.

Potential for relevant customer benefits

- 4.87 Some cloud providers have submitted that egress fee revenues or profits fund and incentivise investment and innovation in data transfer services, and that price controls for egress fees would remove or reduce the incentive to invest in providing high quality networks.¹⁵⁵

¹⁵⁴ Responses to the CMA's information requests [3<].

¹⁵⁵ [3<] submission to the CMA [3<], [3<] submission to the CMA [3<]; [Google response to the Issues Statement](#) dated 17 October 2023, paragraphs 26 and 30(c).

- 4.88 To date, we have seen limited evidence of customers receiving lower prices or greater choice or innovation for data transfer services as a result of investment funded by egress fee revenues.
- (a) Google cites its Cross-Cloud Interconnect service, launched in June 2023, as an example of an innovative new product.¹⁵⁶ However, customers are largely responsible for preliminary set-up of this network connection themselves.¹⁵⁷ Google will assist in the connection process and provide management services on an ongoing basis.
- (b) Microsoft introduced a lower cost internet egress service in March 2021 which routes egress traffic via the public internet,¹⁵⁸ however this does not appear likely to have required significant network investment or innovation, given it is a service based on public internet transit supplied by third parties rather than the use of Microsoft's own global network.
- 4.89 Some cloud providers have submitted that they make significant investments in building networks (which some characterised as high quality or 'premium') for providing data transfer services.¹⁵⁹
- 4.90 However, it seems likely that any such benefits to customers would accrue irrespective of egress fee revenues given that their high quality networks are used to provide other cloud services such as internal data transfers.

Information transparency remedies

- 4.91 The fourth potential remedy we set out in our issues statement was an increase in the visibility and clarity of egress fees for customers, potentially as part of wider requirements on providers to improve the predictability of, and customers' ability to control, their spend on cloud. For example, this could be a requirement for cloud providers to adopt consistent terminology in relation to egress fees, or display egress fee prices prominently on webpages and/or in contracts. This would not tackle the level of such fees directly, but might help to support customer decision-making when choosing a provider or considering switching business away from a provider.
- 4.92 We are continuing to consider the evidence on the extent to which cost predictability impacts on customer decision making. In that context, it is still uncertain whether an information transparency remedy would be needed to remedy any concern on this basis.
- 4.93 Furthermore, we are considering whether an information transparency remedy would be necessary to improve the effectiveness of any price regulation remedy. An information transparency remedy could work in combination with some form of price control to help improve the effectiveness of the price regulation, by helping to make customers aware of pricing changes. This could be particularly important if a price regulation remedy relied on customer self-nomination due to the inability to identify the purpose of data transfers. We note that, as with the price control remedies set out above, we would consider whether the scope of any information transparency remedy should be limited to specific cloud providers.

¹⁵⁶ Google response to the CMA's information request [3<]. [Announcing Google Cloud Cross-Cloud Interconnect | Google Cloud Blog](#), accessed 22 February 2024.

¹⁵⁷ [3<] response to the CMA's information request [3<].

¹⁵⁸ Microsoft response to the CMA's information request [3<].

¹⁵⁹ [3<] response to Ofcom's Interim Report [3<]; Responses to the CMA's information requests [3<].

Summary of potential remedies and invitation to comment

- 4.94 Our thinking on remedies continues to develop, however as set out above, we have further developed our thinking on the approach and design of the potential remedies set out in our issues statement.
- 4.95 We have considered views provided by parties so far but welcome further submissions from parties on the thinking set out above. Additionally, as mentioned above, we will be considering such further remedy design elements in a later working paper on potential remedies.
- 4.96 We would particularly welcome views on the following questions:
- (a) How should we define the scope and duration of any egress fee remedies?
 - (b) To what extent should voluntary commitments made by some cloud providers to provide free switching egress globally be taken into account in our remedies assessment? Given eligibility requirements, customer self-nomination and applicability to switching but not multi-cloud usage, to what extent may these commitments impact UK customers?
 - (c) What would be the estimated costs to cloud providers and/or other impacted parties of implementing each of the potential egress fees remedies?
 - (d) Are there any alternative remedies that would be as effective as those set out in this paper in addressing any barriers to switching and/or multi-cloud for customers arising from egress fees, and that could be less costly and/or intrusive?
 - (e) Could ASNs be a reasonable proxy for determining whether an external data transfer has gone to another cloud provider? Are there ways that ASNs could be feasibly supplemented (eg a step in the data transfer process that requires customers to select their purpose from a menu) to more accurately determine the purpose of a data transfer?
 - (f) In relation to the potential remedy of capping egress fees at cost, how should the relevant costs be determined, and cap enforced? What types of costs should or shouldn't be included in determining a cap at cost?
 - (g) In relation to the potential remedy of capping egress fees by reference to other fees charged by cloud providers, what would be an appropriate reference fee for setting the cap and why?
 - (h) Are there any relevant customer benefits in relation to egress fees that we should consider as part of our assessment of a remedy package?¹⁶⁰
 - (i) Would a complementary information transparency remedy be required in order for a price control remedy to be effective?

¹⁶⁰ Responses to this question should bear in mind the criteria we use for assessing claimed RCBs and that parties will be expected to provide convincing evidence regarding the nature and scale of any RCB that they claim to result from the market feature(s) concerned and to demonstrate that these fall within the Act's definition of such benefits. See [CC3 \(Revised\), Guidelines for market investigations: Their role, procedures, assessment and remedies \(publishing.service.gov.uk\)](#) paragraph 359.

Appendix A: Cloud providers' cost of egress

A.2 In this section we summarise the types of assets and costs that cloud providers' have identified as used and incurred in providing egress data transfer services.

Assets cloud providers use to provide data transfer services

A.3 Cloud providers have identified similar assets used in providing data transfer services – data centres and network equipment (including servers, routers and switches), and connection cables (short-haul and/or long-haul). The key network assets identified are summarised below. We note that some of these may also be used in providing other cloud services such as storage.

Table 4.1: Summary of key network assets used by cloud providers in transferring data

Networking assets	
Servers	Servers are used for receiving and transmitting data. Servers can be origin servers or edge servers – origin servers store the original version of the content, while edge servers store cached versions of the content.
Routers	Routers connect computers and other devices to the Internet. A router acts as a dispatcher, choosing the best route for information to travel. Routing is the ability to forward IP packets – a package of data with an internet protocol (IP) address – from one network to another. Boundary routers forward ingress traffic to internal routers and switches.
Switches	Internal switches send traffic to the cloud computing and storage resources for processing. The return egress traffic that contains customer data is traversed via internal switches to the boundary routers to go to the Internet or to the backbone network. Ethernet switches connect cabled devices, like computers, WiFi access points and IoT devices, and servers, in an Ethernet Local Area Network so they can communicate with each other and to the Internet. Once a device is connected to a port, the Ethernet switch manages the flow of data between the device and other devices, applications, data cloud services, and the Internet.
Optic fibre cables: Backbone	Backbone, also known as long-haul fibre, are optic fibre cables providing connectivity between regions and Points of Presence (POP) around the world. They may be used for inter-region connectivity, providing private direct connection connectivity to clouds, and connectivity to IP transit and peering across regions.
Optic fibre cables: Metro	Metro fibre, also referred to as short-haul fibre, are optic fibre cables which provide connectivity over shorter distances, eg region to region.

Source: Responses to the CMA's information requests [34]

Costs cloud providers incur in providing data transfer services

- A.4 Cloud providers identified a range of costs that are incurred in providing data transfers for customers. The costs identified can broadly be summarised as:
- Costs related to having network infrastructure (eg costs to build or lease data centres and other fixed infrastructure)
 - Costs of connectivity (to the Internet, or to other clouds, third parties or direct customer connections)
 - Shared business/overhead costs, eg sales & marketing, HR, legal
- A.5 Network infrastructure costs can vary substantially between providers depending on the extent to which they invest in building their own backbone network or rely on transit (provided by Internet Service Providers (ISPs) or Network Service Providers (NSPs)) to transfer data over the public internet.
- A.6 However, generally any provider will have some data centre presence which incurs associated costs. Cloud providers cite fixed and variable data centre costs (utilities, wages, rent etc) and costs for network equipment (depreciation and maintenance costs) as data transfer costs.¹⁶¹
- A.7 Connectivity costs are incurred to connect network locations within a cloud and to external parties, and to transfer data out of a cloud. This can include costs to connect networks to edge POPs such as metro (short-haul) cable costs; transit/bandwidth costs (for transfers over internet); and port charges for physical interconnections of the network to third party paid or free transit/peering. We set out connectivity costs in more detail in Appendix B.
- A.8 Cloud providers say they incur different costs based on location/peering provider for transfers and routing.¹⁶²
- A.9 Some cloud providers cite shared and/or overhead business costs such as sales and marketing, customer service, HR, IT and/or R&D as costs of providing data transfer.¹⁶³ However, other cloud providers do not cite general business costs such as sales and marketing as costs of data transfer in their responses.

Cloud provider submissions on egress-specific costs

- A.10 When asked what costs were specific to egress, ingress or internal data transfers:

¹⁶¹ Responses to the CMA's information requests [3<]; [3<] response to Ofcom's information request [3<].

¹⁶² Responses to the CMA's information requests [3<]; [3<] response to Ofcom's information request [3<].

¹⁶³ Responses to the CMA's information requests [3<].

- AWS submitted that, of the [X] costs it allocates to egress, [X] are also relevant to internal transfers within the AWS network.¹⁶⁴
- Microsoft said that for egress it incurs incremental interconnection costs for data transfers over the public internet.¹⁶⁵
- Google identified edge networking infrastructure (eg leased space and power at colocation data centres and points of presence, paid-peering), and the dedicated physical connection installed between GCP and the destination infrastructure for Google Interconnect products as being specific to external data transfers.¹⁶⁶
- IBM identified data centre interconnectivity costs as specific to internal data transfers, and the costs of public internet connectivity from each data centre as specific to egress.¹⁶⁷
- Oracle said it does not typically incur separate costs for ingress, egress, or transferring customer data, but also said it uses both 'core' and edge assets for external data transfers (ingress and egress) whereas internal data transfers only use 'core' assets.¹⁶⁸

Cloud provider submissions on incremental costs

A.11 In relation to the incremental costs of providing egress data transfers to customers, cloud providers identified the following:

- Burst fees: Some cloud providers identified excess fees charged for bandwidth usage above committed capacity levels (also called 'burst fees') as incremental costs for providing egress. They said that if monthly data transfer rates exceed their minimum commitments, then an additional charge is calculated based on an incremental rate.¹⁶⁹ [X] said that if monthly data transmission rates are within the committed limits, then zero incremental and variable costs are incurred related to data egress.¹⁷⁰
- Microsoft identified the incremental cost of interconnection over the public internet for egress.¹⁷¹ It did not specify what these incremental costs were.

¹⁶⁴ AWS response to the CMA's information request [X].

¹⁶⁵ Microsoft response to the CMA's information request [X].

¹⁶⁶ Google response to the CMA's information request [X].

¹⁶⁷ IBM response to the CMA's information request [X].

¹⁶⁸ Oracle to the CMA's information requests [X].

¹⁶⁹ [X] response to the CMA's information request [X].

¹⁷⁰ [X] response to the CMA's information request [X].

¹⁷¹ Microsoft response to the CMA's information request [X].

- IBM identified hardware maintenance and energy consumption as incremental costs incurred when transferring data.¹⁷²
- A cloud provider [§<] data centre operating expenses are only incrementally incurred when attributable to a specific product.¹⁷³
- In contrast, one cloud provider said capital investment in infrastructure should be included as a cost of incremental added capacity and that it does not consider network and infrastructure costs as either fixed or variable costs, as many of these costs have both fixed and variable components.¹⁷⁴ The cloud provider noted that an incremental measure of one unit out of billions of egress gigabytes transferred every year would not be relevant from its business perspective. The cloud provider said it needs to make investment decisions and incur costs to grow the network in a way that can safely and reliably accommodate “incremental” customer egress needs of several billions of gigabytes every year.¹⁷⁵
- One cloud provider argued that inbound data transfers are a one-time cost that can be incorporated into the price for storing the data, whereas this is not the same for outbound data as some customers may not plan to transfer their data out or will transfer data millions of times daily.¹⁷⁶
- One cloud provider said it was unable to estimate its incremental costs. It said it pays [§<] that includes a mix of fixed and incremental components.¹⁷⁷ However, the cloud provider provided information indicating that not all costs being recovered via egress fees are incrementally incurred per egress transfer or specific to egress.¹⁷⁸

¹⁷² IBM response to the CMA's information request [§<].

¹⁷³ [§<] response to the CMA's information request [§<].

¹⁷⁴ [§<] response to the CMA's information request [§<].

¹⁷⁵ [§<] response to the CMA's information request [§<].

¹⁷⁶ [§<].

¹⁷⁷ [§<] response to the CMA's information request [§<].

¹⁷⁸ [§<] response to the CMA's information request [§<].

Appendix B: Cloud providers' connectivity arrangements

- B.1 In this section we summarise the types of connectivity used by cloud providers and related costs identified by cloud providers.
- B.2 There are three main approaches for cloud providers to establish connectivity between workloads running across different infrastructures:
- (a) Over the public internet, using ISPs
 - (b) Through connectivity partner networks, using NSP services
 - (c) Direct connectivity at a colocation data centre facility
- B.3 Internet peering allows cloud providers to exchange traffic directly with other networks and can be done as private peering, using private network interconnections (PNIs), and public peering, through internet exchanges.
- B.4 'IP Transit', via connection to the public internet, allows cloud providers to exchange traffic indirectly with other networks. One cloud provider said this may be used because the other network is not otherwise available for peering or may be for redundancy purposes. IP Transit providers, by connecting with each other to send and receive data, offer access to every network on the internet (as opposed to peers, who offer access only to their own networks and those of their customers).¹⁷⁹
- B.5 Google said private peering (which could be paid or settlement free) is the primary way it exchanges traffic with other cloud provider networks, major ISPs and CDN providers and is purpose built to exchange high volumes of traffic reliably.¹⁸⁰
- B.6 Some cloud providers said each participating member is responsible for the costs of connecting to third party networks, the public internet and/or public peering.¹⁸¹ [redacted].¹⁸² Microsoft said any connection supporting connectivity to another cloud provider's network is shared by both parties equally.¹⁸³ Google said costs incurred in peering between cloud providers are typically shared, but may be negotiated between providers when establishing the connection.¹⁸⁴
- B.7 Oracle said its public network is connected to [redacted] public networks either directly or via internet exchanges (IXs). Oracle has a commercial agreement with Microsoft for connecting OCI and Azure. The interconnect allows customers to build a

¹⁷⁹ [redacted] response to the CMA's information request [redacted].

¹⁸⁰ [redacted] response to the CMA's information request [redacted].

¹⁸¹ Responses to the CMA's information requests [redacted].

¹⁸² [redacted] response to the CMA's information request [redacted].

¹⁸³ Microsoft's response to the CMA's information request [redacted].

¹⁸⁴ Google's response to the CMA's information request [redacted].

private interconnection between their OCI and Azure environments.¹⁸⁵ In both cases, each party pays their own hardware and operational costs (eg data centre space/ power/ cabling, labour).¹⁸⁶

B.8 [REDACTED].¹⁸⁷

Connectivity costs

B.9 On the costs of connecting to the internet and other networks, the cloud providers identified the following types of arrangements with ISPs:¹⁸⁸

- (a) Settlement-free peering: Both parties' exchange traffic freely with no traffic associated fees.
- (b) Paid peering: Cloud provider pays the peering partner a price per mbps based on the amount of traffic exchanged between the networks. That traffic could terminate into the customers of that ISP network.
- (c) IP Transit: Cloud provider pays the transit provider a price per mbps based on the amount of traffic exchanged between the networks. That traffic could terminate anywhere in the world.
- (d) Internet exchange (IX) ports: Cloud provider and peering partner will exchange traffic through IX ports and the cloud provider and peering partner will pay the IX provider's port fee.

B.10 Microsoft cites the following as costs specific to connecting to another cloud provider or any other ISP network: cost of relevant hardware, access infrastructure, cross connects, internet exchange port fees. Microsoft's interconnection arrangement with Oracle means that costs such as cross connect, access fees and internet exchange port fees are not applicable to this service.¹⁸⁹

B.11 Microsoft said it has settlement-free contracts for all its UK ISPs, so it does not pay any fixed or variable costs to ISPs in the UK.¹⁹⁰

B.12 One cloud provider said most of its internet peering is settlement free.¹⁹¹ Another cloud provider said internet peering arrangements with ISPs are mostly on a settlement-free basis.¹⁹² Cloudflare, a CDN that competes with cloud providers in

¹⁸⁵ <https://www.oracle.com/cloud/azure/interconnect/>, accessed 20 May 2024.

¹⁸⁶ [REDACTED] response to the CMA's information request [REDACTED].

¹⁸⁷ [REDACTED] response to the CMA's information request [REDACTED].

¹⁸⁸ Responses to the CMA's information requests [REDACTED].

¹⁸⁹ Microsoft response to the CMA's information request [REDACTED].

¹⁹⁰ Microsoft response to the CMA's information request [REDACTED].

¹⁹¹ [REDACTED] response to the CMA's information request [REDACTED].

¹⁹² [REDACTED] response to the CMA's information request [REDACTED].

offering data transfers to end users, has an open peering policy and peers with more than 12,000 networks around the world on a settlement-free basis.¹⁹³

- B.13 Additionally, the Bandwidth Alliance¹⁹⁴ is a group of cloud providers and networking companies that have agreed to discount or waive data transfer fees for shared customers, for data transfers routed via Cloudflare's network (Cloudflare itself does not charge egress fees). Microsoft and Google allow discounted egress fees within the Bandwidth Alliance, whilst the other group members (such as Oracle, Alibaba, Tencent and Scaleway) provide free egress for shared customers.¹⁹⁵

¹⁹³ [The Cloudflare Blog](#), accessed 26 February 2024.

¹⁹⁴ [The Cloudflare Blog](#), accessed 27 February 2024.

¹⁹⁵ [Bandwidth Alliance | Reduce Data Transfer Fees | Cloudflare](#), accessed 22 February 2024.

Appendix C: Data available for analysis of egress costs

- C.1 In this section we summarise the data used in our analysis of egress costs and its challenges.
- C.2 We collected cost data from cloud providers (including [REDACTED]), on an average costs basis and/or based on the cost reporting used within each business. However, provision of cost data at an egress or total data transfer product level has been challenging for most cloud providers due to the nature of their reporting.
- C.3 AWS does a full cost allocation exercise for its data transfer services [REDACTED].¹⁹⁶ We gathered a cost breakdown as well as revenues and volumes for AWS' data transfer services covering 2019 to 2024 (forecast).¹⁹⁷
- C.4 Other cloud providers said they do not track costs and/or volumes for some types of services at the level of data transfer services.¹⁹⁸ We have included some data from Microsoft internal documents in our analysis, but have otherwise been unable to include other cloud providers in most of our analysis.
- C.5 We have gathered some data on ISP transit rates and costs from a number of cloud providers. However, we note that transit is contracted on a capacity basis, with \$ per mbps rates typically applying. In contrast, egress data transfers are charged to customers based on volume at \$ per GB rates. These volume-based rates are not directly comparable to capacity rates and assumptions for monthly utilisation and application of peak usage percentiles would be required to convert capacity rates into per GB rates.
- C.6 We also note that cloud providers' network infrastructure is used not only to provide customer data transfers, but is also used internally by other business units. Our understanding is that cloud providers apply cost recharges to internal teams for data transfer and/or broader network usage, but there may be differences in how these recharges are determined and recorded in internal financial reporting, and providers' internal recharge policies will impact the level of costs attributed to data transfers services and/or their cloud businesses overall.
- C.7 We focused on gathering global data on egress given that (i) costs are often incurred at a global or regional level (for example, contracts with ISPs for

¹⁹⁶ Responses to the CMA's information requests [REDACTED].

¹⁹⁷ Note: in [REDACTED] internal reporting for data transfer services it also shows tracking of 'non-monetised' costs, which are not included in the cost data provided and used to analyse AWS egress margins. [REDACTED] said that non-monetised values are associated with usage that is not directly monetised by [REDACTED], ie, values which are indirectly monetised by [REDACTED] services such as [REDACTED]. [REDACTED] response to the CMA's information request [REDACTED]. These include, for example, USD per instance per hour values for compute services, and USD per GB per month values for storage services.

¹⁹⁸ [REDACTED] [REDACTED] response to the CMA's information request [REDACTED], [REDACTED]; [REDACTED]. [REDACTED] response to the CMA's information request [REDACTED]; [REDACTED] [REDACTED] response to the CMA's information request [REDACTED]; [REDACTED] [REDACTED] response to the CMA's information request [REDACTED].

bandwidth capacity may be global or for regions such as North America and/or Europe) and as such costs at a country level are often subject to business discretion as to how they are allocated; and (ii) egress revenues can be impacted by the basis on which they are determined (eg based on customer billing location, or by data centre location for data transfers) and as such may not be comparable to the basis on which costs are allocated.