




Cloud Services Market Investigation Qualitative Customer Research

Final Report

May 2024



Contents

- 1. Executive summary..... 4**
- 1.1. Background 4
- 1.2. Research approach 4
- 1.3. Context 5
- 1.4. Theories of Harm..... 7
- 1.5. AI and Future Outlook 11
- 2. Background, objectives and approach 12**
- 2.1. Background to the research 12
- 2.2. Focus of this research 12
- 2.3. Research objectives 13
- 2.4. Research target audience 14
- 2.5. Research method 15
- 2.6. Sample structure 18
- 2.7. Methodological note on qualitative research..... 20
- 3. Main findings – Context..... 22**
- 3.1. Overview 22
- 3.2. Public cloud decision making 22
- 3.3. Why public cloud? 25
- 3.4. Reasons for choosing individual public cloud providers 26
- 3.5. Public cloud operating models 34
- 3.6. Switching 38
- 3.7. Customer examples 40
- 4. Main findings – Theories of Harm 1: Technical interoperability and portability 44**
- 4.1. Overview of technical barriers 44
- 4.2. Implications of technical challenges 51
- 4.3. Exacerbating factors..... 54
- 4.4. Potential enablers for interoperability and portability 58
- 4.5. Conclusions 64
- 5. Main findings – Theories of Harm 2: Data egress and transfer fees 66**
- 5.1. Overall attitudes towards data egress and transfer fees..... 66

5.2.	Egress fees impacting multi-cloud or switching behaviour	69
5.3.	Conclusions	72
6.	Main findings – Theories of Harm 3: Contracts with committed spend discounts	74
6.1.	Overview of different types and levels of discounts	74
6.2.	Implications/impact on behaviour	79
6.3.	Conclusions	82
7.	Main findings – Theories of Harm 4: Software licensing practices	83
7.1.	Overview of software licensing practices	83
7.2.	Implications/impact on behaviour	85
7.3.	Conclusions	88
8.	Main findings – Future outlook	89
8.1.	Artificial Intelligence (AI).....	89
8.2.	Future changes to the cloud market.....	91
9.	Appendix – Field materials.....	94
9.1.	Letter of authority	94
9.2.	Recruitment screener	96
9.3.	Discussion guide	107

1. Executive summary

1.1. Background

- 1.1.1. Public cloud infrastructure services were referred to the CMA for a market investigation by Ofcom, following Ofcom's market study, which identified features that make it difficult for UK businesses to use multiple cloud infrastructure providers and to switch between them.
- 1.1.2. The CMA is investigating the supply of public cloud infrastructure services in the UK. The CMA is considering whether there are competition concerns or not, and if there are what interventions could address those concerns.
- 1.1.3. The CMA's issues statement outlines the proposed initial scope of the CMA's investigation, drawing on the three features of the market Ofcom was previously most concerned about, plus the impact, if any, of software licensing practices on cloud competition.
- 1.1.4. The CMA's investigation is examining:
 - a. Whether **technical barriers** are making switching providers and setting up a multi-cloud architecture difficult for customers and whether these barriers are contributing to the lock-in of customers to a single provider
 - b. Whether **egress fees** act as a barrier to switching and multi-cloud, and hinder competition between cloud providers
 - c. Whether the way certain **discounts** by existing cloud providers are structured acts as a barrier to switching and multi-cloud as well as a barrier to entry and expansion among cloud providers in a way that leads to harm to competition.
 - d. The nature of the **software licensing** practices of the relevant cloud providers and whether these practices disincentivise customers from using rival cloud providers or raise barriers to entry in the market
- 1.1.5. Cloud services are typically classified according to their service models: infrastructure as a service (IaaS), platform as a service (PaaS) and software as a service (SaaS)
- 1.1.6. For the purposes of the CMA's market investigation, cloud infrastructure services include IaaS and PaaS cloud services delivered through public cloud computing.
- 1.1.7. This investigation does not focus on SaaS or on private cloud computing, where computing resources are dedicated to (as opposed to shared between) individual customers.

1.2. Research approach

- 1.2.1. The over-arching research objective for this project was to deepen the CMA's understanding of UK cloud customers' decision-making on their choice of public cloud infrastructure services, their multi-cloud and switching behaviour, and future use of these services.
- 1.2.2. The target audience for this research was medium to large UK public cloud customers, who operate or trade in the UK, paid for IaaS and PaaS cloud services in 2022, and used one of the five largest cloud providers in the UK – AWS (Amazon Web Services), Microsoft, Google, IBM or Oracle, which provide AWS, Microsoft Azure, Google Cloud, IBM Cloud and Oracle Cloud respectively.

- 1.2.3. The goal was to speak to key decision-makers on public cloud services within these businesses to understand their choices in their current role and also draw on their wider expertise and experiences.
- 1.2.4. Jigsaw carried out 60 in-depth interviews with individuals from 50 companies. In 10 companies, where a wide range of individuals were involved in decision-making on public cloud use within the business, two interviews were undertaken to understand different points of view, for example, to gain insight into both strategic decisions (involving significant change or spend, such as migrating to public cloud or switching IaaS provider) and more tactical day-to-day decisions (such as choice of PaaS provider or how to optimise services). Interviews took place from 29th January 2024 to 14th March 2024.
- 1.2.5. Participants were recruited from public cloud customer lists provided by the CMA. The CMA requested the customer lists from the five main public cloud providers to UK businesses – AWS, Microsoft, Google, IBM and Oracle. Companies were selected at random by the CMA from those received, totalling 766 initial contacts provided to Jigsaw, followed by an additional 529 names to complete the target number of interviews required within the timeframe available. A wide range of different individuals were interviewed from the lists, indicative of the overall sample provided and spread across the different public cloud providers, different business sectors and a variety of public cloud spend levels.
- 1.2.6. Qualitative research focuses on understanding underlying issues, experiences, and motivations rather than providing percentages or numbers. The findings are indicative of the market rather than representative. They are based on participants' recall, understanding, and interpretation post-purchase, which can be imprecise and to a degree post-rationalised.
- 1.2.7. Participants in this research have provided their personal views as individuals in their role as an IT professional making decisions about public cloud infrastructure services for their organisation. Each response should be interpreted as being from an individual experienced customer rather than as an official collective submission from their organisation. Individuals who have taken part have done so on the basis that they will remain anonymous, and their details and the company they work for will not be disclosed.

1.3. Context

- 1.3.1. As context, it is helpful to understand the general views and motivations expressed by decision-makers on public cloud use as a backdrop to the discussion around choices and theories of harm:

Public cloud decision-making

- 1.3.2. Approaches to decision-making and considerations around public cloud vary by types of business.
- 1.3.3. Start-ups tend to have tight cost constraints, will work at pace and with less formal decision-making processes, and are more likely to test and switch public cloud provider whilst the cost and risk to operations is low. Established businesses have more formal decision-making structures, and are less likely to switch public cloud provider for IaaS or PaaS services unless there is a strong business case to do so, as they perceive this would be complex, expensive and present a high risk to operations.
- 1.3.4. Those in highly regulated sectors such as Finance and Health, or who are higher spenders on public cloud also face more constraints in terms of cost and risk than those in less regulated sectors or who are lower spenders. IT businesses can often have more flexibility because the level of IT expertise in-house is high and they can more easily adapt or create their own bespoke solutions.

Why public cloud?

- 1.3.5. Participants perceive multiple benefits of public cloud compared to private on-premises or private cloud. For most, the aspiration is to have all their business infrastructure on public cloud, but cost or reduced control present a challenge for a few use cases, including use of GPUs or processes that involve transferring large volumes of data.

Individual public cloud providers

- 1.3.6. The main providers are seen as AWS, Microsoft and Google among participants. For some, this is the main or only consideration set in terms of who might even be on a shortlist of providers in the event of a review of the market or a switch, though most were not aiming to make any changes. They each have an excellent reputation, are seen to deliver a reliable service and offer a wide range of solutions that cover many needs.
- 1.3.7. AWS is perceived as the overall leader, in terms of the level of innovation and range and quality of services provided, followed by Microsoft, and both are used by those working in a wide variety of types of businesses. Google is also thought to be highly reputable, but has a smaller pool of users among those interviewed. Google users in our sample are more likely to be tech-driven businesses such as start-ups, e-commerce, fintech or adtech companies.
- 1.3.8. Many felt that services provided by these main players were '80% the same', but there were distinctions perceived e.g. some see AWS as the first mover and innovator, Microsoft as having superior IAM (identity and access management), and Google as having a strength in analytics (specifically, Google's analytics product BigQuery as having better performance).
- 1.3.9. Very few interviewees use IBM as a sole provider and none use Oracle in this way. Their main use was secondary, for example, for supporting legacy systems. The types of businesses using these providers are also more traditional, such as those involved in energy, health, regulatory support or insurance.
- 1.3.10. A few participants use other smaller providers, such as Digital Ocean or Lambda Labs, for compute services only, as and when needed as a 'commodity'. A few also use them for niche benefits e.g. Digital Ocean as it felt easy for one start-up before it scaled up and Nebius for its green credentials.

Public cloud operating models

- 1.3.11. The main operating models among those interviewed are to use a single public cloud provider for their current company infrastructure or two to three cloud providers with siloed workloads. Other models do exist alongside these for particular use cases, the most mentioned being integrating use of BigQuery from Google or Microsoft IAM alongside infrastructure from a different cloud provider as these services are seen as superior to those offered by their main provider.

Switching

- 1.3.12. Many businesses aspire to having their infrastructure on public cloud. Migrating to public cloud from on-premises or considering whether to host via public or private cloud were key strategic decisions. Once a public cloud provider(s) was put in place, many participants felt less need to switch all or part of their infrastructure to an alternative provider(s), as they cannot see a clear rationale.

- 1.3.13. Switching cloud providers is seen as the equivalent of moving other kinds of infrastructure, such as 'moving house' or moving a business from one country to another. It is not something to undertake lightly or consider at all unless it leads to significant business benefits long term that override the perceived cost and risk of changing. To an extent, dependency on a current provider(s), or a sense of 'lock in', is a factor across all providers as change brings cost and risk.
- 1.3.14. Few of those interviewed had chosen to switch between public cloud providers. Where people have switched either all or some of their infrastructure, this was driven by an anticipated reduction of cost or clear increase in performance. In some cases switching was a requirement following a merger or acquisition to rationalise providers.
- 1.3.15. Switching can be a straightforward process of a few weeks for start-ups, or take up to months or years for more established or complex businesses, followed by a long process of learning and consolidation to optimise and increase efficiency. The process of switching typically involved running two services concurrently, then switching off the legacy system when ready. A few of those who had switched needed to employ third parties to advise and implement.
- 1.3.16. Overall, those who had switched felt it had not only brought cost and operational risk, but took IT staff away from core work and typically ended up being more challenging and time consuming than anticipated.

1.4. Theories of Harm

Theories of Harm 1 – Technical Interoperability and Portability:

- 1.4.1. When asked about technical barriers to switching cloud service provider or using multiple clouds research participants tended to make very general and high-level comments about the overall complexity of re-engineering cloud architectures.
- 1.4.2. The most commonly mentioned tech barriers relate to the portability of certain code environments, specific data formats that are incompatible across providers, difficulties in migrating databases and file storage systems, as well as networking and connectivity challenges due to inconsistent APIs and interfaces.
- 1.4.3. Participants voiced particular concern around data latency, security gaps and the potential for a switch or multi-cloud strategy to absorb so many resources that it would disrupt the IT department's ability to service the rest of the business.
- 1.4.4. Exacerbating factors to these technical challenges are a major skills gap (both in their current team and a difficult hiring environment) and a resource gap, i.e. the engineering resources that switching or multi-cloud use would take away from their core business.
- 1.4.5. While some enablers were mentioned – both technical and service-based – participants don't appear to see them as completely offsetting the technical challenges to switching or using multiple clouds.
- 1.4.6. The most commonly mentioned lever to overcome vendor lock-in was the increasing use of vendor-agnostic services.
- 1.4.7. However, some participants voiced concern around this strategy, which they say may lead to restricting themselves access to cutting-edge innovation provided by proprietary vendor services. A small number of participants even went as far as saying that encouraging greater open-source standards may discourage cloud providers to innovate if they're incentivised to prioritise the 'lowest common denominator' over the most innovative solution. However, this was a minority opinion.

- 1.4.8. Overall, the research clearly shows that there are very significant technical barriers that reduce participants' willingness to consider switching cloud provider or going for a multi-cloud strategy. However, most participants mentioned these barriers in the context of *hypothetical scenarios* in which they would seriously consider a switch or multi-cloud approach. As this research shows, most participants do not see a strong argument in favour of a switch or a multi-cloud strategy in the first place.
- 1.4.9. Furthermore, as some switchers or multi-cloud users in the sample demonstrate, these barriers can be overcome in practice, with enough time, effort and resources. However, there doesn't appear to be sufficient incentive for most participants to invest this time, effort and resources to switch provider or implement a multi-cloud architecture.
- 1.4.10. It is also important to note that, while participants expressed concerns about specific technological challenges, their implications and shortcomings of potential solutions, there is also a more general fear of the 'unknown'. From the research it appears that many cloud users anticipate unforeseen issues if they were to attempt to multi-cloud or switch provider, which introduces risk into the process. That anticipation of unforeseen issues alone can act as a deterrent to multi-cloud or switching.
- 1.4.11. In summary, for most participants, the tendency to reject the idea of multi-cloud or switching provider is based on the view that the benefits of such a strategy are relatively low but the technical barriers standing in the way are relatively high. This view is compounded by the fact that switching or multi-cloud is rarely top of their mind, so participants often have only a hypothetical view of potential benefits while at the same time fearing the anticipated as well as the unforeseen challenges that would have to be overcome.

Theories of Harm 2 – Data egress and transfer fees:

- 1.4.12. Egress fees were rarely top of mind when participants were asked about their openness to switching cloud provider or using multiple clouds, or even when specifically asked about potential challenges for a switch of multi-cloud strategy.
- 1.4.13. In fact, those participants in the research who did switch or use multiple cloud providers for specific use cases, considered the egress fees they incurred a price worth paying to deliver the cloud strategy that makes most sense for their business.
- 1.4.14. However, for some companies, egress fees do disincentivize them from potentially considering a switch to another provider. In particular, companies with high data volumes or a large amount of historical data can be put off by a potentially large one-time 'exit fee'.
- 1.4.15. When asked about the role of egress fees for a multi-cloud strategy, some participants were concerned about the potential costs of keeping databases in sync. While this is usually mentioned in the context of a *hypothetical* multi-cloud scenario, this consideration does suggest that egress fees might act as a barrier or deterrence to even considering a multi-cloud strategy.
- 1.4.16. In a similar vein, even limited multi-cloud use cases such as using a different cloud provider for backup solutions may be discouraged by egress fees.
- 1.4.17. In addition to the amount of egress fees, participants also voiced concern about the lack of transparency and overall complexity of egress fees, which makes it difficult for some to predict costs for their general cloud usage. This is relevant in the switching and multi-cloud context, as the lack of transparency and control makes it difficult for participants to include egress fees in a potential cost-benefit calculation for switching or using multiple cloud providers.

- 1.4.18. In most cases, however, egress fees are rarely the main reason why participants don't consider a multi-cloud strategy or migration to another provider.
- 1.4.19. In summary, egress fees do have the potential to affect cloud users' willingness to switch or multi-cloud; however, the technical barriers, alongside the exacerbating factors and concerns mentioned in section 4, as well as the perceived lack of a strong argument in favour of switching or using multiple clouds, tend to weigh heavier than any potential concern around egress fees.

Theories of Harm 3 – Contracts with committed spend discounts:

- 1.4.20. A wide range of discounts are on offer from public cloud providers. These range from high level committed spend discounts (the most common type mentioned) through to committed use discounts, all the way to specific discounts offered on a service-by-service basis.
- 1.4.21. The research suggests a relationship exists between the size of business/cloud usage strata and the type of financial or contractual relationship companies are prepared to enter into with their cloud provider(s).
- 1.4.22. Larger companies spending heavily with cloud providers (and often with more complex IT infrastructure needs), almost always report having some sort of discounting arrangement in place. Very often these arrangements were reported as being multi-layered and complex, including CSDs (committed spend discounts), CUDs (committed use discounts) or Reserved Instances and other types of discounts.
- 1.4.23. Users thought that this complexity might to some extent be designed to optimise their spend so that the discounting arrangement is based around their actual usage. However, many participants report that it is difficult to assess whether their discounting arrangement is, in fact, the best way to use services offered by their cloud provider. This is due to the complexity of billing and ongoing contracts sometimes not aligning with each other in terms of timing.
- 1.4.24. This reported complexity would also make direct cost comparisons between providers difficult, if any of the participants wanted to do this. However, very few said they did want to do this as they were not looking to switch, and if they were, it would not be a purely financial decision, as many other risk factors come into play.
- 1.4.25. At the other end of the market, smaller, low spend users of the cloud do report switching providers purely based on the financial incentives offered to them. These come in the form of credits – free usage periods or monetary incentives.
- 1.4.26. Companies who are beyond the start-up phase, but who are still in a fast growth phase of their development are more likely to commit to a CUD or Reserved Instance than a CSD as they felt less able to predict their future usage.
- 1.4.27. Finally, a consequence of the contractual discounts in place is that, for some participants, they are clearly driving behaviour in terms of how the company uses the public cloud services they pay for. Minimum spend targets or discounts based on specific services/marketplace usage can and do drive users to potentially use parts of the cloud service that they would not otherwise use.
- 1.4.28. In summary, the research shows that the existence of discounts is widespread within the public cloud market and that these discounts can be seen as complex in how they are put together within contracts. The research also demonstrates that the existence of discounts does have an impact on switching behaviour, although that impact does not exist in isolation from other factors that would be considered if switching cloud providers was being considered.

Theories of Harm 4 – Software licensing practices:

- 1.4.29. It is very clear that original take up of Microsoft Azure is often closely related to an organisation's pre-existing use of Microsoft products and services.
- 1.4.30. The decision to go with Azure as their public cloud provider has often been taken several years ago (over five, sometimes more than ten years ago). The decision to rely on Microsoft's suite of products and services has often been taken decades before the decision to go with Azure. In that sense, participants often struggled to precisely detail why their organisation uses Azure, beyond describing it as a natural choice for both technical and financial reasons. As with the other Theories of Harm, individual factors don't necessarily affect participants' attitudes towards multi-cloud or switching significantly on their own, but rather they add up and shape multi-cloud or switching behaviour in cumulative way.
- 1.4.31. This makes the role of software licensing in the decision to go with Azure difficult to unpick, but participants did not single out licensing as a key factor on its own influencing their decision.
- 1.4.32. Participants were clearer when discussing enterprise agreements with Microsoft and their impact on any potential decision to switch away from Azure or move to a more integrated multi-cloud solution. Here, participants were generally clear that the breadth and depth of their relationship with Microsoft, as evidenced by the existence of an enterprise agreement, means that they struggled to even justify reviewing their Azure usage, so embedded is it within their whole IT approach.
- 1.4.33. The sample does also contain a number of organisations who use Microsoft for all their software and some PaaS needs, but who have made a decision to go with a public cloud provider other than Azure to be their strategic partner. These organisations tend to either be smaller than those with Microsoft enterprise agreements in place, or they have structured their IT in such a way that the IT infrastructure is sufficiently separated from development work so that pure interoperability is less of an issue.
- 1.4.34. Even then, companies using Microsoft who have chosen not to go with Azure, made the decision to do so very early on in their public cloud journey. In that sense, they too have now become embedded with the provider they chose originally.
- 1.4.35. The research did not find many organisations who were Microsoft users, who had initially gone with Azure, and who had then switched away to another provider.
- 1.4.36. We did also speak to a small number of Google Cloud Platform users who had chosen to use GCP because of its interoperability with Google's productivity suite. This highlights the attraction of 'ease of use' as a reason for choosing cloud provider, emphasising that matching software to cloud provider can be purely based on this, and not be at all influenced by any software licensing agreements.
- 1.4.37. In summary, the customers interviewed who used Microsoft nearly always talked about their public cloud choice, and any impact of software licensing, in the context of already being Microsoft users. If an organisation has a historical, strong relationship with Microsoft, it is likely (though by no means certain) that they will also be Azure users. The direct connection between software licensing practices and the decision to go with Azure is generally not recognised by Azure users, beyond a clear general desire to keep everything within a Microsoft ecosystem within their organisation.

Cumulative effect

- 1.4.38. Across the four Theories of Harm, this research explored a wide range of potential barriers to multi-cloud and switching cloud providers (as well as some enablers). As explained in this executive summary, and detailed further in the main findings below, there are a number of factors that clearly show signs of disincentivising participants to consider switching cloud provider or considering a multi-cloud strategy. However, it is important to note that the impact each of the individual barriers has on the participants' cloud strategies, in most cases, ought to be judged in combination with other barriers, be they of technical, financial or contractual nature (and this 'combination of factors' often also includes the lack of a strong business case for multi-cloud or switching to begin with).
- 1.4.39. In other words, the individual barriers have a cumulative effect on participants' willingness to switch provider or adopt a multi-cloud strategy. The extent to which each barrier weighs into these considerations depends on individual circumstances. However, this research sheds light on the specific mechanics through which the individual barriers affect cloud strategies. The main findings provide the detail for each Theory of Harm.

1.5. AI and Future Outlook

- 1.5.1. Artificial Intelligence (AI) features widely in the interviews, especially when participants are talking about future public cloud trends. Expectation is generally that the main public cloud providers will (if they are not already) offer an AI PaaS solution, and that they will compete strongly for this AI PaaS business. Concern does also exist that once an AI solution is chosen by a customer, it will be difficult for that customer to switch away from, creating another opportunity for cloud providers to potentially lock customers into their cloud offering.
- 1.5.2. At the end of the interviews, participants were asked for any suggested areas where they would like to see changes to the public cloud market. These ideas were rarely top of mind or strongly expressed – they were in response to direct questions, and came at the end of an hour-long interview where participants had been given the chance to organise their thoughts in a way that they would not normally do in their day-to-day roles. 4 key areas were mentioned by participants:
- a. Simplify the complex billing arrangements
 - b. Removal, standardisation or reduction in egress fees
 - c. Light touch regulation or independent operating guidelines for cloud providers
 - d. Improvements in services that improve the operation of workloads across clouds – E.g. Kubernetes, containerisation and other abstraction tools

2. Background, objectives and approach

2.1. Background to the research

- 2.1.1. Public cloud infrastructure services were referred to the CMA for a market investigation by Ofcom, following Ofcom's market study, which identified features that make it difficult for UK businesses to use multiple cloud infrastructure providers and to switch between them.
- 2.1.2. The CMA is investigating the supply of public cloud infrastructure services in the UK. The CMA is considering whether there are competition concerns or not, and if there are what interventions could address those concerns.
- 2.1.3. The CMA's issues statement outlines the proposed initial scope of the CMA's investigation, drawing on the three features of the market Ofcom was previously most concerned about, plus the impact, if any, of software licensing practices on cloud competition
- 2.1.4. The CMA's investigation is examining:
 - a. Whether **technical barriers** are making switching providers and setting up a multi-cloud architecture difficult for customers and whether these barriers are contributing to the lock-in of customers to a single provider
 - b. Whether **egress fees** act as a barrier to switching and multi-cloud, and hinder competition between cloud providers
 - c. Whether the way certain **discounts** by existing cloud providers are structured acts as a barrier to switching and multi-cloud as well as a barrier to entry and expansion among cloud providers in a way that leads to harm to competition
 - d. The nature of the **software licensing** practices of the relevant cloud providers and whether these practices disincentivise customers from using rival cloud providers or raise barriers to entry in the market

2.2. Focus of this research

- 2.2.1. Cloud services are typically classified according to their service models: infrastructure as a service (IaaS), platform as a service (PaaS) and software as a service (SaaS)
- 2.2.2. For the purposes of the CMA's market investigation, public cloud infrastructure services include **IaaS and PaaS cloud services delivered through public cloud computing**.
- 2.2.3. This investigation does not focus on:
 - a. SaaS
 - b. Private cloud computing, where computing resources are dedicated to (as opposed to shared between) individual customers.
- 2.2.4. Further detail on technical definitions is provided in the panel.

Public cloud infrastructure services – definitions used in the research

- 2.2.5. **Cloud computing** is defined as the provision of remote access to computing resources (compute, storage and networking) on demand and over a network (public internet or a private connection), instead of a personal computer or local server that are not part of the cloud. The customer buys access to the computing resources as a service and typically does not own the underlying hardware and software.
- 2.2.6. **Public cloud computing** means a cloud deployment model where cloud services are open to all customers willing to pay, and computing resources are shared between them.
- 2.2.7. Cloud services are typically classified according to their service models: infrastructure as a service (IaaS), platform as a service (PaaS) and software as a service (SaaS):
- a. **Infrastructure as a service (IaaS)** are cloud services that provide access to raw computing resources for processing workloads and storing data. These computing resources are in the form of servers and networking equipment owned and managed by the IaaS provider (and typically held on racks in a remote data centre). To allow and manage that access, IaaS also includes some necessary software, including networking (e.g. firewall) and virtualisation. Examples of IaaS offerings include Amazon Web Services (AWS) EC2, Microsoft Azure Virtual Machines and Google Compute Engine – which can be used by business customers, for example, to store data and install software.
 - b. **Platform as a service (PaaS)** are cloud services that provide access to a virtual environment for customers to develop, test, deploy and run applications. These include application development computing platforms and pre-built application components and tools which customers can then use to build and manage full applications. There are many PaaS products, and key categories include databases, analytics, containers, machine learning and IoT (internet of things). Examples of PaaS products include AWS Elastic Beanstalk, Microsoft Azure DevOps and Google App Engine.
 - c. **Software as a service (SaaS)** are complete applications hosted in the cloud. These cloud applications can be offered by the cloud provider that owns the underlying raw computing resources or by an Independent Software Vendor (ISV). The provider of the SaaS service manages all hardware and software. In general, most modern consumer and business facing applications are SaaS, including communications services (e.g. Gmail and WhatsApp), broadcasting video on demand (BVoD) services (e.g. BBC iPlayer), productivity software (e.g. Microsoft Office 365 and Google Workspace) and customer relationship management software (e.g. Salesforce Sales Cloud).

2.3. Research objectives

- 2.3.1. The over-arching research objective for this project was to deepen the CMA's understanding of UK cloud customers' decision-making on their choice of public cloud infrastructure services, their multi-cloud and switching behaviour and future use of these services.

2.3.2. Within this, the research needed to address the following research questions:

- a. To what extent do businesses consider different cloud providers (larger/smaller players) and multi-cloud options when procuring public cloud infrastructure services? How does this differ across different types of cloud customer and IT workload?
- b. How do cloud customers consider the role of public cloud infrastructure services compared to services delivered in other IT environments, such as private cloud (e.g. hosted or on-premises) and traditional IT?
- c. Identify the key factors cloud customers consider when choosing their public cloud providers.
- d. How important are the potential barriers (technical, egress fees, discounts and software licensing) to switching/barriers to multi-cloud in customers' choice of public cloud provider? How do these factors vary by customer type?
- e. Why and when do customers choose to switch public cloud provider? Why and when do customers choose to multi-cloud? How do these reasons vary by customer type?
- f. How satisfied are switchers with their experience of switching? What challenges, if any, do they experience when switching public cloud provider?
- g. How satisfied are multi-cloud customers with their experience? What challenges, if any, do they experience when adopting multiple public clouds?
- h. What are customers attitudes to switching public cloud provider and/or adopting multiple public clouds in the future?
- i. What would motivate customers to switch/multi-cloud in the future? What challenges, if any, would they anticipate when switching public cloud provider and/or moving to adopting multiple public clouds? How do these reasons vary by customer type?
- j. What are customer attitudes to the AI-enabled services being offered by cloud providers? How is this influencing their planning and decision-making on cloud services and choice of provider?

2.4. Research target audience

2.4.1. The target audience for this research was medium to large UK public cloud customers, defined as business customers who:

- a. Operate or trade in the UK;
- b. Paid for IaaS and PaaS cloud services in 2022; and
- c. Used one of the five largest cloud providers in the UK – AWS, Microsoft, Google, IBM or Oracle, which provide AWS, Microsoft Azure, Google Cloud, IBM Cloud and Oracle Cloud respectively.

2.4.2. The goal was to speak to key decision-makers on public cloud services within these businesses to understand their choices in their current role, and also draw on their wider expertise and experiences.

2.5. Research method

2.5.1. An overview of the research process is set out below:

1. Set up	2. Recruitment	3. Fieldwork	4. Analysis
Immersion and set up	Screening	Cognitive testing	Data capture
CMA sampling	Letter of authority	Fieldwork	Analysis and reporting

1. Set up

Immersion and set up

2.5.2. The Jigsaw project team of five interview moderators received a detailed briefing from the CMA on the sector and background to the research. This enabled Jigsaw to design the research to take account of the complexities, ensure we would speak to a good cross-section of people indicative of UK public cloud customers, and have an open and unbiased conversation about their views and experiences.

CMA sampling

- 2.5.3. The CMA requested customer lists from the five main public cloud providers to UK businesses: AWS, Microsoft, Google, IBM and Oracle. Companies were selected at random by the CMA from those provided. The AWS, Microsoft, Google, and Oracle samples were each split into three strata depending on the companies' level of spend (High, Medium or Low) and then the High and Medium were over-sampled.
- 2.5.4. Cloud spend thresholds were tailored to each providers' customer lists. AWS, Microsoft and Google had minimum spend thresholds of \$100,000 in 2022. Oracle customers had a minimum of \$10,000 and no spend thresholds were set for IBM customers. The precise definitions for low, medium and high spend categories varied for each provider but ranged from low: \$100,000-\$700,000, medium: \$500,000 to \$2.2M, large: \$1.6M to \$5M in 2022.
- 2.5.5. As well as companies being explicitly stratified by the spend stratum, they were implicitly stratified by the companies' industry and year they first used the provider to ensure a good range of industries were included. Stratification was not designed to obtain a sample representative of a wider population but to obtain a wide range of companies so Jigsaw could ensure a good spread of respondents. Further detail on sampling is provided in the panel.

Sampling – in detail

Sample 1

2.5.6. CMA initially selected approximately 700 companies at random from the sample frame in order to screen and recruit a target of 60 respondents. This comprised 140 companies chosen at random from each of these providers: AWS, Google, Microsoft and Oracle. The entire IBM sample was selected. The companies' contact details were obtained from the providers; and the combined information was then sent to Jigsaw. Some contacts were then excluded by the CMA before recruitment began, as they were already involved in providing information to the CMA as part of the market investigation. The decision was also made by the CMA to exclude public sector organisations.

Sample 2

2.5.7. During recruitment, it became apparent that the number of contacts was not large enough to guarantee Jigsaw would be able to complete the required number of interviews. The CMA released further contacts from each of AWS, Google and Microsoft to Jigsaw to ensure a sufficient number of interviews with these three largest cloud providers for the UK. This comprised 177 contacts from each of AWS and Microsoft (59 in each spend stratum) and 175 from Google (all in the Low stratum, as the High and Medium strata were exhausted).

Provider	Stratum (level of spend)	Number of companies	Number selected Sample 1	Number selected Sample 2	Number not selected
AWS	High	162	47	59	56
	Medium	459	47	59	353
	Low	2,111	46	59	2,006
	Total	2,732	140	177	2,415
Microsoft	High	170	47	59	64
	Medium	505	47	59	399
	Low	2,091	46	59	1,986
	Total	2,766	140	177	2,449
Google	High	25	25	0	0
	Medium	60	60	0	0
	Low	339	55	175	109
	Total	424	140	175	109
Oracle	High	10	10	0	0
	Medium	29	29	0	0
	Low	316	101	0	215
	Total	355	140	0	215
IBM	Total	206	206	0	0
Total		6,483	766	529	5,188

2. Recruitment

Sampling and recruitment

- 2.5.8. Although the sample was drawn directly from public cloud providers customer lists and selected according to key criteria, the contacts provided were not necessarily all eligible to participate in the research, for example, if the individual or business was no longer a public cloud customer. There were also cases of missing or out-of-date contact information for some businesses. Jigsaw's recruitment partner, Serrula Research, attempted contact with all potential participants and undertook a screening process to identify, verify and recruit the relevant decision-makers. *A copy of the screener questionnaire used for recruitment is in the Appendix section 9.2*

Letter of authority

- 2.5.9. As part of the initial contact made with participants, Serrula Research shared a copy of a letter of authority from CMA that explained the legitimacy and importance of the research, and the sampling process. The CMA also responded to any direct enquiries from individuals who chose to contact them directly to reassure on any questions around validity or any other concerns. *A copy of the letter of authority is in the Appendix section 9.1*

3. Fieldwork

Cognitive testing

- 2.5.10. To ensure the interviews would gather relevant information from participants in an unbiased way, Jigsaw carried out testing of the draft research materials. Jigsaw undertook five cognitive interviews on 22 January 2024. These interviewees met the same eligibility criteria as required for main fieldwork but were free-found (i.e. not drawn from the sample supplied to Jigsaw by the CMA). This exercise tested participant understanding of the interview questions, how well the discussion flowed, and (from participants' perspective) any areas of difficulty or omission. Following the cognitive interviewing, minor adjustments were made to the research topic guide.

Fieldwork

- 2.5.11. The main fieldwork was carried out from 29th January 2024 to 14th March 2024. The first few interviews were observed by the CMA so they could feed in final thoughts on the content covered, leading to further minor adjustments to the interview topic guide. Interim findings were provided to the CMA after 21 interviews had been completed, and this also led to a few additional prompts being added to the guide. Interviews were carried out via Zoom and were up to 60 minutes in length. *A copy of the final discussion guide is in Appendix section 9.3.*

4. Analysis

Data capture

- 2.5.12. The raw interviews were audio or video recorded with permission from the participants and, along with other evidence such as notes, were uploaded and stored securely by Jigsaw for analysis.

Analysis and reporting

2.5.13. Jigsaw set up a Qualitative Analysis Frame in Excel with a row for each participant and a column for each key theme that the individual moderators filled in upon completion of their interviews. This meant the evidence around each topic was collated systematically, was grounded in the views of participants, and no voice was lost. Analysis was carried out collectively with input from all five moderators involved in interviewing. Jigsaw shared interim headline findings after 21 interviews. Once all interviews were complete, Jigsaw presented findings to the CMA in a PowerPoint debrief presentation. This report follows on from the debrief with further iteration and detail based on ongoing analysis of the data.

2.6. Sample structure

Soft quotas

2.6.1. The goal was to include participants from a wide range of different types of customers in the research, so voices indicative of customers of the top five UK public cloud providers would be heard and no key groups inadvertently missed. A good spread of customers was achieved.

Individual and paired interviews

2.6.2. In total, 60 in-depth interviews were achieved with individuals from 50 companies. In 10 companies, where a wide range of individuals were involved in decision-making on public cloud use within the business, two interviews were undertaken to understand different points of view, for example, to gain insight into both strategic and tactical day-to-day decisions.

	Total
Individual interviews – x1 interview within a business	40
Paired interviews – x2 interviews within the same business	20
Total	60

Cloud provider and spend levels

2.6.3. Cloud spend thresholds were tailored to each providers’ customer lists. AWS, Microsoft and Google Cloud had minimum spend thresholds of \$100,000 in 2022. Oracle customers had a minimum of \$10,000. No spend thresholds were set for IBM customers. The precise definitions for low, medium and high spend categories varied for each provider but ranged from low: \$100,000-\$700,000, medium: \$500,000 to \$2.2M, large: \$1.6M to \$5M in 2022.

	Total	Low spend	Med spend	High spend
AWS	20	3	8	9
Microsoft	17	2	7	8
Google	12	4	5	3
Oracle	4	2	2	-
IBM	7	7	-	-
Total	60	18	22	20

Business sector

2.6.4. For some business sectors there were fewer companies in the sample provided to the CMA by the public cloud providers, so it was harder to achieve interviews in these sectors, however, interviews across all target sectors were achieved. Public sector organisations were excluded from the sample.

	Total
Retail, consumer, hospitality	13
Banking, insurance and financial services	11
IT	8
Business and professional services	8
Telecoms, media and entertainment	6
Manufacturing, utilities and primary industry	5
Healthcare and life sciences	3
Other	6
Total	60

Length of public cloud use

	Total
Used IaaS/PaaS for 5+ years	39
Used IaaS/PaaS for <5 years	7
Always used public cloud (Native)	14
Total	60

Egress fees – spend in the last 12 months

	Total
<£50k	31
£50k-£100k	8
£100k-£250k	3
£250k-£500k	0
£500k+	3
Don't know	15
Total	60

2.7. Methodological note on qualitative research

Interpretation and limitations of qualitative research

- 2.7.1. When carrying out qualitative research, the focus is not to quantify but to explore underlying views, experiences, and motivations. Qualitative samples are purposive and quota-driven in nature; they are designed to achieve specific outcomes, such as including a good range of different customer groups relevant to the research. Consequently, they cannot be used to identify proportions of populations holding stated views. For these methodological reasons, it is not appropriate to present qualitative findings in terms of the numbers of participants expressing certain views. The responses should not be seen as representative of the universe but can be seen as strongly indicative. The findings in this report are therefore described in qualitative terms. To give an indication of the strength of a finding from the interviews conducted we have used terms such as 'a small number', 'some', 'many', and 'most'.

Participant recall

- 2.7.2. Some key facts about our participants and their businesses were available through the original sample (business sector, whether the business is a high, medium or low spender on public cloud). Other reported information is based on interview participants' recall and interpretation, and these are not always precise. Sometimes there is uncertainty because decisions around public cloud were taken by previous staff in the company, and therefore the process and rationale may not be fully known. In other cases, participants may simply be uncertain about key facts or events they convey. Their answers reflect their best and subjective understanding of the situation. We asked participants to prepare for the interviews to help encourage them to reflect on choices. However, it was not part of the research objectives to review any documentation, for example, to check the precise nature of public cloud contracts or discounts. It is a finding in itself that some facts are not retained or possible to summarise easily, and this can relate to the complexity of some contractual arrangements and billing or a lack of priority or perceived need to know that information. We comment on this where relevant in the sections on Theories of Harm.

Individual perspective

- 2.7.3. Participants in this research have provided their personal views as individuals in their role as an IT professional making decisions about public cloud infrastructure services for their organisation. They are also drawing on their wider expertise and knowledge, often from long careers in the IT sector and previous roles they have held. Each response should be interpreted as being from an individual experienced customer rather than as an official collective submission from their organisation.

Snapshot in time in a fast-changing market

- 2.7.4. A key feature of the world of IT is that the pace of change is fast. Participants often add caveats when expressing their views, that they are based on the last time they reviewed or used a particular public cloud infrastructure provider or implemented a particular solution. Respondents are most familiar with their current provider or providers and make clear that they may not have an up-to date-view of competitors or the latest solutions or tools given changes can happen month-to-month.

Anonymity

- 2.7.5. Individuals who have taken part have done so on the basis that they will remain anonymous, and their details and the company they work for will not be disclosed. This is standard practice in market and social research (and an industry requirement of the Market Research Society Code of Conduct) and aims to encourage open and honest answers from those who take part without fear of any repercussions.
- 2.7.6. Verbatim quotes are annotated with the individual's business sector, main cloud provider(s) and spend level only, to indicate the type of respondent but preserve their anonymity.

Technical terms used

- 2.7.7. Technical terms used in the research were provided by the CMA and are consistent with the glossary of terms used in the CMA's Market Investigation.

3. Main findings – Context

3.1. Overview

3.1.1. Before looking at the Theories of Harm, it is helpful to understand the perceived nature and drivers of decision-making around public cloud among those interviewed and how these related to choices. An overview of key themes is provided below, with detailed references where relevant throughout the report.

3.2. Public cloud decision making

Decision-making varies by type of business, including maturity, industry sector and nature of the work. A common driver underpinning choices across all types of business is the tendency to ‘stick to what you know’ as a default as this is perceived to be the path of least cost and risk.

3.2.1. At the start of the interviews, participants were asked to provide a brief overview of how they made decisions about use of public cloud.

3.2.2. Approaches to decision-making and considerations around public cloud vary by types of business.

	More flexibility	More constraints
Maturity of business	Start-up businesses (may have evolving operational models)	Established businesses with legacy systems
Public cloud spend	Lower spenders – decisions carry less cost and risk	Higher spenders – decisions carry more cost and risk
Business sector	IT/technology-driven businesses with strong IT skills/capacity in-house	Businesses in highly regulated sectors e.g. Finance, Health

Differences in decision-making by maturity of business

3.2.3. The nature of decision-making varies most markedly by how established a business is:

3.2.4. **Start-up businesses:** Newer businesses, by definition, tend to be smaller and still evolving in their business model and operations. In many cases, the key decision-maker around public cloud is a founder or co-founder of the business as well as being CTO (Chief Technology Officer). Many start-ups are cloud native with technology core to the function of their businesses, and those we spoke to tended to have a lower public cloud spend per annum (in the region of £1,000s if still on start-up credits to £100,000s).

- 3.2.5. Decisions tend to be driven by the IT leader's own previous experience, and potentially one or two other early IT employees, plus advice and support from investors, incubators or key partners. Start-up credits from public cloud providers, either directly or indirectly via their deal with their investors or incubators, are also a major influence on choice. Start-up businesses often make use of multiple cloud providers in their early years to take advantage of credits from each, as this keeps cost low, allows experimentation, and means they can draw on a range of provider input and advice. Over time there is consolidation to one or two chosen public cloud providers once their operating model, funding or client base become more established, at which point the best provider to use has become clear and/or the cost and risk of any change is perceived to increase.
- 3.2.6. In summary, start-ups tend to have tight cost constraints, will work at pace and with less formal decision-making processes, and are more likely to test and switch public cloud provider whilst the cost and risk to operations is low.
- 3.2.7. **Established businesses:** Those in more mature businesses have often gone through the experience of moving from on-premises infrastructure to public cloud (or their predecessor had), either 'lifting and shifting' in one go or piecemeal over time. Those we spoke to tended to have higher public cloud spend per annum (£10,000s, for example, among hybrid on-premises/public cloud users to £10,000,000+).
- 3.2.8. In many cases, there are multiple decision-makers and more formal processes than in a start-up. Those involved can include the CTO plus leaders from Infrastructure, DevOps (responsible for both IT development and operations) and Security, and internal IT policy and decision-making groups. Major decisions involving significant strategic shift or spend, such as migrating to public cloud or switching IaaS provider, can be scrutinised by the CFO (Chief Financial Officer) and Executive Boards. More tactical decisions, for example around best PaaS provider or how to optimise services, receive less scrutiny and can be managed less formally by technology leaders and managers. Decisions tend to be driven by the need to consider existing legacy infrastructure and software often coupled with existing skills and experience with these systems within the businesses. Major changes to infrastructure are considered high cost and a high risk to operations in established businesses, so tend to be rare, for example, only being reviewed every 5-10 years or in some cases not really on the horizon at all unless something significant shifts in the market.
- 3.2.9. In summary, established businesses have more formal decision-making structures, and are less likely to switch public cloud provider for IaaS or PaaS services unless there is a strong business case to do so, as they perceive this would be complex, expensive and present a high risk to operations.

Differences by public cloud spend

- 3.2.10. Whether a business is newer or more mature, choosing a provider or switching between them is more feasible if their use involves less spend or less complexity. For example, a business that still has its own data centre and only uses certain PaaS or compute services for specific tasks may be more open to consider switching provider. This does not necessarily mean that those that spend less switch more. A low spend on relatively basic cloud infrastructure can mean less engagement and scrutiny.
- 3.2.11. Those with higher public cloud spend tend to have more resistance to change or switching provider as the complexity or cost is high. For example, a business has a high spend on one provider for IaaS, PaaS and potentially SaaS too, may be more reluctant to consider switching or will need to see a major shift in performance or cost of competitors to do this.

Differences by business sector

- 3.2.12. In some industries, whether a business is a start-up or more established, the sector they operate within can influence the range of choice open to them.
- 3.2.13. IT or more technology-driven firms have more flexibility because of the skills in the business:
- a. **IT (or technology-driven companies)** – often have the skills and capabilities and/or capacity to take advantage of more options e.g. to select and customise open-source PaaS services or to develop their own solutions rather than using cloud providers more ‘off the shelf’ proprietary services
- 3.2.14. Some sectors have more constraints as there are policy or regulatory considerations:
- a. **Finance** – businesses operating in the finance sector must abide by financial regulation, and strict criteria on data security and retention, which may include rules across other markets such as the EU. Their choice of public cloud provider may be driven by factors such as what services can offer appropriate levels of security and be able to meet audit and reporting requirements.
 - b. **Health** – as well as rules around privacy of health data, NHS policy dictated that many services migrate to the cloud a few years ago, meaning private sector companies in healthcare that interact with the NHS have also had to follow this approach. Their choice of public cloud provider tends to be driven by factors such as security and what services health partners are using.

Influence of individuals and their skill-set on decision-making

- 3.2.15. Whatever the process of decision-making, a common factor across all types of company is the influence of existing skillsets within the business on choices around public cloud. Sticking with what is known is the path of least resistance. It also avoids the cost and risk of a steep learning curve, retraining or hiring new staff, which are only deemed worthwhile if there might be a significant gain in performance or cost reduction long term.

Low use of competitive tendering in decision-making

- 3.2.16. For major infrastructure decisions, across all types of company, some participants mention a process of competitive tendering or holding conversations with multiple providers. However, in many situations the choice of public cloud provider is seen as a given based on internal expertise and coding languages used in the business, or the need to use the same provider as key partners and clients. In these cases, there is little time spent considering alternative options as they do not appear relevant or realistic.

“Public cloud usage has been more accidental and driven by developer familiarity rather than a strategic decision. The barriers to entry to using these service providers is very low. Some developers will be more familiar with one toolset than another, and when they’re asked to do something, they typically select the cloud service with which they are most familiar.”

IT/Retail, Microsoft main provider (also AWS, Google), Medium

“Because we have these 60 people that all have been working in Microsoft for a period of time, we took the decision to go with Azure (from on-premises).” **Retail, Microsoft, High**

3.3. Why public cloud?

There are multiple perceived benefits of public cloud compared to private on-premises or private cloud. For most, the aspiration is to have all their business infrastructure on public cloud, but cost or reduced control present a challenge for some use cases.

- 3.3.1. Participants were asked to describe their use of public cloud infrastructure services and the rationale behind their choices.
- 3.3.2. Many of those interviewed were only using public cloud for their infrastructure, having always used this approach or having migrated fully to that position from an on-premises set-up in the past. Others were using a mixture of private and public cloud.
- 3.3.3. An overview of the perceived benefits of public cloud and some of the challenges is below:

Benefits

Requirement of doing business: In many cases, those interviewed worked in businesses that could not operate at all or as effectively without use of public cloud. This can be driven by:

- Public cloud being core to business function e.g. firms selling their own PaaS/SaaS services.
- Industry/partner requirements e.g. NHS 'cloud first' policy meaning private healthcare businesses also need to use the same approach.
- Customer requirements e.g. customers insisting their data is in their geolocation, which is cheaper to achieve via public cloud than via private cloud or on-premises servers.

Lower cost (for most use cases): The barrier to starting to use the public cloud is low. Participants mentioned start-up and migration credits, pay-as-you-go services and easy basic set-up for many public cloud providers. These were especially important for start-ups who cannot afford the initial capital expenditure needed for their own data centre or the staff to build their own bespoke IT services.

Reliability: Most public cloud providers are reputed (and had proved) to be highly resilient, with less risk to continuity than private or on-premises data centres.

Scalability: If a business was in growth (or required flexibility in volume of workloads), it was seen as easier to scale up via public cloud than via a private data centre or on-premises set-up.

Easier use and maintenance: On-cloud infrastructure was perceived to enable easier set up, integration and software updates. Essentially, the public cloud provider was seen as sharing the burden of architecture design, upgrades and maintenance rather than the business having to do this on its own.

Innovation: Cloud providers are great innovators. Many businesses valued sharing the burden, where the providers innovate on infrastructure and the user can focus on their core business function. The public cloud is also seen as a good test environment for research and development, for example exploring use of AI.

Support and advice: The main public cloud providers, especially AWS, Microsoft Azure and Google Cloud, are seen as experts in their field. They are valued for great advice, their ability to problem solve and to keep their customers up to date with new developments.

Challenges

High cost for some use cases: Some felt they would like to use public cloud more for some use cases, but the high cost of doing this had been a significant driver to set up their own private data centres. Examples of these use cases include:

- **Machine learning via GPUs** – a couple of participants mentioned the high cost of using co-located, high performance GPUs meant that they had instead set up their own private data centre using ‘cloud bursting’ when required i.e. only accessing additional public cloud GPUs when their data centre was at capacity.
- **Data transfer** – a couple of participants mentioned that the high cost of transferring large volumes of data from one public cloud provider to another had meant that their current business had set up its own private data centre

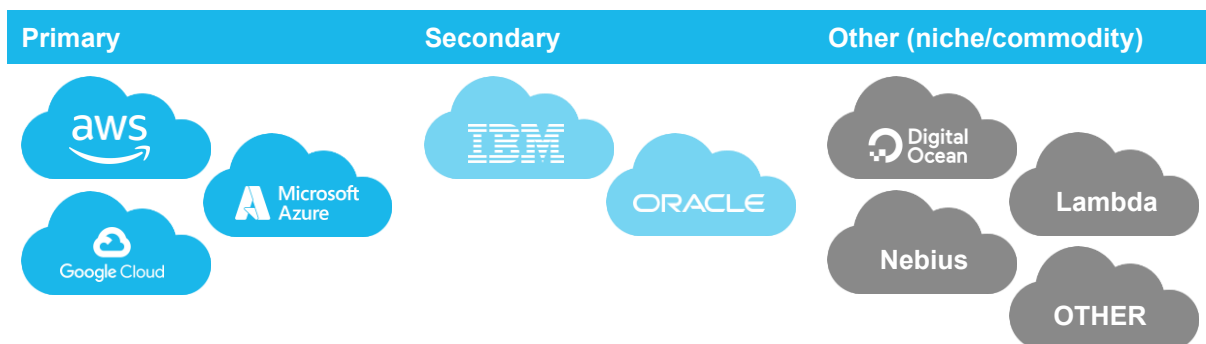
“Decisions around public cloud vs private cloud are made modularly based on factors like performance needs, cost, security, and being able to avoid vendor lock-in.”

Retail, Microsoft, Medium

3.4. Reasons for choosing individual public cloud providers

For many in our sample, AWS, Microsoft, and Google are the main providers that they would consider using, with other providers perceived as being for secondary use or only for certain use cases.

- 3.4.1. Participants were asked which public cloud infrastructure providers they used, the nature of that use and why they chose each relevant provider. They were also asked about their general opinions about providers.
- 3.4.2. Participants often had a clear perception of who they felt the leading providers are, plus some understanding of others available.
- 3.4.3. Note that, by definition, participants were most familiar with their current provider or providers, and make clear that they may not have an up-to-date view of competitors given the pace of change is very fast. Their view of other providers may be based on the latest technological news, word of mouth or learnings from the last time they reviewed or implemented use of a particular provider or solution, rather than current experience. They also state that, where there is innovation by a provider (or via open source or a third party), others quickly aim to follow.



Primary providers

- 3.4.4. The main providers are seen as AWS, Microsoft and Google among participants. For some, this is the main or only consideration set in terms of who might be on a shortlist of providers in the event of a review of the market or a switch, though most were not aiming to make any changes. They each have an excellent reputation, deliver a reliable service and offer a wide range of solutions that cover many needs.
- 3.4.5. AWS is perceived as the overall leader, in terms of the level of innovation and range and quality of services provided, followed by Microsoft, and both are used by those working in a wide variety of types of businesses. Google is also thought to be highly reputable, but has a smaller pool of users among those interviewed. Google users in our sample are more likely to be tech-driven businesses such as start-ups, e-commerce, fintech or adtech companies.
- 3.4.6. Many felt that services provided by these main players were '80% the same', but there were distinctions perceived e.g. some see AWS as the first mover and innovator, Microsoft as having superior IAM, and Google as having a strength in analytics (specifically, BigQuery as having better performance).
- 3.4.7. Spend on main providers could be up to millions of pounds a year.

Secondary providers

- 3.4.8. Very few interviewees used IBM as a sole provider and none used Oracle in this way. Their main use was secondary, for example, for supporting legacy systems. The types of businesses using these providers were also more traditional, such as those involved in energy, health, regulatory support or insurance.
- 3.4.9. Spend on secondary providers could be up to hundreds of thousands of pounds a year.

'Commodity' or niche providers

- 3.4.10. A few participants used other smaller providers, such as Digital Ocean or Lambda Labs, for compute services only, as and when needed as a 'commodity'. A few also used them for niche benefits e.g. Digital Ocean as it felt easy for one start-up before it scaled up, and Nebius for its green credentials.

"We tend to focus on the big three providers as they have the scale, the products, investment they put into their platform, the constant innovating and armies of developers, so you can stand on the shoulders of giants. And we can find the skills. Smaller providers are just like another data centre."

Media, Microsoft and AWS main providers (Google legacy use), High

AWS – First mover offering innovation and scale



Common types of users (among participants): Tech driven businesses e.g. fintech, adtech, e-commerce, start-ups
Spend: Up to £10,000,000s per annum

Benefits

Innovator: AWS is commonly seen as a 'first mover'. There was a sense that other public cloud providers are catching up but there is still a reputational advantage.

Requirement of doing business: AWS was often a 'default' choice as it is used by partners and clients e.g. it is used as infrastructure for ticketing so used by firms in the transport sector.

Reliable: Many participants state it is resilient and secure, and a few mention that it also supports a range of regulatory needs and reporting requirements.

Ease of use and maintenance: AWS is seen as having a wide range of solutions that can help reduce workload for users, and is seen by some as more flexible and configurable than the competition.

Ease of set-up: There are a variety of ways mentioned in which AWS enabled easy set up and/or migration from on-premises to the cloud: They have a migration support service, credits for start-ups and switchers that can be worth £100,000s, and fast set up using Elastic Beanstalk.

Good for some uses/workloads: Some saw AWS as best for Linux servers/databases, voice solutions, and for use of JavaScript.

Support and advice: AWS is seen by some as a provider of expert advice, for example helping customers to optimise to reduce cost and keeping them informed about new innovations. They are also seen by some as proactively trying to keep their business.

Challenges

Greater risk of 'lock in' than competitors: A sense of dependency can be inherent with any cloud infrastructure provider, however the wide range of proprietary services from AWS made life easier for some customers, and led to a reduced IT headcount but also a sense of reliance.

Complex billing: Some felt bills from AWS are detailed and take time to review, with a few larger businesses having dedicated resource to support this activity. An example is one participant who said their bill can be a '100-page PDF'. It can take time to manage costs and optimise for efficiency.

Cost: Some said AWS is relatively expensive compared to competitors, for example, the egress fees can be higher. This can be a frustration but also accepted given the quality of service.

Amazon can be a competitor: A couple of firms could see a risk of Amazon developing services that compete directly with their own business. Their use of AWS brought benefits, and kept them close to how AWS was innovating, but was also a risk.

"AWS is the gold standard. Cost aside, it has the best reputation for being reliable, the best availability in terms of geographic coverage, and our clients use AWS, and we use Redshift which is proprietary – apart from the cost there's no real reason to move. They sting you on egress fees."

Retail, AWS, Medium

Microsoft – Whole-enterprise solutions



Common types of users (among participants): 'Microsoft Houses', traditional sectors e.g. energy, insurance, law, hospitality, tech B2B businesses where their clients use Microsoft

Spend: Up to £10,000,000s per annum

Benefits

Familiarity and ease of set up: Microsoft Azure is a default choice for many customers who refer to themselves as 'Microsoft Houses' or 'Microsoft Shops' due to their long-established relationships with Microsoft, its legacy software and use of its software development framework .NET. There is also great support from Microsoft and a range of third parties for 'lift and shift' migration from on-premises to public cloud. This can mean they migrated from on-premises to using Microsoft Azure for all their public cloud services, or in some businesses, for internal infrastructure only with a competitor for external-facing infrastructure/core business function.

Requirement of doing business: Microsoft Azure is often a 'default' choice as it is used by partners and clients e.g. it is commonly used in some industries requiring high degrees of security such as the legal sector.

Good integration with other Microsoft products: Microsoft Azure is seen to offer better interoperability with products such as Office 365, Dynamics 365, Windows Servers, security platforms and more. Using Azure has also allowed some participants' firms to access extended support for legacy Microsoft products.

Good for some uses/workloads: Some mentioned Microsoft Azure as best for analytics, disaster recovery (at a cost), location of servers to help ensure continuity, and offering better authentication services than AWS.

Good value for money via Enterprise Agreements: Microsoft Azure is often sold as part of wider agreements that provide access and support for a bundle of Microsoft products. Microsoft Azure is seen by some as good value compared to the cost of other software licences and services.

Challenges

Risk of 'lock in' from Enterprise Agreements and historical relationships: Microsoft is a reliable and respected provider, but some perceived a risk of dependence from use of a wide range of services, long-term legacy use and the in-house skills and experience that they had developed.

"The business was already a Microsoft shop, so Microsoft PCs, Microsoft 365/or Office at the time and a Microsoft Active directory, which means you already kind of have a relationship with Microsoft and Azure. But I think probably more importantly, our code was built in .NET which is Microsoft's programming framework. So, operating in Azure, it was probably a little bit simpler and more familiar for the teams at the time in 2015-16."

Retail, Microsoft, High

Google – Strong on analytics and AI



Common types of users (among participants): Tech-driven businesses e.g. fintech, adtech, e-commerce, start-ups.
Spend: Up to £1,000,000s per annum

Benefits

Familiarity: Google Cloud is seen as a sensible and familiar choice for those who use other Google services (in a similar vein to businesses that are Microsoft Shops feeling Microsoft Azure is the rational choice, but less marked). Google Cloud is seen to offer smooth integration with products such as Google Workspace and Google Marketing (Ads and Analytics)

BigQuery: Some participants felt BigQuery (Google PaaS analytics service) is a superior service to the equivalents from other cloud providers, primarily as it has faster retrieval on large databases. It is sometimes used by participants even though they have a different IaaS provider.

Good integration with other platforms: Some participants mentioned that Google Cloud allows easier integration with other tools used such as Kubernetes and Databricks.

Easy to use and manage: Some felt Google has better 'off the shelf' proprietary tools and apps.

Great support for start-ups: Some start-ups felt that Google offers better advice and help with optimisation and problem solving and proactively suggests new features. Google Cloud are also known to offer good start-up credits, one example being £300,000 credits for two years.

Good for AI: Google is seen as superior in particular to AWS in terms of its AI offering including for tasks such as machine learning and translation. It also offers AI powered Google Marketing for Google Cloud users.

Challenges

Complex billing: Some felt bills from Google Cloud are detailed and take time to review, with a few larger businesses having dedicated resource to support this activity. In turn, it can take time to manage costs and optimise for efficiency.

Less configurable: A few felt Google apps and tools can be user friendly, but not as configurable as those from AWS e.g. cannot be as easily adjusted or customised to the participant's needs.

"The level of technical support we get has been much better from Google Cloud compared to other providers ...and products like BigQuery and Bigtable provide performance and capabilities that are light years ahead of everybody else."

Professional services, Google, Low

IBM – Mostly secondary use



Common types of users (among participants): Traditional e.g. energy, health, regulatory compliance.

Spend: Up to £100,000s per annum

Benefits

Flexibility: A few said IBM is highly configurable and one person mentioned that IBM also charged less than AWS and Google (previous providers they had used) to run their own or third party tools e.g. layered architecture.

Easy financial management: One participant said IBM's pricing was more consistent and predictable than with AWS, one of their previous providers (although one other participant held a directly opposing view and said IBM pricing lacked transparency).

Good governance: One participant mentioned that IBM has a good reputation in the financial services sector and with more traditional companies, and that IBM is part of the governing council for Web3.

Challenges

Falling behind competitors: IBM was described as falling behind by a few participants. Some users had or were planning to switch away from using IBM for a variety of reasons including their services being less easy to manage, less scalable and not commonly used by customers or partners.

“So, in 2019 the decision was made to switch to AWS and we exited IBM. Three key reasons for that – 1. was the relative immaturity and instability of the services provided by IBM, 2. was a lack of pricing transparency in IBM, it was unclear where the money was going and on what and they didn't really support us in figuring that out, and 3. our customers and partners at the time were using other clouds, AWS principally, and so we were kind of gravitationally pulled in that direction. And finding developers who have experience with IBM is a lot harder than finding those people with the skills in AWS.”

Manufacturing, AWS, Low

Oracle – Secondary use



Common types of users (among participants): Traditional e.g. manufacturing/insurance, legacy systems.
Spend: Up to £100,000s per annum

Benefits

Supporting Oracle software and legacy systems: A few mentioned using Oracle for specific legacy systems that only run on Oracle cloud infrastructure or work more effectively if they do e.g. Oracle ERP, Insuretech SaaS solution, databases.

Challenges

Falling behind competitors: Oracle is seen as only relevant for certain workloads. It was not used as a main provider by any participants interviewed.

“We did a big study before we started to look at the target cloud that we should aim for, looking at each application in turn and based, generally speaking, on the technology that it was running on. So, we’d be a bit unusual, I think in having about 50% of our stack in Azure and about 50% of our stack in Oracle Cloud. That is because the technology that we’re migrating, a lot of that was written using Oracle technology. So for databases, in particular, Oracle has this technology RAC that basically means you’re running two databases in parallel. And so if one goes down, the other one keeps running. It’s very hard to run an application that’s got that service underneath it in any other cloud other than Oracle’s. So there are technology constraints that have guided us towards Oracle Cloud. What Microsoft are better at is coming in and sitting with us and going OK, what’s the next business problem you want to solve?”

Utility, Microsoft main provider (also AWS and Oracle), Medium

Other cloud providers – mostly secondary use



Common types of users (among participants): AI/high load/'cloud bursting' or very early start ups
Spend: Up to £100,000s per annum

Benefits

Niche benefits: One business used Digital Ocean as their only provider in their early days as a start-up as it has a good reputation in the financial sector. Otherwise, it is used by a few for compute services only. A few mentioned Digital Ocean has fair pricing and an easy interface for cost management, and clear upgrade and down grade options. As such, it is seen as being good for start-ups who value ease of set up. Nebius was used by one participant as it is a sustainable cloud provider.

GPU access: Other small cloud providers like Lambda Labs, Together AI and NextGen, were used for capacity and price by a few participants that needed access to extra processing capacity in addition to their private data centre.

Challenges

GPU costs: Scarcity plus growing demand for GPU access means costs for public cloud GPU access are very high for a few AI/high load users – a couple of participants said they had chosen to use their own private data centre in their current business to make this work affordable, with use of public cloud as and when needed to top up capacity.

Lack of scalability: Only a few participants considered smaller providers other than for added process capacity, as they did not see clear benefits of doing so. Lack of scalability is a key barrier.

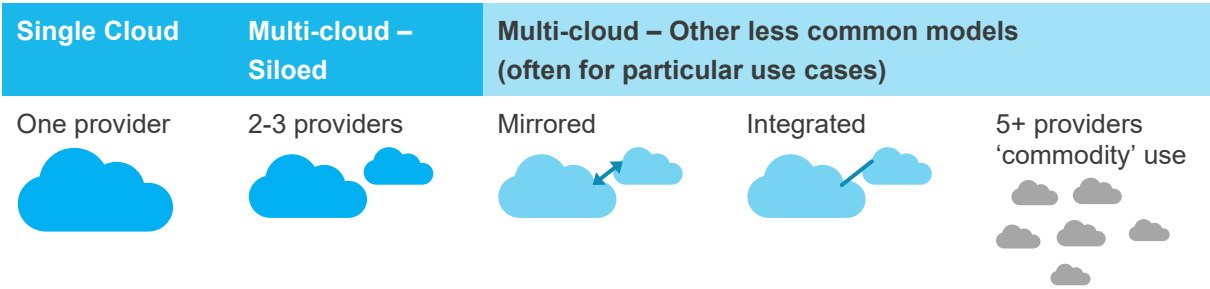
“So Digital Ocean is kind of designed for getting start-ups off the ground. So to create a database, it's just a couple of clicks. And then to create your API, you just sort of tell it what you want and then place your code inside it, and then it effectively boots it up for you, which is quite nice.”

Fintech, AWS (previously Digital Ocean), Low

3.5. Public cloud operating models

Single cloud use and siloed multi-cloud are the main operating models among our participants' current firms.

- 3.5.1. When speaking to participants about the public cloud infrastructure providers they used, moderators explored their rationale for using one provider (single cloud) or more than one (multi-cloud), and where more than one provider was used, the level of integration across different public cloud providers.
- 3.5.2. The main operating models among those interviewed were to use a single public cloud provider for their current company infrastructure or two to three cloud providers with siloed workloads. Other models do exist, but were far less common and mainly for particular use cases rather than being the core approach of a business to public cloud. An overview of these is below.



- 3.5.3. **Single Cloud:** Some participants' companies used one public cloud provider as participants regarded this as efficient and effective in terms of cost and staff time given the expertise developed in the IT team with this provider's systems and coding language. They were mostly satisfied with the service they received, and it was rare that there was a reason to opt to use a multi-cloud approach or to switch.
- 3.5.4. **Multi-cloud – Siloed:** Some participants had a multi-cloud infrastructure set-up, commonly involving one dominant cloud provider where most infrastructure was located, plus one or two other providers for specific separate workloads. This was driven by internal or external factors, or both:
 - a. Internal factors included e.g. legacy systems only working or working most effectively on one provider and all other services being elsewhere, mergers where two firms that joined together historically used different providers, continuing to use Microsoft Office/365 on Azure as this was familiar for staff across the organisation and provided good functionality, but all other workloads being on AWS or Google.
 - b. External factors included e.g. matching public cloud providers used by key partners to enable easier interaction, or matching the provider used by clients, and some of their clients use AWS and others use Microsoft.

3.5.5. **Multi-cloud – Other models:** A few participants used other infrastructure models to support particular use cases.

- a. **Mirroring** – There were no cases of cloud architecture being simply duplicated on two or more public clouds, but there were some examples of plans being in place to easily transfer to another service in the case of significant change in their business need or provider failure e.g. through use of ‘open-source’ rather than provider’s proprietary services to enable easier portability if needed, use of tools that help define infrastructure as code, or actively maintaining disaster recovery plans and/or low-level relationships with other cloud providers (sometimes as a regulatory requirement).
- b. **Integrated** – There were a few specific cases of choosing to integrate services across public cloud providers. The main reason for this was the perceived performance benefits, with the most mentioned examples being use of BigQuery from Google or Microsoft IAM alongside infrastructure from a different public cloud provider as these services were superior to those offered by their main provider. Most businesses avoided integrating across public cloud providers as this presented unnecessary challenges, such as added complexity, security risk and the need for extra training or staff.
- c. **Commodity use** – There were a few examples of firms using five or more public cloud providers at once. This was another form of siloed use, but where the behaviours were price-driven, sometimes temporary or for short periods, and considered separately to their main provider(s). The two kinds of use uncovered were start-ups using services because they were free or heavily discounted via start-up credits, even though they did not intend to continue to use these suppliers once credits ended, or companies ‘cloud bursting’ (using public cloud only when their private data centre was at capacity, allocating work automatically based on provider availability and price for uses like machine learning).

3.5.6. Some key perceived benefits of the two most common public cloud operating models are below. The technical enablers and barriers to multi-cloud use are explored in detail in the main findings on Theories of Harm.



Single cloud

Benefits

Familiarity: In many cases, participants and their wider teams have built their expertise in the coding language and systems of their single cloud provider, and are increasingly able to optimise and find efficiencies.

Reduced complexity and risk: Many single cloud users stated that having one provider meant there are fewer unforeseen technical or security issues, for example, from data transfer between providers, that it is easier to assure regulatory and security compliance and easier to co-locate services.

Cost efficiency: Related to the reduced complexity, there is no need to set up or manage integrations across providers and reduced data transfer fees.

Satisfaction: Many said they were happy with their provider, and have all their current needs covered, so do not have any reason to adopt a multi-cloud approach (or to switch). They also did not currently perceive enough differentiation between providers for their use cases to justify using more than one provider

Leaner in-house team: Some participants mentioned that the provider shares responsibility for maintenance, rather than their current business on its own, meaning the business does not need such a big in-house IT team. There is also only one system to learn and train the IT team on.

Fit with clients/partners: For some participants, matching the public cloud provider used by key partners or clients made it easier to interact or serve them.

Easier billing: A few participants mentioned that it is easier to manage one billing system, including recharging relevant fees to their clients.

"I only have one infrastructure manager that handles Azure. I would maybe need a second one to handle AWS also and that's not in our budget to hire more people for that."

Retail, Microsoft, Medium

"I can't think of an architecture where multi-cloud would make sense for us. The only appeal would be mirroring, but actually the service is pretty reliable so why bother. And even if the service goes down, it goes down for everyone, so you can always 'blame' Google."

E-commerce, Google, Low



Siloed multi-cloud

Benefits

Putting a workload with the best provider: Some placed different workloads with different providers to use what they consider best in class, for example using Microsoft Azure for internal facing systems like Office 365, but using AWS for external facing infrastructure.

Alternative provider in case need a Plan B: For some, maintaining relationships with alternative providers meant there is back up in case business needs change and in case of disaster recovery. For certain sectors e.g. utilities it is a regulatory requirement to have more than one provider.

Fit with clients/partners: For some businesses, matching the public cloud providers used by key partner or clients makes it easier to interact or serve them, and so if their clients used different providers their business would do the same.

Cost effective: A few mentioned that having more than one provider meant you can get the best price for some use cases e.g. compute, use of GPUs, and also helps with price negotiation with each provider so there may be a reduced risk of price inflation.

“For the AI infrastructure stream that our research team uses, for example to develop new models, that is completely agnostic. Wherever we find the cheapest GPUs we will go. And we have developed our own management system to deploy across any of your usual suspects plus a lot of the ones that you probably haven’t heard of, whereas our production system that has up-time requirement SLAs now runs on Google Cloud...For the research part we use Google Cloud, AWS, Lambda Labs, Digital Ocean, Together AI and NextGen Cloud, I think.”

IT, Google, Low

3.6. Switching

Public cloud is infrastructure – to an extent, dependency or a sense of 'lock in' is a factor across all providers as change brings cost and risk. Switching is not a common choice and only occurs if there is a significantly strong business case and return on investment.

Perceptions of switching

- 3.6.1. As part of the interviews, moderators asked participants if and why they had ever switched public cloud infrastructure providers and about their experiences. Moderators also explored if participants would consider doing this or not and why.
- 3.6.2. Many businesses aspire to have their infrastructure on public cloud. Migrating to public cloud from on-premises or considering whether to host via public or private cloud are key strategic decisions. Once a cloud public provider(s) was put in place, many participants felt less need to switch all or part of their infrastructure to an alternative provider(s), as they could not see a clear rationale.
- 3.6.3. Switching cloud providers is seen as the equivalent of moving other kinds of infrastructure, such as 'moving house' or moving a business from one country to another. It is not something to undertake lightly or consider at all unless it leads to significant business benefits long term that override the perceived cost and risk of changing.
- 3.6.4. There is a preference towards consolidation, using as few public cloud providers as possible. Ideally, many favour the simplicity of one provider that covers their business needs, rather than seeking to use multiple providers, but there were benefits of multi-cloud for some (as seen in the previous section).

Perceptions of 'lock in'

- 3.6.5. Given public cloud providers provide essential infrastructure and switching was not perceived as straightforward or desirable for many, this does bring a sense of dependency or 'lock in'. In most cases this was an accepted risk and mitigations were put in place in case of any significant change in business need or failure of their public cloud providers.

"It's like trying to pick your house up and move it right, it's that complicated. So, you've got to pick up the whole thing in terms of the physical architecture, the data, but also the contracts the commercials... So that's why people sign up to massive contracts with one provider... You decide where you're going to put it. You put it there and you leave it there."

IT, AWS main provider (and Microsoft), High

"The more complicated your business the more locked in you are. When we decided to switch, we expected 1 month of work, but it was more like 3... one thing is terminology – you have to learn an entirely new language."

E-commerce, Google, Low

“We have vendor lock in with Amazon, you know that’s a risk...we’ve accepted the risk and we’re very aware of the risk.”

Retail, AWS, Low

Experience of switching

- 3.6.6. Few of those interviewed had chosen to switch between public cloud providers.
- 3.6.7. Where people had switched, either all or some of their infrastructure, this was driven by an anticipated reduction of cost or clear increase in performance. In some cases switching was a requirement following a merger or acquisition to rationalise providers.
- 3.6.8. Switching could be a straightforward process of a few weeks for start-ups, or take up to months or years for more established or complex businesses, followed by a long process of learning and consolidation to optimise and increase efficiency. The process of switching typically involved running two sets of services concurrently, then switching off the legacy system when ready. A few of those who had switched needed to employ third parties to advise and implement.
- 3.6.9. Overall, for those who had switched, the process not only brought cost and operational risk, but took IT staff away from core work and typically ended up being more challenging and time consuming than anticipated.
- 3.6.10. The technical enablers and barriers to switching are explored in detail in the sections on Theories of Harm.

“We started in parallel using IBM and we have various what we call micro-services which are like big objects of programming. We just ran another copy on IBM, checked it’s working fine and then we had a certain level of confidence that we could turn AWS off, and we just carried on doing that.”

Financial, IBM, Low

“We moved everything from Amazon, where we had been for 2 years, to Google Cloud. Although the migration process was challenging, we moved our entire workload from AWS across to Google because of the significant cost savings. The migration caused headaches, especially when reconfiguring some of our AI services to run on Google’s architecture instead of AWS. But with the lure of £300,000 free credits for 2 years, we pushed through the problems and were able to get fully moved over to Google Cloud.”

IT, Google, Low

3.7. Customer examples

Example 1: Satisfied Microsoft Azure single-cloud customer

Business sector: Retail, consumer, hospitality

Public cloud provider(s): Microsoft

Length of time using public cloud: 5+ years

Annual public cloud spend: High

Annual egress fees spend: Negligible

–

Context

This senior engineer works for a leisure business. Customers can sign up, manage their memberships, access, and book their services through their mobile app and website.

Cloud infrastructure decisions are made collaboratively between the engineering, architecture and product teams and signed off by the CIO (Chief Information Officer). The main considerations are cost and the existing skills and capabilities in-house.

Cloud operating model

This company has a Microsoft Enterprise Agreement – a master services agreement with Microsoft which provides a discount on top of list prices across Microsoft software and cloud services used. They are billed monthly for cloud services, with some services pay-as-you-go based on consumption, plus some Reserved Instances where they get a discounted rate based on commitments to certain level of usage for a couple of years. Their egress fees are negligible.

Their digital platforms and apps are hosted on Microsoft Azure and they use services like storage, SQL, Functions and App Services. They store large volumes of data in the cloud including member and sales information, usage patterns, cancellation predictions, geographic and demographic data.

Switching

This firm already considered themselves a 'Microsoft Shop' when they switched from on-premises to public cloud as they were using products like Office 365 and Active Directory and had an existing relationship with Microsoft. More importantly, their code was built in .NET, Microsoft's programming framework, so operating in Azure was easier and familiar for the IT teams.

This customer is not considering switching provider. He recently evaluated other major cloud providers, like AWS, and could not see any major business benefit to switching given the complexity and cost of switching and the team's lengthy experience in Azure. He feels cloud costs with Microsoft Azure are overall 'pretty good' and there is a lot more value in optimising their current consumption than moving to another provider.

"The cost difference would have to be material (to shift cloud providers) because of the cost to move. It's probably a good seven figure project to move from cloud to cloud. You'd have to save a lot of money within a period to make that worthwhile. Functionality-wise we could do it. Obviously, there's a bit of a revolution going on with generative AI, so if one cloud provider came up with a capability that we wanted to exploit that would have a business benefit, we would act. But we could just have a multi-cloud strategy, right? We wouldn't need to move everything. Or if Microsoft was bought by the Chinese and they shut it down. It would have to be something really dramatic."

Example 2: Siloed Multi-cloud user

Business sector: Transport

Public cloud provider(s): Microsoft and AWS

Length of time using public cloud: 5+ years

Annual public cloud spend: High

Annual egress fees spend: Unaware

–

Context

This participant works for a transport business. He reports to the Group CIO (Chief Information Officer) and influences decisions but ultimately the CIO would have final say. He and the business have no immediate desire to relook at public cloud providers, but spend on public cloud is becoming an issue

so will be more heavily scrutinised by the Finance team in future.

Cloud operating model

They began transferring from on-premises 5 years ago. Everything is basically now in the cloud, but culturally he thinks the organisation is very behind/basic in its public cloud adoption. There is no formal public cloud strategy.

His firm uses AWS for all public facing operations (mainly ticketing software and where/how that's hosted), and Microsoft Azure for all their old on-premises data centres – these were 'lifted and shifted'

to Azure. Despite different usage, both AWS and Azure are used for basic services only – so Compute, Storage, Databases and a bit of Web/Mobile (AWS only).

His firm has a siloed model, with both AWS and Azure being used for legacy/network effect reasons (all ticketing software needs to be hosted on AWS to match the rest of the industry, for example). The two clouds do not need to talk to each other, and no thought has been given to how this would happen if they did.

Switching

He is not looking to switch at all. He sees their organisation as relatively basic users of the cloud – if in the future their usage needs became more sophisticated, he could imagine they might look at issues to do with interoperability, but not in the foreseeable future. He is concerned at rising costs, and at the complex structure around cloud billing. His AWS invoice runs to over 100 pages and needs full time resource in the finance team to stay on top of it.

"When I get the AWS bill, its well over 100 pages long. And you know it's super detailed in every service and every data point. And it's then some poor soul in finance that manages that."

Example 3: Start-up switching between cloud providers

Business sector: Business and professional services

Public cloud provider(s): Google main provider, plus a range of others for AI

Length of time using public cloud: Native cloud user

Annual public cloud spend: Medium

Annual egress fees spend: Negligible

–

Context

This participant's start-up firm provides forecasts for business customers through use of AI. He is one of the founders and the CTO (Chief Technology Officer). He makes decisions around public cloud along with input from the lead engineer and their team.

Cloud operating model

Their core infrastructure is hosted in a third party private cloud data centre and they also use public cloud services. They use Google Cloud as a main provider, using Compute Engine, Cloud Storage, BigQuery, Bigtable, Vertex AI, Machine Translation, with a pay-as-you-go billing arrangement.

For research and developing new AI models, they 'rent' GPUs from multiple cloud providers – Google Cloud, AWS, Microsoft Azure, Digital Ocean, Lambda Labs and others. They use these public cloud providers for 'cloud bursting', when their private data centre is at capacity, using an automated cluster management process to allocate to providers based on availability and price.

Switching

The company accessed start-up credits from the five main cloud providers in their early days. They used credits from Microsoft, then switched to AWS then to Google, deliberately in that order, as they wanted to end up with Google. They did not really use IBM and Oracle credits in the end, so closed those accounts.

This customer sees Google as a superior service in terms of advice, proactivity and problem-solving. He also feels Google provides better performance, with BigQuery having 'by orders of magnitude' faster retrieval than rival services. AWS and Azure have offered free support and credits to switch, but these are not of interest. They have set up infrastructure as code as a fallback, so, theoretically, in the event of a disaster or need to switch, it would take two to three people a week of work to switch, but there is no current desire to do so.

This firm switched from only using public cloud to using private cloud for GPU use to save significant cost, but would switch back to public cloud if the cost of using co-located GPUs would come down.

"If I'm going to spend two or three weeks of engineering time migrating things around – our infrastructure has grown significantly since the last time we migrated away from AWS – I would probably want to use those days of engineering for something else – the product."

Example 4: Switching presented some challenges

Business sector: Telecoms, media and entertainment

Public cloud provider(s): Google main provider, plus AWS

Length of time using public cloud: Native cloud user

Annual public cloud spend: High

Annual egress fees spend: £500K+

–

Context

This participant works for a company that provides media solutions for a range of large multinational brands. As CTO (Chief Technology Officer), he has the final decision on major decisions on cloud infrastructure, but needs approval from the CEO (Chief Executive Officer). There is no formal process, but he still has to provide supporting evidence and a cost-benefit analysis.

Cloud operating model

The customer considers the company as having one cloud provider, Google, although they do use a small amount of AWS services if a client specifically requests this. Single cloud is their preferred strategy rather than multi-cloud or other hybrid options as it offers simplicity and means that they do not have to train employees to use different platforms. It also means they can benefit from CSDs.

The company mostly uses IaaS services 'close to the metal' as they develop their own custom applications as it suits them better (e.g. Google AI stack doesn't meet their needs) and can be more cost efficient. The customer has a good relationship with Google stemming from a previous company and is satisfied with the current set-up.

Switching

When the participant first started working for the company over 5 years ago, they were only using AWS cloud services. The participant made the decision to switch all workloads over to Google because the customer service received from AWS was poor and Google were hungrier for the business and offered more favourable rates. Google services were also seen as much more open source/agnostic which suited the business better.

The respondent had some prior experience of switching from a previous company and understood that it would be challenging. The Google account manager suggested using a third party to facilitate the migration and paid towards the cost of this – which was critical to the whole process. The switch did have its challenges, including in terms of the quality of support from the third party. It took longer than expected, around 4 months in total and they lost 2 months of person effort on the core business too. Despite this, the customer feels it was worth it in the long run.

“It was a bit more painful than I anticipated (switching). I think the third party consultancy, as is often the case with third party consultancies, weren't as good as they promised they would be. So there were random hiccups. I mean, I kind of expected a level of this...and there was a level of things that weren't to a high standard.”

4. Main findings – Theories of Harm 1: Technical interoperability and portability

4.1. Overview of technical barriers

Most research participants point to technical barriers such as lack of code portability, especially in the context of PaaS and databases, when asked about a multi-cloud approach or switching cloud providers.

- 4.1.1. During the interviews we explored the potential role of technical barriers in multiple ways. For example, participants who had never switched cloud provider were asked to explain the reasons why they hadn't done so. Single-cloud users were asked if they had ever considered multi-cloud and what factors were part of that consideration. Likewise, those who did switch provider or do use multiple cloud providers were asked about their experience. When technical barriers were brought up by participants themselves the moderators then probed around these barriers in more detail. When participants didn't mention technical barriers themselves the moderators specifically probed on any potential technical barriers and/or enablers they may have encountered or that they would expect to encounter when it comes to multi-cloud or switching provider.
- 4.1.2. Most research participants in the sample pointed to technical barriers when asked about their consideration of a multi-cloud approach to their cloud infrastructure, or when asked about their likelihood or willingness to switch cloud service provider.
- 4.1.3. However, participants sometimes struggled to identify specific technical challenges that would have to be overcome to achieve a well-functioning multi-cloud infrastructure or a smooth switching experience. This is because for many companies in the sample multi-cloud or switching is not an active consideration, so concerns about specific technical barriers or instances of interoperability were not front of mind and often of a more general nature.
- 4.1.4. As a result, many participants cited more high-level concerns about the perceived difficulty of multi-cloud or switching, such as the amount of coding time and resources that would have to be invested in this effort, and the opportunity costs and disruptions to the delivery of digital infrastructure a multi-cloud or switching exercise would mean to the rest of the business.
- 4.1.5. Nevertheless, the research identified a range of specific technical barriers that participants expect to face were they to attempt a multi-cloud approach or to switch their cloud service provider.
- 4.1.6. Therefore, for many participants these technical challenges were anticipated barriers, which means the expectation of these barriers alone acted as a deterrence to seriously entertain the prospect of switching cloud provider or embarking on an integrated multi-cloud journey.
- 4.1.7. Additionally, the sample included a number of companies who did in fact switch their provider or work with multiple providers in a multi-cloud set-up (to varying degrees), and these participants were able to report about specific barriers they encountered during the migration (or still encounter in the case of multi-cloud).

Portability of code

- 4.1.8. The most common response from participants, when asked about technical barriers to multi-cloud or switching, referred in some way or another to the amount of recoding such a move would require, i.e. the need of the companies' software engineers to rewrite a large amount of code to recreate applications for a different cloud environment or to integrate applications across different cloud environments.
- 4.1.9. The lack of portability of code manifests itself mainly in various ways:
- Differences in interfaces, features and capabilities across providers require different lines of code that has to be learned and written. For example, there is the concern that API Gateways are different and would require a considerable amount of re-work (although there is acknowledgement that APIs tend have been standardised 'on platform level')
 - Authentication methods are seen to be different across providers, suggesting that portability of proprietary Identity Access Management (IAM) tools is particularly difficult
 - Development platforms work differently across providers, resulting in the need for a lot of code to be rewritten. For example, one participant explained the difficulty of migrating from the .NET environment to the IBM equivalent, specifically the fact that certain components exist in one environment but not the other, as well as different dependencies between components.
 - Unforeseen errors resulting from a multi-cloud approach or migration between cloud providers. This concern is often rooted in the different interfaces that cloud providers used and the various dependencies that individual cloud services need to perform error-free.

"The practicalities of multi-cloud are mind boggling, in a sense, because while the different cloud providers have similar services, they're not exactly the same. So you could say, 'well, let's create one code base'. And then in a dream world, we normalise it and then pass it on to [...]. I mean just the practicalities of thinking about how to do that. Maybe there are vendors out there who will offer to do this stuff, who knows?"

Retail, AWS, Medium

"But all the infrastructure is totally and utterly different. The authentication methods are different. The audits, the logging, all the other bits around it, the performance, how you see the performance, how you do the monitoring. All of that is different."

Financial Services, Azure, High

"The API Gateway in AWS and the equivalent in Azure are very different. They are very similar in the core, but they are very different in terms of how the APIs are exposed. The problem is that the surface comes with a different range of APIs. That means that there is an API to do B and C but on another platform they may have APIs to do XY and Z. On the platform side it is easy."

Transport, AWS, High

"What does my application use? Does it use S3? Does it use EC2? Does it use MKS? What's the translated piece where we're going to? And it's not an automated process to then say pick up an application and drop it from one cloud into the other. There's effort needed to go down that pathway because invariably there are dependencies on some of the system services."

Fintech, AWS, High

"They all have different interfaces to start with. So, you know, the calls that you make to execute a particular function are going to be different in the two environments. So you've got that level even if you do a temporary phase migration like we did, moving some .NET code from Microsoft Azure over to the IBM cloud. What we found was, some of the dependencies were actually different on the two platforms, so something that would work on the Microsoft cloud wouldn't work on the IBM cloud. Certain components would need upgrading that didn't need upgrading on the other cloud environment. So moving from A to B, even if it's in theory, the same set of software didn't work."

Healthcare, IBM, Low

Specific interoperability issues with PaaS

- 4.1.10. Participants pointed to the incompatibility of data formats as a barrier even when considering a limited multi-cloud approach, i.e. whereby the company uses one cloud service provider for their main IaaS and another provider for a single PaaS.
- 4.1.11. For example, one participant mentioned Google's analytics PaaS BigQuery as a service they would like to use alongside their main IaaS provider (Azure) but pointed to difficulties in migrating data between the two systems. The participant suggested that there are limitations or inflexibility with BigQuery that make the data less usable for outside the Google framework, which means they were having to actively move and transform the data to make it more usable.

"We're using Google Analytics, and that uses BigQuery, which is in GCP. We're having to look at how we integrate with that, so take the data from BigQuery and make it more accessible to our users. It has some limitations, should we say or inflexibility for our users. It's a case of having to take the data from one platform and put it into another to make it more usable."

Retail, Microsoft, High

- 4.1.12. Some participants went a step further and specifically say that the adoption of PaaS itself is a key factor towards vendor lock-in, citing the different data formats and protocols that PaaS products from different cloud providers require.

“The fact that we now move to PaaS has probably tied us in even more so with Microsoft because, you know, were now utilising their specific service with the files on. So how do you migrate back out of those into something like AWS or Oracle. I’m not quite sure if I’m honest. Whether that has tied us in, I suspect it probably has to a certain degree.”

Financial Services, Microsoft, High

“I think the fact that someone made a decision to use Microsoft’s development tools and .NET that was the main cost issue for us. If they’d have used a development tool that worked across multiple back ends or multiple clouds it would have been a much easier migration.”

Healthcare, IBM, Low

- 4.1.13. Furthermore, while specific design decisions such as building abstraction layers via containers and Terraform modules are seen as a means to overcome lack of code portability (for example, by reducing the amount of re-coding required due to incompatible data formats), these types of approaches have some disadvantages. Specifically, participants said that they have a choice between accepting lack of code portability (and hence greater vendor lock-in) or not taking advantage of vendor-specific PaaS solutions¹. Building abstraction layers while still using PaaS, on the other hand, is considered to require a substantial amount of effort. However, it’s worth pointing out that writing abstraction layers with tools like Terraform may have benefits beyond just portability and reducing vendor lock-in. For example, according to one participant, IaC allows for more standardization and consistency in complex cloud architectures.

¹ Sentence updated to correct a previous proofing error where the word ‘not’ was missing before the phrase ‘taking advantage of vendor-specific PaaS solutions’ in this sentence.

“All the infrastructure is totally and utterly different. The authentication methods are different. The audits, the logging, all the other bits around it, the performance, how you see the performance, how you do the monitoring. All of that is different. Essentially, if you want to be multi-cloud, you also have to be on-prem private cloud and public cloud compatible. So you need to be able to run virtual machines anywhere or containers anywhere. That's the only choice. Which means you can't take advantage of PaaS really, or you've got to duplicate your effort. And that means you need two lots of expertise. People who've got experience in Cloud A and people who've got experience in cloud B.

They sound like they should be easy to move between, but they're actually really difficult to move between, I think because they're so complicated. They have this database service or that computer service or this serverless product. All these kinds of things which are all great products. But by the time you've set that up, then maybe you have some security risks. You put some firewalls and other network configuration around it. Then you have some IP addresses that you have to attach so that your customers can configure their networks and all this stuff, and you start to automate it and the more complex your environments become the more locked in you are really. Because you've invested so much effort to build a system within this ecosystem, that you have to pretty much completely rebuild it if you move to a different system. And so it becomes a really huge cost blocker. It does mean that they have a very sticky customer base.”

E-commerce, Google, Low

“There are ways of writing particular database queries that are cloud specific. We try and keep that level of abstraction high enough that it would work on all cloud providers and don't use anything proprietary and use all of our coordinations as the services based on APIs and Infrastructure-as-Code. So we could take that code and change those APIs to any other provider.”

Transport, AWS, High

Specific interoperability issues with databases and file storage systems

- 4.1.14. Many participants specifically pointed to the technical challenges around migrating database or file storage solutions (or attempting to integrate them across cloud providers in a multi-cloud approach). Some of the challenges most commonly mentioned include:
- a. Different ways to write code (for example differences in writing database queries). While it is possible to work out how to rewrite code, and eventually to do the rewriting, this can take significant amount of time and effort.
 - One participant also mentioned that databases have specific features that are not available elsewhere. Specific examples that participants provided relate to the difficulty of integrating or migrating DynamoDB and SQL databases.
 - Similarly, there is the perception that file storage services such as Azure Files is deemed difficult to migrate or integrate with equivalents from other providers, while storage services such as AWS S3 and equivalents in other clouds also require completely different code across providers.
 - Related to the need to learn and write new code, there was also some concern that new nomenclature would have to be learnt, as providers use different terminology to name and describe their range of cloud services.

- b. Many services don't have automated testing capabilities, so participants said that a migration would entail a very time-consuming manual testing effort
- c. Migrating data from database A to database B may lead to service downtime. Furthermore, participants were concerned that the moment of data migration between database creates a significant security gap as sensitive data are transferred via the public internet
- d. One participant said that a data transfer between databases might potentially be disrupted due to IP geo-blockers (e.g. if one database is located in a country from where data transfers are restricted)

"We use a proprietary database from AWS that is Dynamo DB. And there is no other database with the same properties in other cloud providers that will not imply migrating the data, migrating the codes and changing the codes in order to deal with that data. And some of the features are not available because with this database, there is no match between the features. So this is an example that is not only about the outage, the availability, but also some features that we are using that are specific to [the cloud provider]."

Mobile gaming, AWS, Medium

"There are, for example, ways of writing particular database queries that are cloud specific. We try and keep that level of abstraction high enough that it would work on all cloud providers and don't use anything proprietary. So we could take that code and change those APIs to any other provider."

Transport, AWS, High

"You can set up several database nodes and they always are kept up to date. That might be more difficult if you have two completely different database systems like AWS and Azure. For example, if I spin up a second Azure data centre, they would be able to work perfectly well together, just because they are exactly the same platform. But I don't think it would be the same if I let the second one be on AWS instead, because they have their own type of databases. They might function slightly different so they wouldn't be able to communicate seamlessly."

Retail, Microsoft, Medium

"Digital Ocean have an S3 equivalent. But down to the API, it's literally the same thing, so you just have to tell it to point to somewhere else. But I think if you're a small team, managing two is quite cumbersome because it's just double maintenance."

Fintech, AWS, Low

“You need to transfer data from one place to another. They might block it, they see an IP from, let's say... There are a lot of service places in Ukraine or something else, and they might see an IP in a jurisdiction they don't like it might be blocked. You're not able to transfer your data, especially with the cheaper providers. There are some cheap cloud providers.”

IT, IBM, Low

“There's a lot of testing to make sure that it is resilient and stays where it is. One of the challenges for that platform is it's not all got automated testing, so that there is a manual testing effort.”

Fintech, Microsoft, Low

Data latency

- 4.1.15. Several participants said they would not seriously entertain a multi-cloud configuration in part because of data latency concerns, i.e. the concern that data streams between services within their cloud infrastructure would be slower if these services were hosted by different cloud providers.
- 4.1.16. Data latency in some sectors can also determine initial vendor choice (and by extension the likelihood to migrate away from that vendor). This is most relevant when there is a so-called 'network effect', i.e. due to considerations around data transfer speeds, a company might be more inclined to choose a cloud service provider that is already being used among most clients and/or suppliers within the same ecosystem of companies. This qualitative research sample is too small to make inferences about specific sectors where this 'network effect' plays out more significantly. Concerns about a potential network effect were more likely to be expressed when the company was reliant on very fast data streams and/or when one particular cloud provider was particularly dominant in the company's supplier- and client ecosystem. In this research, concern about the network effect was voiced by participants from the financial industry and a provider of transport data services.

“Every hop in and out of a public cloud provider adds time, it adds latency. It adds potential for things to go wrong. But for us where we have very strict latency requirements on some of our services, we just wouldn't consider a multi-cloud solution because it would just add too much latency to that.”

Transport, AWS, High

“We need to have very fast machines which operate best when they are co-located. And so what we didn't want was the potential for network latency that is inherent in cloud services. Cloud services are excellent. There's lots of cloud services. And like I said, we're looking at cloud services that operate with SLA'S numbered in the microseconds. But when you're dealing with operating at tens of thousands of operations per second, microseconds, data latency is important.”

Retail/IT, Microsoft, Medium

“By virtue of the nature of our trading, we’re always seeking for low latency, and one of the aspects for low latency is to be co-located with the exchanges to which we provide liquidity. And many of them are on one specific cloud provider. So because of that, we’re with that cloud provider. So that way, the physical infrastructure is closest to that lane.”

Fintech, AWS, High

“AWS definitely has a moat to their market share and that the network effects are very strong for our use case because of the market share. And I know Google try is trying to shift market share to them, but we wouldn’t be the ones to lead in that. Because we have to go with where the exchanges are.”

Fintech, AWS, High

“It’s just rules of physics. If you have Amsterdam, which is West Europe, which is a partner with Dublin, the time for a network packet to go from one to the other is 25 milliseconds, which doesn’t sound like a lot, but actually it is in terms of computing and that slows applications down by a factor of 100 or 1000-fold.”

Financial Services, Microsoft, High

4.2. Implications of technical challenges

The implications of multi-cloud and switching that participants are most concerned about include data latency, security and potential service disruption from a switch or multi-cloud adoption.

- 4.2.1. When asked about potential challenges of multi-cloud or switching many participants point to the huge effort involved in developing a multi-cloud approach or switching cloud provider. But regardless of the time and financial investment this effort would entail (or the benefits that multi-cloud or switching provider would bring in return), participants were also concerned about the potential implications of multi-cloud and, to some extent, switching cloud service provider. These implications were mostly brought up by participants during the discussion about potential challenges without the moderators having to specifically ask for them.
- 4.2.2. The concerns about potential implications can be described as ‘second-order effects’, i.e. they arise from the technical difficulties with multi-cloud and switching provider that were explained in the previous sections.
- 4.2.3. Second-order concerns of multi-cloud and switching can be grouped into two categories:
 - a. Security
 - b. Disruption to the IT service

Security

- 4.2.4. Many participants in the sample voiced security concerns, specifically in relation to multi-cloud. Their concern that moving to a multi-cloud model has the potential to undermine the security of their IT infrastructure is rooted mainly in a general sense of loss of control in an architecture design that they aren't familiar with and that they perceive to be complex and difficult to maintain.
- 4.2.5. They were concerned that multi-cloud due to the complexity inherent in a cloud infrastructure with distributed services provided by different vendors and managed by different in-house teams is more prone to security and compliance gaps. For example, one participant raised 'Information Governance' as one potential security risk in a multi-cloud design, as sensitive data flows through a wider array of services by different providers, with teams potentially struggling to keep track of this flow of sensitive information.
- 4.2.6. On the technical side, some participants pointed to the challenge of building secure lines between cloud providers – which is essentially in the 'public internet' – through which their data can flow without being at risk of hacking attacks.
- 4.2.7. Exacerbating both of these risk factors is a skills gap in the industry, i.e. some participants doubted that their engineers would be able to design these network connections between different cloud providers in a secure way, and that finding that type of talent or skill is difficult.
- 4.2.8. One participant also pointed to the security risk of multi-cloud in case of a major event, i.e. when disaster recovery strategies suddenly become pertinent. In this particular example, the participant voiced concern around the lack of an automated failover strategy if parts of their data were located across IaaS and PaaS from different providers.

“My preference would be we use one main cloud as much as possible because we have people who use it. We have processors. It makes our security and compliance a lot easier. [...] As you start to use multi-cloud, the complexity in the environment increases exponentially, as does things like security, compliance standards. It's an exponential growth in terms of the complexity and the effort to manage multi-cloud.”

Transport, AWS, High

“Also, there's a whole load of IG problems there because now you've got to have sensitive data transiting between different platforms. So that's a complexity. [...] Obviously, if you've got very private data, sending it to the cloud requires them to do certain things to make sure it's well protected. In a world where that wasn't an issue, there's no risk because you're still staying within the AWS secure environment.”

Healthcare, AWS, High

“Once you start spreading yourself across multiple clouds, we also have a bit of a skills gap in the networking side because we don’t necessarily have the people who have the experience of how to network clouds together in a secure way. What we do need is we need proper sort of point-to-point VPN connectivity and point-to-point tunnels between Azure and AWS, and that’s proving to be a project challenge. And consequently, it’s a cost. You know, we have to either hire people, network architects or we have to hire contractors with these skills to actually help us over this hurdle. Well, it’s almost like two barriers in one, isn’t it? Because one there’s a security issue, and two, there’s the skills to address the security issue.

If you have the database and the software in the same cloud, you can be really confident that a hacker can’t hack into your database directly because it’s not a publicly available database. Whereas if you have a database in one cloud provider and software in another, then already there has to be a way for your software to talk to the database across the Internet. And if that’s the case, then maybe a hacker can pretend to be your software and gain access to it. So in this moment, when they’re sort of transmitting data and communicate with each other, that’s where the security risk is for potential hacking.”

E-commerce, Google, Low

“For me, the challenge [of multi-cloud] would be around security. I don’t want to introduce extra ways of signing on to other things.”

Manufacturing, Microsoft, Low

“A lot of banks insist on 100 kilometres [between data centres]. So, the idea being if you get a tornado or something that hits you, they’re far away. So, essentially, if a data centre fails, we still have things running and it’s automatically failed over. That’s provided by the platform, both IaaS and PaaS, and we don’t need to do it, it’s easy for us. But if we needed to do it across regions the latency would guarantee data loss so we can’t do an automatic failover. So, it’s a manual failover scenario with a guaranteed data loss. You’re looking at queuing and all sorts of extra complicated engineering and design replaying transactions that are in flow. Even if it’s only five seconds. You know, that’s a load of financial transactions that have been lost. You can’t just lose them.”

Financial Services, Microsoft, High

Disruption to the IT service

- 4.2.9. Participants in the sample were keen to point out that the cloud infrastructure they look after is essentially a service to the rest of the business and any changes that may inhibit their ability to provide this service even for a short period of time needs to be justified with a strong business case. As such, participants were extremely wary of the implications multi-cloud or switching activities might have on overall service availability.
- 4.2.10. Concerns around a potential disruption to the service focuses primarily on potential outages during the data transfer. Some participants said that during the moment of data transfer the service may be unavailable for a short period of time. While this can be managed (e.g. partial transfer, or transferring during the weekend) participants were wary of unforeseen complications that restrict access for example to data bases for longer than anticipated.
- 4.2.11. Participants were also concerned that disruptions are more likely to occur in a multi-cloud setting because of the various dependencies between services that may no longer work if certain updates from one of the providers affects the functioning of said dependencies.

- 4.2.12. Another aspect of concerns about potential disruptions to the IT service relates to the resources a switch or multi-cloud approach would take away from the IT department. This is not a direct technical challenge but rather an exacerbating factor, making a switch or multi-cloud attempt more difficult. This consideration is explained in more detail in the following section.

"When you do this kind of integrated multi-cloud, some things would go wrong. Some things you would not anticipate. There is a level of complexity in the details that is really difficult to anticipate."

Media, AWS, High

"The issue we saw with having multiple cloud providers is you've got to build and maintain all of those integrations. And of course, the minute something changes something else falls over. Also the difficulty of testing in a multi-cloud environment. And given that this is literally mission critical because you're looking after people's lives, you can't afford any outages at all."

Healthcare, IBM, Low

"It's quite easy to end up with terabytes of stuff and shuffling that around is very time consuming. There's some moment where the last request is served from your old infrastructure and the first request has to be cut across and a lot of services these days it's very hard to have maintenance windows. You must be constantly serving traffic. Even a few minutes is considered bad to be down. So that cut across process when you move from one environment to the next is quite difficult to achieve. Because of this, partly because of this database synchronisation even if you have the software running in both environments, there comes a point when you have to say OK, we're going to stop using the database on this cloud system, migrate all of the database content across, start up the database over here, and then start serving requests. That's actually a very tricky process."

E-commerce, Google, Low

4.3. Exacerbating factors

While most participants believe the technical challenges of multi-cloud and switching providers are surmountable in theory, two key factors exacerbate these barriers in practice. First, there is a skills gap – most engineers are only expert in one cloud environment, and talent with experience networking multiple clouds is rare, necessitating expensive retraining or multiple teams. Second, there is a resource gap – participants fear multi-cloud or switching would absorb too much engineering capacity for too long, hampering their ability to work on other business priorities and innovations.

- 4.3.1. During the discussion about technical challenges, participants often also referred to certain factors that are not technical barriers in their own right, but they amplify the effect of these technical barriers.

- 4.3.2. The technical challenges that participants point to in the research in relation to multi-cloud and/or switching cloud service provider were, in most cases, not considered to be theoretically insurmountable. (NB: the impacts on data latency are an important exception to this statement, i.e. many participants see increased data latency as a physical necessity of multi-cloud technology due to the likely geographical distance and lack of connectivity between data centres of different cloud providers.)
- 4.3.3. However, apart from the lack of a strong argument in favour of multi-cloud or migrating to another provider from a business or IT point of view, there are factors that exacerbate the technical barriers discussed above, and thereby make it more difficult to overcome these barriers. The two main exacerbating factors are:
 - a. The skills gap
 - b. The resource and opportunity costs

Skills gap

- 4.3.4. Many respondents employ software engineers with cloud-specific coding skills. This means that a considerable amount of re-training would be required to enable switching or multi-cloud. That training may or may not be included in a negotiated contract with a cloud provider and it takes time away from engineers.
- 4.3.5. Especially in the case of switching cloud provider, the human knowledge that a company's engineering team has built up over many years working within a specific cloud environment may be seen as 'stranded assets' by decision makers – i.e. the company may have invested a lot of time and money into building a body of knowledge with one particular cloud provider which makes it even harder to justify a switch into another environment where a considerable amount of that knowledge is no longer applicable.
- 4.3.6. The research shows that companies with very complex and sophisticated cloud infrastructures that grew over many years are particularly reluctant to consider a switch. But even if such companies were willing to consider a multi-cloud strategy or a wholesale switch of their IaaS provider, participants pointed to the difficult talent and hiring environment that is contributing to the technical barrier to switching or using multiple clouds. They believe they might struggle to find enough qualified engineers and/or would have to invest a lot of money into the hiring process.

- 4.3.7. As for the skills gap for using multiple clouds, participants said that it is extremely rare to find software engineers who are not only expert enough in more than one cloud environment, but also have experience in networking multiple clouds together, considering all the security implications of multi-cloud discussed above. As a result, participants said that they would have to hire two separate cloud teams and essentially run a distributed architecture team, which in turn, raises security concerns.

"It's untypical to find people who are experts in each cloud provider. That's just not feasible and not feasible today. They could learn, you know, if we decided we wanted to move from this cloud provider to another, we'd have a learning curve. We'd work out what all the equivalent services were. We'd do a load of training, you know, we'd put the time in, but in general, you know, someone's a deep expert in one cloud vendor. So the group of people that I gathered around me happened to be experts in AWS."

Retail, AWS, Medium

"You end up having technical expertise within one provider. So we know how to use Amazon tools. We know how to use their platform. You build up a certain knowledge level. So using other tools means you have to kind of unpick that. I'm sure it's absolutely fine. But you'd have to unpick what you've got, people here they've known AWS and their tools so that's the downside, I suppose, it's learning some other way of doing things."

Entertainment, AWS, Low

"We've got eight years of delivery under the belt [with Azure]. We're not going to start that mission of switching. The costs would be absolutely, horrifically large compared with the savings to be made. You know, we were spending what, about £3.5 million a year on heads to build stuff. So, you can say that's maybe a £30 million investment. So, if you said to rebuild the whole lot from scratch, it would take me 30 million quid."

Retail, Microsoft, High

"We have very limited AWS knowledge and experience and zero Google cloud experience. So based off that background, that's why we've chosen Azure over AWS."

Financial Services, Microsoft, High

"I think the thing here is we are so embedded into AWS. It's almost a cultural thing, and our people know it. Our processes are built around it that there needs to be a big impetus for us to move. And I don't see that being simply price."

Transport, AWS, High

Resource and opportunity costs

- 4.3.8. Related to the skills gap is what we refer to as 'resource and opportunity costs', which prevents some companies from considering a multi-cloud approach or switching cloud provider. At the heart of this argument is that in addition to lacking the required skills, participants were also concerned about the sheer resources a multi-cloud approach or migration would occupy.
- 4.3.9. Many research participants expressed concern that switching provider or adopting a multi-cloud architecture would absorb too much of their software development resources.
- 4.3.10. Similarly, migrating cloud provider is believed to take multiple months and occupy most of their software development capacity. This is an opportunity cost as it takes away resources from commercial priorities, such as improving customer-facing services, fixing issues, building innovations, supporting business initiatives etc.
- 4.3.11. As already discussed above, multi-cloud, whether siloed or integrated, is believed to be achievable only with multiple and, according to some participants, separate cloud teams due to the lack of multi-cloud talent. This not only leads to distributed workstreams and a perceived lack of control and security gaps but also potentially requires considerable investment given that the cloud architecture team would almost certainly need to be larger in size.
- 4.3.12. Some participants were also concerned about the practicalities of having to manage the relationships with multiple cloud providers; especially when it comes to the billing side of things, which participants found challenging already with just one provider.
- 4.3.13. Participants were also concerned about the 'unknown effort' – i.e. without expert knowledge in both clouds (in case of a switch) users are unable to predict just how much effort in re-architecting would be required, and what the implications for the rest of the business would be, in terms of potential disruptions to the service (as discussed above).
- 4.3.14. The concern about the resource cost is partly confirmed by those participants in the research who did switch. They said they underestimated the switching effort and ended up working on the architecture re-design much longer than anticipated.

"We thought maybe we could move across in a month of engineering time. One person could work on this for a month or something, but it took us more than three months. It took really a long time. I think once you're on one of these platforms as a medium sized company or, you know, in the very early stage, as I say when we moved across, we were still a company of 10 people and it took us like, three months. Our systems were not that complicated. If we had to move now, it would be even more complicated."

E-commerce, Google, Low

"It took four months thereabouts, five months and again, that was with the consultancy, the third party consultancy that they deployed and funded. We minimised the distraction to my team because we still had our own roadmap to get through. But it obviously did take time out of a lot of people to understand what needed to get lifted and so forth. We probably lost two months of person effort in my team to facilitate that."

Business services, Google, High

“The primary thing was that we don't have the skills to learn or to know how to manage the systems. I only have one infrastructure manager that handles Azure. I would maybe need a second one to handle AWS and that's not in our budget. I just didn't see any possibility to be able to manage that extra workload.”

Retail, Microsoft, Medium

“Maybe you could take some services and put them in one cloud and then take some other services, put them in another cloud and some other services, and put them in the other cloud. You're going to have relationships with each of those vendors which you need to manage. You're going to have to have people who are experts in billing and, securing each of those which is probably number one and then in managing and being efficient in each of those. So you'd have to do that three times the amount that I do it today. Just the complexity of all of that, is just a not a practical proposition.”

Retail, AWS, Medium

4.4. Potential enablers for interoperability and portability

The main potential enablers include using vendor-agnostic services instead of proprietary PaaS to reduce lock-in, containerization and orchestration tools to improve portability, and Infrastructure-as-Code solutions like Terraform. Some service-based enablers like third party consultants and training from the cloud service provider are also mentioned. However, these can have downsides around innovation, costs, and compatibility and many organizations lack expertise to effectively implement vendor-agnostic cloud architecture.

- 4.4.1. While most participants indicate in their responses that the technical barriers to multi-cloud and switching cloud provider weigh heavier than the potential benefits, many do have a view on potential levers for technical challenges that might make multi-cloud or switching easier.
- 4.4.2. The discussion guide used for the interviews, included a set of questions specifically about 'enablers' that might help overcome the barriers to multi-cloud or switching.
- 4.4.3. The bulk of the discussion in the interviews around 'enablers' focused on technological solutions, but a small number of participants also reported experience with 'service-level enablers', i.e. services that were provided to ease the process of migration.

Approaches to overcome tech barriers to multi-cloud or switching cloud service provider

- 4.4.4. By far the most commonly mentioned technical enabler to reduce vendor lock-in is to use vendor-agnostic solutions instead of proprietary services from cloud providers. As mentioned earlier in this report, participants considered PaaS solutions in particular to be lacking in portability so using agnostic alternatives to proprietary PaaS solutions was seen as the most important decision to reduce vendor lock-in.
- 4.4.5. Related to this is the increasing use of containerisation software such as Docker. These tools, that package software applications along with their dependencies into standardised 'containers', make the applications more portable across computing environments and are widely seen as a way to reduce vendor-lock in.

- 4.4.6. Similarly, many participants reported usage of orchestration tools such as Kubernetes to manage containerised applications. This suggests that many users of cloud services are aware of potential dependency and 'lock-in' to a specific cloud provider, with some using such solutions to improve portability (not all participants use containers specifically for portability reasons but the portability benefits were widely recognised).
- 4.4.7. Infrastructure-as-Code (IaC) was also seen as a technological solution to help with the portability of applications across computing environments. In particular, some participants mentioned Terraform as an IaC solution they use, or are at least aware of. Applications like Terraform were also seen by many as a tool that could make multi-cloud more feasible. It is used by companies to automate and standardise the building of their infrastructure environment, and often cited as either a reason why they wouldn't operate a mirrored cloud model, because IaC solutions enable users to recreate critical workloads reasonably quickly, or it was cited as a way cloud users would overcome some of the technical barriers associated with wanting to switch cloud provider or even using multiple cloud providers. For example, one participant described Terraform (i.e. Infrastructure-as-Code) as a way to customise cloud architecture according to their needs by rewriting code into a more generic format, which would then help during a migration to a different cloud provider.
- 4.4.8. One participant mentioned deploying open security standards to connect Azure's IAM solution to their main IaaS provided by AWS. The specific solution that was used in this case is called SAML.
- 4.4.9. Another participant mentioned using virtual machines as a helpful tool to allow for a 'phased migration'. Their cloud infrastructure had been entirely built in Microsoft's .NET environment, and they were using Azure's proprietary Identity Access Management (IAM) tool, which meant that the company had to completely rewrite the code when they switched to IBM Cloud. However, the participant explained how moving the IAM into a virtual machine hosted by IBM allowed them to do a 'phased migration', i.e. they rewrote the code bit by bit, without disrupting the IT service, as the relevant cloud services were still available via the virtual machines.

"If I had the benefit of starting building tomorrow, I might choose agnostic but you know that wasn't really an option back then. We've got 8 years of delivery under the belt. We're not going to start that mission"

Professional Services, Microsoft, Medium

"We increasingly moved away from cloud native orchestration and gone to third parties like terraform because then we can use the same technology for Azure as AWS. Because although we are much bigger in AWS than we are in Azure, we want to make decisions that do not bind us to any particular cloud provider if we can avoid it."

IT, IBM, Low

"We got some practises where we want to be a bit more independent for flexibility, like infrastructure as code. So we use Terraform. We started doing some of our build work in GCP. With that we do a lot of AWS build work with Terraform, and that's not tying us into cloud formations or the Google equivalents or whatever. So we're keeping some independence on how we deal with this going forward, because the complexity is just increasing, it makes more sense."

Construction, Google, Medium

"I think it's easier than it has been to do multi-cloud and that I think there's quite a lot of tooling out there now that helps with multi-cloud, and I think that it's a more advantageous environment than it was."

Transport, AWS, High

It [Terraform] is a recipe that says 'I make my cloud work as I want'. I want to have this type of servers and these types of containers and this type of storage and so on. [...] Obviously, when you make things generic you lose some of the advantages of one cloud over another. There are ways of doing that in Terraform, so that you can say, if you're AWS, then I want you to use these kind of features, and if you're Azure, I want you to use these kind of features. [...] It's particularly good because you can parameterize what you're doing, so you can say I would like this application to run in Big Mode and it will say 'that means you need 10 of these and 20 of these and 500 of these', or [you choose] Small Mode, which will then dynamically change the sizes of everything. So Terraform is sort of almost like a language in that sense. [...] The perfect answer is that we could run half of our application in AWS and half Azure, and you could dynamically move between the two. You would have to be using something like containers, something like Kubernetes, something like Terraform, something very cloud agnostic. Otherwise, you'd be building it twice, and you'd have to be able to have that really tight, very reliable, very low-cost integration between the two things. If you start out with the mindset [using Kubernetes and/or Terraform], so 'I want to be able to run this on any cloud', you can build it like that from the get-go, and you shouldn't need to redevelop it when you want to use a new cloud provider.

Healthcare, AWS, Medium

"Security is all through open standards because Azure is good at that, and AWS provide our main cloud services. But we don't call that multi-cloud deployment. We simply just use different services from different cloud providers. They are either APIs or some open standards for security. There are some other protocols like SAML. They are not APIs. It's just a way of open standards integrating security systems."

Transport, AWS, High

"There was quite a lot of code that needed to be rewritten, components that are available on IBM that weren't mirrored exactly on Microsoft. For instance, the security layer had to be rewritten. Using Microsoft development tools and the .NET architecture means you've got very limited choices where you can go without a complete rewrite. So we actually did a phased migration. Some of the .NET code was migrated over to, I suppose you call them virtual machines, that were hosted at IBM. And over time, some of that code was rewritten. So we didn't use .NET again with the identity management layer. Initially, it was moved over and hosted in a virtual machine at IBM, which gave the company time to rewrite that layer completely, and it took months and months. The switch happened over a weekend, and it that was pretty seamless, but there was a huge amount of development resource spent actually rewriting the code to get off Microsoft's proprietary tools.

Healthcare, IBM, Low

- 4.4.10. A small number of participants also referred to non-technological enablers that may help overcome technical barriers to switching cloud provider. Two types of 'service-based enablers' emerged in the research:
- a. Third party consultants providing support during the migration. These may come recommended by the 'new' cloud provider. While they do a lot of the heavy lifting, participants said that the effort is still considerable on their part
 - b. Support offered by the 'new' cloud provider, mainly in the form of training
- 4.4.11. Participants mentioned these service-based enablers exclusively in relation to switching cloud provider, however, not when discussing ways to overcome tech challenges related to multi-cloud.

"AWS and Azure have migration services, and usually they'll do it for free to try and get you off of another provider. They'll say we'll spin up the right infrastructure that you need, we'll set it up, and then you can test it, and you can deploy your software onto it. The point is that you still have to deploy your software onto that infrastructure. Ultimately, you have to put in the work."

Professional services, Google, Low

"As well as the credits, I think we got some training. And I believe that a couple of the people did go and do some of the Google cloud training courses and they did get some advice from some Google architects for a part of what they were doing. And I think that that was OK, well received."

E-commerce, Google, Low

"What made it easier was, being lent some development resources by IBM. I guess that's the kind of sales incentive to move away from our competitor. So, yeah, those kind of migration services were helpful."

Healthcare, IBM, Low

"I may have eventually done it without that [using a Third Party service], but it would have been a very protracted migration over quarters or years. Because there's just so much work that had to go into first understanding how things happen in Google versus Amazon. The consultancy had that experience. They came in and normalise it to open-source technologies, do a lot of the heavy lifting so that we could deploy the same service in both environments. There were some changes that we needed to make to our code because some of the code had been calling Amazon specific APIs. So they generalised that so it would be open-source technology. It was pretty important. Even when the seller first reached out, my team was like 20 people at the time, 25 people. I had absolutely no bandwidth to bite off something of that scale. They came recommended by Google"

Professional services, Google, High

Downsides and limitations of technological enablers to mitigate barriers to multi-cloud or switching cloud service provider

- 4.4.12. Several participants raised concerns and pointed to limitations and potential downsides of these technological enablers.
- 4.4.13. Notably, some participants spoke about the potential loss of innovation inherent in avoiding vendor proprietary services. Participants expressed the worry that by limiting themselves to vendor-agnostic solutions, they would risk forgoing access to cutting-edge innovation from the market leading cloud providers, potentially impeding their competitive edge within the digital sphere.
- 4.4.14. Related to this, some participants talked about a perceived trade-off between standardization/increased interoperability on the one hand and performance optimization on the other. These participants cautioned that in striving towards greater interoperability between cloud services, companies (or regulators) risk targeting the 'lowest common denominator' in terms of functionality of these services, potentially at the expense of innovative freedom from cloud providers.
- 4.4.15. Financial considerations are another potential downside of expanding the use of vendor agnostic services, as some participants expressed reservations regarding the long-term cost implications of avoiding proprietary vendor services. Specifically, they fear that avoiding proprietary solutions will make it more difficult to achieve committed spend discounts and, as a result, potentially pose challenges to the financial sustainability of their cloud strategies.
- 4.4.16. The research also suggests that there are resource constraints involved with avoiding proprietary services. Not all organizations possess the requisite software engineering expertise to navigate the complexities of vendor-agnostic solutions effectively. The increased coding and development efforts some participants associated with implementing and maintaining such solutions can make the use of proprietary services the more appealing strategy despite their implications for vendor lock-in. This is a particularly salient barrier for start-up companies who, aside from having limited resources, may prioritise agility and a faster road-to-market over avoiding vendor lock-in. The challenge then is to 'break out' from this proprietary ecosystem at a later stage, when the cloud architecture has matured and teams have grown familiar with these services.

- 4.4.17. Furthermore, some participants doubted the efficacy of Infrastructure-as-Code in facilitating multi-cloud environments. While proponents highlighted its potential to streamline deployment and configuration processes, concerns persist regarding its complexity and compatibility across disparate cloud platforms. Secondly, given participants' concern about the 'unknown challenges' of designing a multi-cloud architecture there is some reluctance to considering IaC as a significant enabler. Some participants said that Infrastructure-as-Code solutions often require provider-specific configurations, which means that the main technical barrier to multi-cloud or switching persists (i.e. the need for a major recoding effort).

"Using things like Kubernetes or not allowing people to use Lambda in certain ways is a lot more difficult. So I think it is a lot more difficult for organisations of a certain size or maturity, and I don't think it's necessarily size. I think its maturity as well."

Transport, AWS, High

"So you've heard of Terraform? OK, so if you're going to do that kind of thing, if you'd want to deploy your say IaaS or PaaS service in AWS and in Azure. So you deployed the same server to both and you'd want really the same deployment code to do that with Terraform. Thing is, Terraform says it's multi-cloud, but actually it's entirely different. I don't see any benefit in using Terraform across a multi-cloud deployment, because it's not just different parameters to the same commands. It's totally different commands. So you essentially have to develop everything twice, all your infrastructure is designed twice."

Financial Services, Microsoft, High

"Well, people call it like Infrastructure-as-Code, so basically code to set up your IaaS. And so then this Infrastructure-as-Code stuff is coupled to the particular cloud provider because that has particular services and particular ways that you orchestrate it. If you want to move, you can't just take your code that you have to build one production environment and just build it straight into a different cloud provider. You have to rebuild this kind of stuff and that's effectively like writing a new piece of software which will then build your infrastructure in this new setup."

E-commerce, Google, Low

"But you're in the lowest common denominator scenario. Which is you can't take the innovations of any of the clouds if you do that because you'd have to do something that all the clouds could do."

Retail, AWS, Medium

“If you want to be able to deploy into multi-cloud, you have to pick services that are common everywhere, at which point you are working with the lowest common denominator, as opposed to taking advantage of some of the specialist services that exist in a particular cloud. So, for example, if you were to look at the high throughput, high availability database stuff in Azure, it doesn't exist and replicate into GCP. So you couldn't use that service if you want to be multi-cloud. Suddenly you're losing the advantage of something that's really key for you as a business. You'd have to start building that yourself that you can then deploy in multiple places to keep that commonality. And then you've added another layer of cost onto that and something that isn't going to be as good as something that they've invested massive amount of R&D in. And so you're making compromises as well as adding layers of complexity, layers of testing, layers of cost.”

Fintech, Microsoft, Low

4.5. Conclusions

The research demonstrates that technical barriers have a tangible effect on behaviours around switching and multi-cloud use.

- 4.5.1. When asked about technical barriers to switching cloud service provider or using multiple clouds research participants tended to make very general and high-level comments about the overall complexity of re-engineering cloud architectures.
- 4.5.2. The most commonly mentioned tech barriers relate to the portability of certain code environments, specific data formats that are incompatible across providers, difficulties in migrating databases and file storage systems, as well as networking and connectivity challenges due to inconsistent APIs and interfaces.
- 4.5.3. Participants voiced particular concern around data latency, security gaps and the potential for a switch or multi-cloud strategy to absorb so many resources that it would disrupt the IT department's ability to service the rest of the business.
- 4.5.4. Exacerbating factors to these technical challenges are a major skills gap (both in their current team and a difficult hiring environment) and a resource gap, i.e. the engineering resources that switching or multi-cloud use would take away from their core business.
- 4.5.5. While some enablers were mentioned – both technical and service-based – participants don't appear to see them as completely offsetting the technical challenges to switching or using multiple clouds.
- 4.5.6. The most commonly mentioned lever to overcome vendor lock-in was the increasing use of vendor-agnostic services.
- 4.5.7. However, some participants voiced concern around this strategy, which they say may lead to restricting themselves access to cutting-edge innovation provided by proprietary vendor services. A small number of participants even went as far as saying that encouraging greater open-source standards may discourage cloud providers to innovate if they're incentivised to prioritise the 'lowest common denominator' over the most innovative solution. However, this was a minority opinion.

- 4.5.8. Overall, the research clearly shows that there are very significant technical barriers that reduce participants' willingness to consider switching cloud provider or going for a multi-cloud strategy. However, most participants mentioned these barriers in the context of *hypothetical scenarios* in which they would seriously consider a switch or multi-cloud approach. As this research shows, most participants do not see a strong argument in favour of a switch or a multi-cloud strategy in the first place.
- 4.5.9. Furthermore, as some switchers or multi-cloud users in the sample demonstrate, these barriers can be overcome in practice, with enough time, effort and resources. However, there doesn't appear to be sufficient incentive for most participants to invest this time, effort and resources to switch provider or implement a multi-cloud architecture.
- 4.5.10. It is also important to note that, while participants expressed concerns about specific technological challenges, their implications and shortcomings of potential solutions, there is also a more general fear of the 'unknown'. From the research it appears that many cloud users anticipate unforeseen issues if they were to attempt multi-cloud or switch provider, which introduces risk into the process. That anticipation of unforeseen issues alone can act as a deterrent to multi-cloud or switching.
- 4.5.11. In summary, for most participants, the tendency to reject the idea of multi-cloud or switching provider is based on the view that the benefits of such a strategy are relatively low and the technical barriers standing in the way are relatively high. This view is compounded by the fact that switching or multi-cloud is rarely top of their mind. As a result, participants often have only a hypothetical view of potential benefits, while at the same time fearing both the anticipated and unforeseen challenges that would have to be overcome.

5. Main findings – Theories of Harm 2: Data egress and transfer fees

5.1. Overall attitudes towards data egress and transfer fees

Egress fees are not a major concern for most participants, whether this relates to setting up their cloud infrastructure initially or their ongoing use of cloud services. However, several participants note a lack of transparency and control over data transfer fees, and some have begun making changes to their cloud architecture design when the scale of egress fees became apparent. There is also some concern around the role of egress fees in – mostly hypothetical – switching or multi-cloud scenarios.

- 5.1.1. A substantial amount of time during the interviews was spent discussing the financial arrangements participants have with their cloud provider(s). As part of this discussion, the guide included a series of questions specifically focusing on egress fees. In most cases, the moderators needed to probe on this as only a few participants brought up egress fees themselves.
- 5.1.2. For many participants egress fees are not a major cause for concern in terms of provider choice, both in terms of their initial provider choice or their ongoing use of cloud services. In fact, many participants were unable during the interviews to recall details of the amount of egress fees they pay on a yearly basis and could only give broad estimates, confirming that egress fees are not top of their mind. Attitudes among participants towards egress fees ranged from a 'necessary evil' to a 'cost of doing business'. Few participants considered egress fees to be an unfair practice.

"It's such a low fee that we pay now. If it was higher, yes, it probably would be [influencing our choice]. It's maybe just a few percent of our spending."

Retail, Microsoft, Medium

- 5.1.3. Furthermore, across the sample, almost no participant considered egress fees to be a factor in their cloud provider choice. This is unsurprising for companies where egress fees are so low that decision-makers consider them negligible compared to their overall cloud spend. Additionally, egress fees are barely a factor in provider choice even among companies who do accumulate considerable egress fees. This is partly because all major providers charge egress fees and cloud users tend to view them 'cost of doing business'.

"We're a B2B company. Our traffic volume is not the same as a B2C. On the business side of things, it really varies. So egress fees, because of the typical style of business we are, don't necessarily worry me too much. Eventually we'll get to some scale, but I think right now it's a bit too low for that."

Fintech, AWS, Low

“It honestly isn’t looked at. It’s just an accepted cost. And I don’t believe the fees are that much. About 11% to 12% of the total charges for the month.”

Fintech, AWS, High

- 5.1.4. The finding that egress fees don’t play a major role in participants’ cloud strategy does not mean, however, that they do not matter at all, including when it comes to switching cloud provider or considering multi-cloud. Section 5.2 gives more detail about the ways egress fees may impact switching or multi-cloud behaviour. In addition, there are some issues around egress fees beyond just switching and multi-cloud.
- 5.1.5. In particular, several participants voiced dissatisfaction with the lack of control cloud users have over egress fees. The perception is that providers can set them at whatever level they want, without justification or warning, and they are incurred in ways which are outside the control of the company. This suggests a general sense of a lack of transparency and a lack of control, which is compounded by the billing complexity. Some participants said that they lack the ‘FinOps skills’ in their team (i.e. skills in Financial Operations) to really understand the scale of egress fees, while others have managed to reduce their egress fees through some changes in the way they use cloud services.

“I think the problem is that egress fees are very hard to calculate upfront. [...] At the moment we’ve got three main databases, two of which are in Azure and one’s in AWS, all hosting customer data. [...] So, there’ll be a constant interchange of data. Now, can I put a number on what that’s going to cost me? Simple answer is no. I don’t have the technical skills to understand the data transfers across the network. And I don’t think any of our network team really have that level of skill and knowledge in order to be able to cost it out and work out what the Volumes will be.”

Retail, Microsoft, High

- 5.1.6. However, this is not a uniform perception with some participants talking about helpful cost tracking tools that are provided by the cloud provider. There was also some mention of cloud providers making helpful suggestions to cut egress fees, while other participants said that various cost optimisation strategies on their part manage to reduce the egress fees substantially.

“There’s decent cost management in there. It’s not opaque, so that you suddenly get a big bill and don’t really have much choice but to pay. I think we can predict it well and we know what it is. We budget for it, and we don’t go over.”

Professional services, Microsoft, Medium

“I know we can do a lot on the Microsoft platform; they have a great tool. I can look at my spend over the past day, week, month, year, you know, various different options. I can see exactly what I’m spending my money on. It also provides me with an interface to look at cost savings as well.”

Charity, Microsoft, Low

"We've cut our costs with Azure from £350,000 a month to about £55,000 a month, and that is because we were incompetent, didn't have the governance and all the other appropriate ways of deploying things. And now we've done it intelligently and looked at the cost and picked the right things and designed it right, you know, seven times cheaper."

Financial Services, Microsoft, High

- 5.1.7. Related to this is also the perception among some participants that egress fees are hard to predict and can lead to unforeseen costs. This opinion was particularly strong among companies with high data volumes, for example when consumers download data out of their cloud, which would typically incur egress fees.

"It's challenging. A good example is this week. We looked yesterday. The bill was £237,000. We look today. It's £227,000. So it dropped £10,000 overnight. Why has it changed by £10,000? As you progress through the market changes to the predictions aren't great. So it's very hard. And I think when you were talking £2 million it's quite worrying. That is definitely a lot of money to not have visibility over for sure."

Financial Services, Microsoft, High

- 5.1.8. Some participants also said that comparing costs between providers' egress fees is extremely difficult, which is one reason why the egress fees (i.e. the specific unit prices) barely factor into the cloud provider choice.

"When we compared the prices at that switch over, we did not look at egress fees at all. We didn't think of egress fees. I believe we just compared storage prices and CPU prices. Now we get charged quite a bit of egress fees [...] What's very hard, though, is to understand some months, it's up, and some months it's down. It's quite significant percentage changes and is very hard to understand exactly what's causing it."

E-commerce, Google, Low

- 5.1.9. It is worth pointing out, however, that there is also a degree of acceptance of egress fees, i.e. the view that data transfers do naturally incur costs and cloud providers are expected to charge customers for this. One participant expressed the concern that mandating the removal of egress fees between cloud providers would only benefit a small number of multi-cloud customers who incur ongoing egress costs as a result of their multi-cloud strategy. In response, cloud providers would likely push the costs for data transfers into other cloud charges, thereby disadvantaging single-cloud companies who make up the vast majority of cloud users.

"They have to cover their costs somehow, and it's a waterbed. So if you lower costs in one place, they're going to come up elsewhere. So by lowering or reducing egress fees, there's a very high risk that you benefit the really large platforms, the ones who can have multi-cloud and those costs ends up falling to us. It's smaller providers who can't benefit from those advantages and therefore have to wear costs in other areas. I do worry about the unintended consequences of looking too closely at egress charges."

Media, AWS, High

5.2. Egress fees impacting multi-cloud or switching behaviour

While egress fees are rarely the primary barrier to switching or adopting multi-cloud, they do contribute to some companies' reluctance in certain cases. Typically, this involves high data volume companies or those with accumulated years of cloud usage worried about one-time exit fees or ongoing syncing costs. The rise of AI has the potential to move egress fees more into the spotlight.

- 5.2.1. After asking participants about the role of egress fees in their cloud usage in general, moderators asked specifically about their views on egress fees in the context of switching cloud provider and multi-cloud use. In very few cases, participants brought up egress fees as a potential barrier themselves.
- 5.2.2. Some participants in the research said that egress fees play a part in disincentivising them from potentially considering a migration to another cloud provider or considering a multi-cloud approach. However, it is important to note that in almost no case were egress fees the main or even one of the main barriers to switching or multi-cloud.

"It's expensive to move data out of a cloud platform, and the issue with that is that it impacts the decision as to whether or not you were going to move wholesale from a cloud provider."

Media, AWS, High

- 5.2.3. For most participants, the technical challenges (alongside exacerbating factors such as the skills- and resource gap), combined with the lack of a clear business case, were the main barriers to switching and using multiple clouds, with egress fees in some cases contributing to the reluctance to consider a switch or multi-cloud strategy.
- 5.2.4. 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 of IaaS and PaaS usage.
- 5.2.5. Some participants voiced concern about one-off 'exit costs', i.e. when switching cloud provider and transferring a large amount of data into the cloud of the new provider, while others comment on the ongoing egress costs in relation to a potential multi-cloud strategy. These participants were particularly concerned about the costs of keeping large databases in synch across different cloud providers.
- 5.2.6. It is worth noting that most research participants commented on *potential* implications of egress fees for switching or multi-cloud in *hypothetical* terms, i.e. they explain how egress fees *might* impact their behaviour if they *were* to switch or use multiple clouds.

"I don't think we look at that egress fees so much...I don't think it's often a forefront of mind conversation. We look at the more obvious charges of compute and storage and things like that. [...] But when moving to another cloud there's a lot of data behind [our] platform because it's a lot of video of user journeys on people's mobile banking apps etc... So, there's a lot of data that would need to move as part of that. There would be a fairly hefty egress cost to move that amount of data."

Fintech, Microsoft, Low

"No, it wouldn't be things like that [egress fees]. If they did suddenly start, if they said, Oh, well, the prices we are going to charge you for this. Then again, that might force our hand"

IT, IBM, Low

"[Hiring talent] dwarfs the cost for us. Maybe for other people you know who are spending millions on egress fees for the migration. Only if you have a huge amount of historic data and that's a one off. We don't have anywhere near that. And our egress costs by running multi-cloud, as we are at the moment, between AWS and Azure are 10 times higher than what it would take to migrate one off."

Financial Services, Microsoft, High

- 5.2.7. In a few cases, participants reported specific ways egress fees impact on their cloud behaviour. For example, one participant specifically spoke about their plans to migrate from GCP to AWS and how egress fees affected their decision on which data to transfer over.

"We've got our data warehouse and everything in Google, and now we're basically moving our data centres and putting all the source software in AWS. And now there's going to be egress charges at some point in the future going into Google and at this stage, I don't know the impact of that, but it's obviously going to be a key decision maker. And actually, you know, how do we minimise and control that? It would be nice if as Gartner predicted, these charges disappeared. And we could just consume the best of whatever cloud in the way that we want to do it, and just tap off that ultimate flexibility with it. But at the minute we've got to really think about moving the data around. So it does affect decisions."

Retail, Google, Medium

- 5.2.8. One participant spoke about his company's use of Google's BigQuery and the limitation of integrating this with the rest of their cloud infrastructure. The main reason for keeping BigQuery completely separate are egress fees.

"I talked about us using BigQuery. When we looked at the cost of pulling it out to put it in something else, that fee was seen as not logical. So it's kept separate. So as with 365 and Microsoft, Google Analytics plays a lot better and is a lot cheaper if you use BigQuery. And there are other services around that, which is partly why we have some Google estate because it is cheaper. Because if we were to take that data out of Google Analytics, put it in AWS, that is not an insignificant amount of money to do that egress on a daily or more frequent basis."

Transport, AWS, High

- 5.2.9. One participant reported their dissatisfaction with having to use an AWS-native back-up solution for their databases that are hosted within AWS. They would have preferred using a back-up solution from a third-party cloud provider (essentially a multi-cloud for mirroring purposes) but was discouraged from doing so because of the prospect of egress fees.

"You've got to take all of your data out of your AWS environment and put it into their [backup provider] environment. If we're talking 60-70 terabytes, that's a lot of data. That's continual data every single month. So you've got to pay for your backup solution with the backup provider. You then got your licencing costs with that vendor. Plus the storage costs associated with that. So the egress and ingress fees that they are also charged are directly passed on to me under the guise of licencing costs. So I'm now having to look at AWS native backup, and I have no other choice but to go with AWS native because of the egress fees. And that's just one example of how you get locked into these platforms. The backup solution that's being offered is OK, it's fit for purpose. Does it come with all the bells and whistles and the nice to haves that I would want? No. I should be able to have more, but cost is now prohibitive."

Retail, AWS, High

- 5.2.10. Another participant said that their company wanted to explore using Azure OpenAI and integrate this with their IaaS on AWS but was put off by the potentially high egress fees this might incur. In fact, there was a general concern among some participants that the increasing importance of AI, alongside the expected larger data volumes, will put egress fees more into focus in future.

"We'd really like to use the Azure AI service particularly that powers ChatGPT, the OpenAI service. But to do that on top of AWS documents means us sending the documents to Azure to get them processed to then send the result back. Now that is going to cost us a lot of money in the egress fees if we do that."

Healthcare, AWS, Medium

“If you had a situation where you were heavily Microsoft orientated, it might be very advantageous to drop your ERP into Azure. But likewise, you wanted to use the data science capabilities in Google. The egress charge will probably stop you from doing that in terms of trying to take all the data from there and move it to there to process it. To get some AI would get very, very expensive. You could as an enterprise have the Microsoft benefit over here, but use the Google benefit over here. It will be those egress movement charges that probably stop you from doing that, because it would probably then make it cheaper to run your AI in Microsoft.”

IT, AWS, High

- 5.2.11. Beyond the egress fees that companies incur when moving data between cloud providers, some participants also voiced concern about egress fees to the public internet, and other types of data transfer fees such as those incurred during inter-region and inter-availability zone data traffic. This affects cloud architecture decisions beyond switching or multi-cloud.

“Our total data transfer charges would be something like \$110k for the year, probably 7-10% of our total spend. So that includes egress to the public internet but also inter-region and inter-availability zone traffic within our own private accounts or network. It's a very big number, unreasonably big number. It is making us consider compromises to our architecture [...] And we are looking at introducing Cloudflare in front of CloudFront because they don't have these very high egress charges [...] but that's more complexity so it brings some other costs to us apart from just the actual cost of the infrastructure. And maybe it's less secure.”

Retail, AWS, Medium

5.3. Conclusions

The research demonstrates that egress fees have the potential to affect switching and, to a lesser extent, multi-cloud behaviour. However, they are rarely the most important part of the switching or multi-cloud consideration.

- 5.3.1. Egress fees were rarely top of mind when participants were asked about their openness to switching cloud provider or using multiple clouds, or even when specifically asked about potential challenges for a switch of multi-cloud strategy
- 5.3.2. In fact, those participants in the research who did switch or use multiple cloud providers for specific use cases, considered the egress fees they incurred a price worth paying to deliver the cloud strategy that makes most sense for their business
- 5.3.3. However, for some companies egress fees do disincentivize them from potentially considering a switch to another provider. In particular, companies with high data volumes or a large amount of historical data can be put off by a potentially large one-time 'exit fee'
- 5.3.4. When asked about the role of egress fees for a multi-cloud strategy, some participants were concerned about the potential costs of keeping databases in sync. While this is usually mentioned in the context of a *hypothetical* multi-cloud scenario, this consideration does suggest that egress fees might act as a barrier or deterrence to even considering a multi-cloud strategy.
- 5.3.5. In a similar vein, even limited multi-cloud use cases such as using a different cloud provider for backup solutions may be discouraged by egress fees

- 5.3.6. In addition to the amount of egress fees, participants also voiced concern about the lack of transparency and overall complexity of egress fees, which makes it difficult for some to predict costs for their general cloud usage. This is relevant in the switching and multi-cloud context, as the lack of transparency and control makes it difficult for participants to include egress fees in a potential cost-benefit calculation for switching or using multiple cloud providers.
- 5.3.7. In most cases, however, egress fees are rarely the main reason why participants don't consider a multi-cloud strategy or migration to another provider
- 5.3.8. In summary, egress fees do have the potential to affect cloud users' willingness to switch or multi-cloud; however, the technical barriers, alongside the exacerbating factors and concerns mentioned in section 4, as well as the perceived lack of a strong argument in favour of switching or using multiple clouds, tend to weigh heavier than any potential concern around egress fees.

6. Main findings – Theories of Harm 3: Contracts with committed spend discounts

6.1. Overview of different types and levels of discounts

A wide variety of discounts are either in place or on offer from different public cloud providers. Where companies have chosen to benefit from a discount or discounts from cloud providers, these discounts could have an impact on how those companies use the public cloud services on offer

- 6.1.1. Part of the depth interviews included an extensive discussion about the contractual arrangements the participants have in place with their cloud service provider(s). The moderators asked about contractual arrangements in an open-ended way, and participants typically focused their initial response on the contract duration. The discussions then tended to focus on the financial aspects of their contractual arrangements, with the moderators specifically asking about any discounts included in the contract. Where needed, - i.e. where this wasn't brought up by participants themselves – moderators asked specifically about Committed Spend Discounts (CSD), the monetary value of their CSD, and to what extent it shapes their decision-making regarding their cloud strategy.
- 6.1.2. However, it is clear from the language used by the participants that not only are there a variety of different discounts on offer from public cloud providers, but that in some cases the same type of discount can have different terminology to describe it depending on the public cloud provider in question.
- 6.1.3. When we factor in the possibility that these different types of discounts can be, and are, offered in combination with each other to individual customers, it provides some explanation for the complexity that some customers described in their contractual arrangements with cloud providers. We will talk more about this in the next section, as this perceived complexity can impact on behaviour.
- 6.1.4. Later in this section we will also outline why some of the participants in the research were not taking advantage of any form of discount. At this stage, a key point is to note that a substantial majority of the research sample were enjoying some form of discount.
- 6.1.5. Firstly, it is important to outline the different types of discounts that the participants described as being offered by public cloud providers.

Committed Spend Discounts

- 6.1.6. Not always referred to by this term, but when prompted by this term within the interview, it is instantly recognised by participants as a discount which is offered based on the contracted spend levels with a particular provider over a specified period of time. Across the interviews conducted, customers of Azure, AWS and Google all reported receiving some form of committed spend discount, and, while the specifics of the discount can vary widely from customer to customer, their existence in the public cloud market appears to be extensive. While we did not speak to any customers of Oracle and IBM who reported being offered committed spend discounts, we cannot be certain that IBM and Oracle do not also offer committed spend discounts to customers.

- 6.1.7. Spend discounts seem to be in the range of 5% – 20%, with other types of discounts tending to be much more variable. Where these discounts are in place, participants reported that they are generally agreed following negotiation with the public cloud providers. While individual participants were not able to compare relative discounts across companies, the general pattern seems to be one of higher discount levels for the higher spenders. It is clear to participants that higher discounts are on offer the longer their company is prepared to commit to a particular provider. Contract lengths are typically 1-3 years, with almost no one in the sample saying their company had committed to anything longer than this.
- 6.1.8. Companies who have entered into a discounting contract with their public cloud provider (or providers) tend to be the higher users and these companies were also the larger companies in the sample.
- 6.1.9. In some cases, participants talked about these discounts purely on the basis of their total spend levels with their public cloud provider, whereas in other cases, they were described as relating to spend on data storage with their provider, and in that sense distinct from the other forms of discount on offer. It is not clear from the interviews whether this difference is a reflection of the role of the individual spoken to (some participants in very large organisations admitted they were not familiar with all the detail of the contract their employer had entered into), or whether it reflects real differences in how discounts and contracts are structured.
- 6.1.10. It is likely that both of the above explanations are true, and that what is clear from the research is that bespoke deals are on offer to public cloud users, based upon the negotiations that take place.

Reserved Instances or Committed Use Discounts

- 6.1.11. Less common, but still referred to by a substantial minority within the sample, are Reserved Instances or Committed Use Discounts (CUDs). Reserved Instances are most commonly associated with AWS, although Azure users will also use this term. CUDs is a term used by Google to describe the same thing as a Reserved Instance.
- 6.1.12. Both terms refer to public cloud users' usage of compute or resource within the public cloud that they are using. Companies reserve processing or compute power with their provider for a set period of time, and by doing so they receive a discount off the standard rates published by the cloud providers.
- 6.1.13. Participants reported that the discount rates applied within a CUD or Reserved Instance may well be different to the discount rate applied within a CSD. This means cloud users may be managing contracts with different discount rates for compute/processing (CUDs/RIs) compared to the CSD which covers, among other things, data storage (which those same users often reported to be the largest component of their spend with their public cloud provider).
- 6.1.14. Participants based in larger companies with high cloud usage, spoke of these arrangements alongside other types of discounts, describing a number of strands or layers to their discounting arrangements, which some participants admitted is hard to monitor and keep on top of.
- 6.1.15. Across the sample, there were differences in the way types of discounts were perceived, with CSDs being seen as bigger commitments than CUDs or Reserved Instances. Companies who are beyond the start-up phase, but who are still in a fast growth phase of their development appeared more likely to commit to a CUD or Reserved Instance than a CSD. Many of these companies in the sample tended to be either cloud native businesses (E.g. mobile gaming, shared economy etc...) or simply still relatively small but with ambitious plans for growth.

- 6.1.16. Their reasoning is that they are established enough to know that they will need to use their public cloud provider's virtual machines at least as much as they are currently and probably to a greater degree, but not so clear on their growth trajectory that they are prepared to commit to spend levels over a period of time.

"There was quite a big push [from AWS] for us to enter into an enterprise discount plan. Basically, it's another one where you commit to a certain level of spend generally for a longer period of time. So in 3 years, supposedly, you start getting decent savings. But in return, we had to purchase their enterprise support, which would have basically wiped out a large portion of that discount. So we tend to go for quite short commitments and pay nothing upfront. Our standard commitment is a year. The benefit to us is that we make some quite significant savings, particularly on reserved instances"

Retail, AWS, Medium

Credits

- 6.1.17. Credits were primarily mentioned by businesses in the start-up phase of their lifecycle. These credits can be defined either in financial terms (E.g. \$100,000 credit) or as a fixed time period (E.g. free for first two years).
- 6.1.18. Participants from start-up businesses described the cloud providers as quite active in trying to win their business. One business spoke about using one cloud provider for 6 months (IBM) and then switching to another (Google) because the latter provider offered them a 2-year free deal if they moved everything to Google.

"Google approached me six months ago and they gave us \$300,000 credits for two years. Google approached me and they said you know this is free for two years I say let's do it. They have this programme for start-ups and I say OK great"

IT, Google, Low

- 6.1.19. In addition to credits for start-ups, there is also evidence that providers offer credits for migrating more of a company's on-premises workloads into the cloud. These credits are offered to mature organisations who have not yet fully migrated everything onto the cloud and can be offered in addition to a CSD. One participant, working at a large global financial services company which operated on a siloed multi-cloud basis, spoke of how they decide which cloud provider to migrate these additional workloads to, based largely on the value of the credits being offered by the provider. The implication here was that there was very little differentiation between the providers, other than the monetary value of the credits they were offering.

"There may be some services that we get even bigger discounts on. There may also be credits to help us migrate new services into the cloud. We've got a package with both of them that helps us to move things. And also, we've got consistent pricing from a discount perspective."

Financial Services, AWS and Microsoft Azure, High

- 6.1.20. The quote above clearly illustrates the different types of arrangements that public cloud users enter into with the providers and the multiple layers of discounts/credits on offer, particularly for the more mature businesses, with sophisticated, complex needs from their IT infrastructure.

Other discounts

- 6.1.21. In addition to the discounts and credits outlined above, participants also mentioned other types of discounts which will form part of the overall contract a cloud user has in place with the provider. While these discounts tended not to be mentioned as widely as broader spend based or usage-based discounts, it should not be assumed that these discounts are unusual. It is just as likely that the relative infrequency of them being mentioned in our sample is due to how knowledgeable the individual participants were about the specific details of the contracts their company had in place.
- 6.1.22. Our analysis does suggest that the multi-layered, more complex discounting structures tend to be experienced by higher cloud users, although even some start-up participants reported difficulty in being able to fully understand the financial arrangements they have with their provider.
- 6.1.23. One example of a more specific discount is illustrated by the quote below:

"We found that even each of those divisions has the ability to negotiate the discount level separately. So at AWS, for example, their compute division might give you more discount than their database division than their storage division. So especially with our AWS contract, we've got different discount levels, depending on the division even."

Financial Services, AWS and Microsoft Azure, High

- 6.1.24. Given the wide variety of IaaS and, in particular, PaaS services available from public cloud providers, it is easy to see how, if different discounts can be negotiated on a service-by-service basis, public cloud contracts can become very complex to manage and monitor.
- 6.1.25. There is no evidence to suggest that this complexity is any better or worse with any one provider compared to another. As the quote below illustrates, another research participant has experienced multiple different discount levels with Google.

"I think we get a 22% discount on the normal compute instances, we get a 15% discount on networking, 15% thereabouts on BigQuery. On our commitment if we promise to buy this certain processor or memory it's a 40% discount on a three-year commitment."

Advertising, Google, High

- 6.1.26. As well as service-by-service discounts being offered by providers, the research also shows that discounts can be offered on other elements, such as egress fees. As with the service-by-service discounts, egress fee discounts are offered as part of an overall package, not in isolation.
- 6.1.27. An example of this can be seen in the quote below:

"When you commit to something with AWS, I think that's where you get discounts up to 70% – for example, where you commit to a three-year agreement, you agree to spend a certain amount, and then you get 40% to 70% of discount on traffic [i.e. egress fees]."

Mobile Gaming, AWS, Medium

- 6.1.28. The final area of discounts mentioned by participants relates to incentives offered by the cloud providers to use their marketplaces. In particular, these were mentioned in connection with AWS and Google, but it is possible they are also offered by other providers.
- 6.1.29. One participant mentioned that AWS offered their organisation credits if they bought third party services through the AWS Marketplace. These purchases count towards the overall spend with AWS, but the incentivisation (in the form of credits) is different to a straightforward discount rate on spend levels (as with a CSD). This particular participant went on to describe how they track this aspect of their spend alongside all other aspects. As they were a sophisticated Fintech business providing software to the financial services industry, they had the in-house capability to build their own Fin Ops solution to track their spend with AWS and better optimise their public cloud usage. This option (of developing in-house solutions to monitor cloud spend) was not common or practical for most research participants and goes some way to explaining why many of the participants said they find it difficult to be fully on top of what they are spending with their provider(s).
- 6.1.30. As well as saying monitoring cloud spend suffers from a lack of transparency, some participants also spoke about the multilayered contractual arrangements not necessarily aligning with each other. This can mean that public cloud users find it hard to know exactly when they are technically 'out of contract' if, for example, one element of a contract has come to an end, but another remains in place. Are they better off renegotiating the element that has finished, or waiting until all elements have lapsed and they are fully out of contract? This is just one example of the complexity faced by public cloud users. This is illustrated by the quote below:

"I could go into Google's billing system now and show you what it's recommending. We change to get a better discount but I know from looking at it probably about a month ago it's got somewhere in the region of 400 recommendations. 300 of them will save me 15 cents a month or less. If you look at the time that it would take to sit there every month and go through those – that doesn't pay somebody's salary."

Healthcare, Google, Medium

Why no discounts?

- 6.1.31. A number of participants, especially those in the Low spend strata (who tend to also be smaller organisations) reported that they have either been offered discounts and decided not to take advantage of them or that they have not yet been offered discounts.
- 6.1.32. Those who have not yet been offered discounts are often start-ups in the early phase of their cloud usage and may well be benefitting from credits offered to them at the start of their public cloud usage journey, as described earlier.
- 6.1.33. Those participants who have actively decided not to accept their cloud provider's offer of a discount, are often doing so for broader reasons relating to where their business is in its development, and what their cloud strategy is at this point in time. These businesses, often in a steep growth phase, are less able to predict where their business would be in the next 1-3 years and so do not want to commit to a contract with a cloud provider at this stage. The mindset of these same businesses also tends to be that they wish to retain their cloud strategy independence for as long as possible. It is clear from listening to these businesses that, although they can see the financial benefit of the discount, they see the decision as being difficult to back out of, so that once committed, it would be likely to lead to a long term, deep, relationship with that provider. This sentiment is illustrated in the quote below:

“As an individual user of Amazon we are wholly pay as you go. We’ve looked into savings plans and the reason we haven’t bitten the bullet on a saving plan yet is because we’re constantly moving and we’re constantly growing – I think we’re doubling revenue every single year for example so we can’t commit to a year because in six months that might be useless to us.”

Retail, AWS, Low

6.2. Implications/impact on behaviour

There are two main ways in which discounts have at least some impact on cloud customer behaviour. Firstly, participants say that it does influence propensity to stay with their current provider – at least to some extent. This can be down to the complexity of their contractual arrangements making comparison with other providers difficult, or it can simply be due to the certainty that a committed spend discount gives a business, disincentivising looking for any alternative. Secondly, there is evidence from across the sample that the existence of behaviour-based discounts influences how companies actually use the services offered by their cloud provider. This section explains both of these effects.

Impacts on potential cloud switching behaviour

- 6.2.1. After participants explained the types and value of discounts they have in place with their cloud provider(s), moderators shifted the conversation to the role committed spend discounts in particular played in shaping their willingness to consider switching provider and/or attempting a multi-cloud strategy.
- 6.2.2. There is evidence from a broad range of participants that the contractual arrangements they have entered into with their public cloud provider(s) do make it harder for them to contemplate switching from those providers. This is rarely, if ever, spoken about in a way which suggests public cloud users are locked into contracts so that they are denied the opportunity to shop around for a better deal. Instead, participants are much more likely to talk about the existence of discounts disincentivising looking at other public cloud providers for their business.
- 6.2.3. This view should be read in the context that very few of the medium and high strata users are looking to switch or perceive any potential benefit in switching their public cloud use to another provider. When we take that important context into account, it then goes a long way towards explaining how the existence of a committed spend discount over a fixed period not only provides a positive benefit of being able to budget spend over that time period, but also means there is essentially no reason to look elsewhere while that contract is in place. One participant – a high strata user of AWS’s services – describes this view as follows:

“We operate within a budget, and we have to make sure that we’re being cost effective. If some of those discounted services disappeared, then we would absolutely think about whether we would use another cloud provider. If they offered us the best discount you would then have to factor into that the cost of migration, the cost of operation, the cost of upskilling because it would be the total cost of completing that piece of work, it wouldn’t just be that’s \$50, this is \$20 or whatever.”

IT, AWS, High

- 6.2.4. In this sense, the existence of the discounts, in conjunction with other disincentives as outlined by this participant, do contribute to a sense that there is little incentive to switch. This view becomes even more pronounced when the view held by the participant is also that there is little difference between the main cloud providers. Therefore, once a choice of provider is made, and the relationship is deep enough (both in terms of the breadth of services used and the length of time with that provider) that contractually based discounts are on offer, the business case for spending time even looking at alternative providers becomes very weak.
- 6.2.5. Related to this point, the wide range of discounts (as outlined in the previous section) and the way in which these discounts are interrelated either within a master services/enterprise agreement, or individually negotiated between user and provider, also tend to disincentivise any attempts to compare offers from providers. A complex, multilayered discounting structure is very hard to compare to another provider's potential offer.

“The commit spend discount was a considerable factor [in provider choice]. And it is also potentially one of the reasons why we don't look elsewhere. Because we do get that deal. And also, there are things wrapped into that that aren't just around commercials. So, for example, we've put all of our architects, testers, engineers, data engineers through comprehensive AWS training for free. That is something we get under that agreement. That is a factor in that decision and isn't something that has been matched elsewhere.”

Transport, AWS, High

- 6.2.6. This point is further illustrated when we look at the opposite end of the spectrum in terms of complex cloud usage – i.e. those users (mainly in the Low strata) who are using their public cloud services in a relatively simple way. Participants who fall into this category stated that the existence of credits early on in their public cloud journey were instrumental in their decision to switch from one provider to another.

“We moved everything from Amazon, where we had been for 2 years, to Google Cloud. Although the migration process was challenging, we moved our entire workload from AWS across to Google because of the significant cost savings. The migration caused headaches, especially when reconfiguring some of our AI services to run on Google's architecture instead of AWS. But with the lure of £300,000 free credits for 2 years, we pushed through the problems and were able to get fully moved over to Google Cloud.”

IT, Google, Low

Impacts on cloud usage behaviour

- 6.2.7. One aspect of behaviour which many of the research participants reported is affected by the contractual discounts in place is the way in which these contracts affect cloud usage within the organisation.
- 6.2.8. Some research participants described how their companies use certain cloud services, not because there is a business or an IT need, but for the sole purpose of meeting committed spend targets.

"If we were close to [not meeting] the minimum spend, there would be a push to try and move as much as you could into the cloud to ensure that you're getting the best value."

Outsourcing, Microsoft, High

"Let's say you commit to \$10 million over five years. You've got to get there by the way, because if you don't hit it, they take money back off you, which is incredibly prohibitive."

IT, AWS, High

- 6.2.9. Other participants spoke of relocating more workloads than they might otherwise need to, simply in order to meet committed spend targets. As the quote below illustrates, providers can then use the current discount structure to influence what they are prepared to offer when the contract comes up for renewal:

"Then when you're looking at renewal because you're then locked into that provider, you are dependent on the functionality they offer and you invest in resources and manpower to support that functionality. So then the discounts become very important because what Amazon tends to do is at the renewal stage, they will link their existing discount structure to an increase in commitment, which means that we are obliged to increase our usage of AWS via more capacity, more data."

Media, AWS, High

- 6.2.10. When talking about the ways in which discounting contracts affect how their business uses the cloud, participants did sometimes point out an inherent contradiction in the relationship with their cloud provider. Specifically, some participants acknowledged that providers such as AWS and Google, are very good at offering to optimise their customers' usage of the cloud so that their spend is not wasteful or unnecessary. However, it was pointed out that even when ways of reducing spend or at least spending more efficiently are identified, it was not always in the customers' interest to act on those recommendations because to do so risked missing the minimum spend level needed in order to trigger a discount. An example of this can be seen below:

"We did a whole load of testing and discovered that there's a form of compute that's 27% cheaper and that performs equally. But the problem was because we put a committed use discount on top of the contractual discount for a whole load of that type of compute, and we're locked into that for three years, whether we use it or not, we're going to pay the bill, so we may as well carry on using it [the more expensive type of compute]. So come this time next year, we'll be shifting if we haven't already completed the shift from the more expensive type of compute to the cheap type of compute."

Healthcare, Google, Medium

6.3. Conclusions

Overall, the research has shown that the existence of discounts is widespread within the public cloud market and that these discounts can be seen as complex in how they are put together within contracts. The research also demonstrates that the existence of discounts does have an impact on switching behaviour, although that impact does not exist in isolation from other factors that would be considered if switching cloud providers was being considered. Rather, the impact of CSDs on switching behaviour is only one element of a larger cumulative set of factors that affect cloud users' willingness to multi-cloud or switch cloud provider.

- 6.3.1. A wide range of discounts are on offer from public cloud providers. These range from high level committed spend discounts (the most common type mentioned) through committed use discounts all the way to specific discounts offered on a service-by-service basis.
- 6.3.2. The research suggests a relationship exists between the size of business/cloud usage strata and the type of financial or contractual relationship companies are prepared to enter into with their cloud provider(s).
- 6.3.3. Larger companies spending heavily with cloud providers (and often with more complex IT infrastructure needs), almost always reported having some sort of discounting arrangement in place. Very often these arrangements were reported as being multi-layered and complex, including CSDs, CUDs or Reserved Instances and other types of discounts.
- 6.3.4. Users thought that this complexity might to some extent be designed to optimise their spend so that the discounting arrangement is based around their actual usage. However, many participants reported that it is difficult to assess whether their discounting arrangement is, in fact, the best way to use services offered by their cloud provider. This is due to the complexity of billing and ongoing contracts sometimes not aligning with each other in terms of timing.
- 6.3.5. This reported complexity would also make direct cost comparisons between providers difficult, if any of the participants wanted to do this. However, very few said they did want to do this as they were not looking to switch, and if they were, it would not be a purely financial decision, as many other risk factors come into play.
- 6.3.6. At the other end of the market, smaller, low strata users of the cloud reported switching providers purely based on the financial incentives offered to them. These come in the form of credits – free usage periods or monetary incentives.
- 6.3.7. Companies who are beyond the start-up phase, but who are still in a fast growth phase of their development are more likely to commit to a CUD or Reserved Instance than a CSD as they felt less able to predict their future usage.
- 6.3.8. Finally, a consequence of the contractual discounts in place is that, for some participants, they are clearly driving behaviour in terms of how the company uses the public cloud services they pay for. Minimum spend targets, or discounts based on specific services/marketplace usage, can and do drive users to potentially use parts of the cloud service that they would not otherwise use.

7. Main findings – Theories of Harm 4: Software licensing practices

7.1. Overview of software licensing practices

Software licensing and the extent to which licensing practices affect public cloud choice is always talked about by participants in the context of organisations being Microsoft users. If an organisation has a historical, strong relationship with Microsoft, it is likely (though by no means certain) that they will also be Azure users. The direct connection between software licensing practices and the decision to go with Azure was generally not recognised by Azure users, despite a clear general desire to keep everything within a Microsoft ecosystem within their organisation. A small number of Google Workspace users also reported how their decision to choose Google Cloud is linked to their use of Google Workspace applications, but there are no licensing agreements in place – this is a business decision based on ease of use and consistency.

- 7.1.1. As part of the research interviews, participants were asked to explain the role software license agreements might have played in their initial choice of public cloud provider. The importance of any software license agreements was explored with all participants, not just Azure users.
- 7.1.2. Among participants using Azure as either their sole public cloud provider, or one of their main providers – especially as a provider of IaaS – the research clearly shows a link between organisations who use a range of Microsoft products and their usage of Microsoft Azure. The research uncovered multiple examples where this link is present, with participants often referring to it as though it was the only sensible business decision available to them, when their choice of public cloud provider was first made. The quote below summarises this point:

“We’re a Microsoft shop. So we use Microsoft productivity tools... the whole 365 suite. We used to have on prem exchange and email servers that we ran ourselves. We moved to 365, and then we’re also moving from on prem Windows servers to Azure 3M Windows servers. It’s as simple as that. The internal support engineer is much more familiar with the Microsoft Stack, so it’s easier to move to an Azure stack.”

Transport, AWS and Microsoft Azure, High

- 7.1.3. As many of the participants in the research sample worked for organisations that had been Microsoft users for many years (often decades), the decision to go with Azure was often described as ‘historical’. This means it predates the interview by at least five years, but often more than this. A practical outcome of this is that many of the research participants were not in their current role when the decision to choose Azure was taken. The perspective of our participants tended to focus on high-level reasons behind the decision to go with Azure and then explaining the organisation’s current thinking in relation to their current and future use of Azure.
- 7.1.4. Reasons behind the historical decision to go with Microsoft Azure tended to focus on the technical aspects of the decision – in particular the skills base within their organisation and the interoperability of Microsoft software products with Microsoft Azure IaaS and PaaS offerings.

“We were largely a Microsoft shop, so that turned us down the Microsoft journey in terms of the early explorations. And then once you build capability, you build confidence. You build throughput on that. It snowballed on its own in a way, because it just made it easier to leverage the capability and the platform that a company like ourselves had invested in. [...] Then you’re tied into the software agreements. You’re then tied into platforms that are built that consume those software agreements. Then there’s a natural and an easier journey to migrate from there into the Azure cloud versus switching over to something else.”

Business Services, Microsoft, High

- 7.1.5. As this quote clearly demonstrates, the decision to go with Azure was, in this participant’s words, part of their ‘Microsoft journey’. In this example, this was a journey that had started over 35 years prior to the interview. In that sense, as the participant said later in the interview, the decision to go with Azure (which was itself around 10 years ago) was inextricably linked to the pre-existing decision to use Microsoft products. The often-predetermined nature of the decision to go with Azure was a rationale the participants gave repeatedly, but not with any negative connotations attached – it was, in their eyes, a simple statement of fact.
- 7.1.6. As well as a decision linked to other historical decisions within an organisation, it is also clear that Azure users see a benefit to their continuing relationship with Microsoft. As the quote below shows, participants were able to identify current practical benefits alongside the ‘natural’ choice to go with Microsoft Azure in the past.

“I think it really was all the synergy. We’re a heavy Microsoft customer and we have a pretty good relationship with Microsoft. We already had an element of skills within the organisation in that space. So for us it seemed to be the natural fit. Some of the services that we wanted to use in the cloud Microsoft make it cheaper to use it in their cloud than they do if you use it in somebody else’s cloud.”

Financial Services, Microsoft Azure and AWS, High

- 7.1.7. Participants also made multiple references to ‘enterprise agreements’ in place between Microsoft and their organisation. These agreements cover all products and services offered by Microsoft to their organisation, of which Azure is only a part – in some cases, a relatively small part. It is clear that organisations which are heavy Microsoft users see these enterprise agreements as being beneficial to them. The quote below illustrates the frustration one participant feels by not being able to benefit from a Microsoft enterprise agreement:

“We used to have an enterprise agreement. We were a slightly larger company at the time, and they’ve raised the limit for the number of seats that you need to 500 employees before you can make an enterprise agreement. Now, we’re not close to being it, which means we can’t purchase licences from Microsoft directly. We have to go through a cloud solution provider, which can be a painful experience.”

IT, Microsoft Azure, Medium

- 7.1.8. The effect of these enterprise agreements upon Azure customers’ view of their Azure service is that it is seen as inextricably linked to all the other Microsoft products and services they use.

- 7.1.9. Some participants reported that it is quite difficult to unpick the precise charges they incur for their Azure use as they are part of a complicated enterprise agreement which covers all products such as the full Office suite (365), Windows Servers, Servicing, Software Updates, Dynamics 365, Power BI etc. This situation is summarised well by the quote below:

"I'm not sure of the exact figure but we do have a significant amount of committed spend with Azure as part of our Microsoft contracts and that will include Azure, It will include SPLA [Service Provider License Agreement], it'll include Office 365 – it will include all of those things."

Business Services, Microsoft Azure, High

- 7.1.10. However, despite this lack of transparency reported by some in being able to accurately assess their specific Azure spend, the overall interoperability and familiarity with the .NET ecosystem reported by heavy Microsoft users means that virtually all the Azure users in the sample feel that, on balance, the decision to go and remain with Azure is beneficial to the organisation. This same interoperability and familiarity also means that the Azure users within the research sample are very unlikely to be contemplating switching away from Azure, as the quote below illustrates:

"Because we were already in an enterprise agreement with Microsoft, we were a Microsoft house, we use Microsoft development tools, so in all honesty we didn't review another cloud provider. We didn't go 'right what are our options today? Which one gives us the best functionality?' Because we already had everything in one bucket so it was a much more time effective option to stay where you are."

Energy, Microsoft Azure and AWS , Medium

7.2. Implications/impact on behaviour

Despite the clear connection between widespread Microsoft usage within an organisation and the decision to choose Azure as their public cloud provider, participants generally do not see a direct causal link between their organisation's decision to choose Microsoft for its software needs (a decision which often predates the public cloud decision by several years, possibly decades) and its decision to choose Azure. There is, however, some recognition that it would be very difficult to unpick their organisation's reliance on Microsoft from their present-day situation.

- 7.2.1. The conversation about software license agreements didn't focus exclusively on the role they played during the initial vendor choice, but the interviews also explored the role they might continue to play in the company's decision to stay with that provider and the extent to which such agreements might contribute to any multi-cloud considerations, if there were any.
- 7.2.2. The main way in which software license agreements impact upon behaviour is through their inclusion in enterprise agreements struck between Microsoft and the customer. The overall benefit of these enterprise agreements is to offer large Microsoft customers discounts on the standard pricing across a range of services, of which Azure is a part. Participants stated that the benefits offered by these enterprise agreements mean that they rarely, if ever, review their Azure usage. It is important to note that these are the terms in which they talk about the impact of the enterprise agreements –that they are a symbol of how deeply embedded their organisation is with Microsoft, and that organisationally, there is no appetite to review that relationship in the near future.

- 7.2.3. As well as enterprise agreements, a small number of participants referred to licensing reasons that exist within cloud providers' ecosystems which also discourage both switching or a potential multi-cloud strategy.

"If you have lots of Microsoft orientated users using a workload on AWS, you have to re-license some of their Microsoft workload again to get them to use it on AWS. Whereas Microsoft allow you to port your licences over, effectively the money you've already paid for the license if you use it on Azure. [...] There were one or two types of workloads where the Microsoft overall licencing model made Azure the most competitive, because obviously you could use your Microsoft licences on it."

IT, AWS and Microsoft Azure, High

"If your economic model is, I want to charge you for what you have in the cloud, and I'm going to benefit whenever you have high performance workloads sitting in the cloud, which could be run much more cheaply by a smaller on-premises infrastructure. What would be my benefit of letting you move between those two things seamlessly?"

Business Services, Microsoft Azure, High

"Software like operating systems and database servers are very expensive, they have a material impact. Because Microsoft owned those licence agreements, they were in control. They were able to heavily discount down. [...] We have some of those older operating systems. You know, they go end of support and Microsoft were offering incentives where they would extend that support. [...] That meant that we would get things like critical security patches for longer if we went to Azure."

Media, AWS, High

- 7.2.4. It follows that while there appears to be very little appetite among Azure users to review their Microsoft (including Azure) reliance, there is some recognition that there are aspects of the technical and financial structures inherent in Azure that make the idea of reviewing Azure usage unattractive from a business cost/benefit analysis point of view.
- 7.2.5. It should be noted that similar comments to those above (which are in relation to Azure), were also given in relation to other cloud providers, especially AWS. However, comments relating to the AWS ecosystem incentivising users to move more and more workloads into the AWS cloud, do not include reference to software licensing or enterprise agreements. In that sense, AWS users are more willing to entertain the idea that they could switch, even if they do not want to, while for Azure users the stakes are too high for that (switching) to be a realistic scenario.
- 7.2.6. It is important to point out that not every Microsoft based organisation in the sample had chosen to use Azure. While it is true that the majority did, there were one or two exceptions and the reasons behind their decision not to choose Azure are relevant to understanding how organisations can take an alternative decision.
- 7.2.7. One participant described a historical decision to go with AWS in 2015, despite describing themselves as '.NET development house'. Crucially, this participant highlights the 'separation between development and operations' when the decision regarding cloud

provider was made. This separation meant that a decision over infrastructure could be taken entirely independently of the organisation's reliance on Microsoft for its development work.

- 7.2.8. This participant then went on to describe how, given their nine-year relationship with AWS, the decision to move away from AWS now would be difficult, even though they remain a '.NET development house'

"Although we use .NET in the years since then, AWS supports .NET much more so now you can write Lambda Functions in .NET as well. Because we're baked into those AWS services it would be a rebuild now. How it would work is that everything's defined through terraform which is an infrastructure as code language... it certainly wouldn't be like a few hours work, it would be quite a heavy investment from a team to redesign everything for Azure."

Charity, AWS, Medium

- 7.2.9. This example illustrates that it is not axiomatic that a Microsoft reliant organisation will automatically choose Azure. However, for many of the participants in our research who worked at large organisations using Microsoft, their IT operations were far more intertwined and therefore the decision to go with Azure was the default choice.
- 7.2.10. It is also worth noting that there were instances within the sample where pre-existing use of Google's productivity suite (Google Workspace in particular) influenced an organisation's decision to choose Google Cloud as its cloud provider. The participants did not make reference to enterprise agreements or software licensing agreements in place with Google when discussing their decision, which demonstrates a wider finding within the research, namely that businesses tend to make IT decisions based on what is easiest for them to implement. The quote below explains this point in more detail:

"Workspaces made sense for us, and that was our first step into that Google area. It makes sense to us to have a cloud-based office environment. And at the time that was more mature because of how our branches operate. We don't want to be constantly having to update the licencing and everything in the branches. Now they can just log on to a Web browser, so it simplifies our IT. We've got the cloud storage associated with Google Drive. So we're getting that collaborative element across different people. With regards to GCP, then it made sense. We've got a relationship with them. We're happy, you know, and that's how that started."

Retail, Google Cloud and AWS, Medium

- 7.2.11. This illustrates how an organisation's choice of software can influence cloud provider choice even when, as in this case, there is no licensing contract in place. The path of least resistance can be an attractive one to choose when businesses are making IT decisions and having to weigh up risk as well as the costs/benefits of different decisions.

7.3. Conclusions

Overall, software licensing practices are not seen, in isolation, as influencing the choice of public cloud provider. This conclusion refers solely to companies using Microsoft as these were the only participants where discussion of this point was relevant. In many of these companies, enterprise agreements are in place, of which software license agreements are a part. Participants say that these enterprise agreements and the relationships they symbolise mean they are very unlikely to consider switching from Azure, but they do not see them as directly inhibiting switching – just that the status quo would make switching an illogical business decision, given the benefits these agreements provide and the effort required to disentangle their infrastructure from the Microsoft ecosystem.

- 7.3.1. It is very clear that original take up of Microsoft Azure is often closely related to an organisation's pre-existing use of Microsoft products and services.
- 7.3.2. The decision to go with Azure as their public cloud provider has often been taken several years ago (over five, sometimes more than ten years ago). The decision to rely on Microsoft's suite of products and services has often been taken decades before the decision to go with Azure. In that sense, participants often struggle to precisely detail why their organisation uses Azure, beyond describing it as a natural choice for both technical and financial reasons.
- 7.3.3. This makes the role of software licensing in the decision to go with Azure difficult to unpick, but participants did not single out licensing as a key factor on its own influencing their decision.
- 7.3.4. Participants were clearer when discussing enterprise agreements with Microsoft and their impact on any potential decision to switch away from Azure or move to a more integrated multi-cloud solution. Here, participants were generally clear that the breadth and depth of their relationship with Microsoft, as evidenced by the existence of an enterprise agreement, means that they struggled to even justify reviewing their Azure usage, so embedded is it within their whole IT approach.
- 7.3.5. The sample does also contain a number of organisations who use Microsoft for all their software and some PaaS needs, but who have made a decision to go with a public cloud provider other than Azure to be their strategic partner. These organisations tend to either be smaller than those with Microsoft enterprise agreements in place, or they have structured their IT in such a way that the IT infrastructure is sufficiently separated from development work that pure interoperability is less of an issue.
- 7.3.6. Even then, companies using Microsoft who have chosen not to go with Azure, made the decision to do so very early on in their public cloud journey. In that sense, they too have now become embedded with the provider they chose originally.
- 7.3.7. The research did not find many organisations who were Microsoft users, who had initially gone with Azure, and who had then switched away to another provider.
- 7.3.8. We did also speak to a small number of Google Cloud Platform users who had chosen to use GCP because of its interoperability with Google's productivity suite. This highlights the attraction of 'ease of use' as a reason for choosing cloud provider, emphasising that matching software to cloud provider can be purely based on this, and not be at all influenced by any software licensing agreements.

8. Main findings – Future outlook

8.1. Artificial Intelligence (AI)

Unsurprisingly, AI features widely in the interviews, especially when participants were talking about future public cloud trends. The expectation is generally that the main public cloud providers will (if they are not already) offer an AI PaaS solution, and that they will compete strongly for this AI PaaS business. Concern does also exist that once an AI solution is chosen, it will be difficult to switch away from, creating another opportunity for cloud providers to potentially lock customers into their cloud offering.

- 8.1.1. AI is mentioned by the majority of participants in the research, either because it is a field their business is directly involved in, or because it is seen as an emerging technology which they are looking at and expecting to make use of in the near future. Regardless of the context, these participants all said that AI is likely to be the most significant change to public cloud services over the next five years or so.
- 8.1.2. AI was also discussed in the context of potentially being, in the future, the most common reason why an organisation who is either operating a single cloud model, or a siloed multi-cloud model, may wish to move to a more integrated multi-cloud model. The reasons for this would be if their sole/main public cloud provider is different to the provider of the AI PaaS service they wish to use. The downside, however, is this would necessitate large amounts of data transfer as well as challenges in addressing potential technical barriers preventing interoperability. In that sense, AI was talked about in a public cloud context in terms which are wider than the expected technological advantages and efficiencies which AI can bring. The quote below is from a participant who was speaking about the theoretical challenges of using an AI solution from one provider (Microsoft in this case) when their main infrastructure is coming from another (in this case AWS):

“...to do that [using OpenAI] on top of AWS documents means us sending the documents to Azure to get them processed to then send the result back. Now that is going to cost us a lot of money in egress fees if we do that. Also, there’s a whole load of Information Governance problems there because now you’ve got to have sensitive data transiting between different platforms.”

Healthcare, AWS and Microsoft Azure, Medium

- 8.1.3. As the quote hints at, this participant does not currently pay egress fees due to operating a totally siloed multi-cloud model (mostly AWS, with a small amount of Azure usage). Adopting Microsoft’s AI solution (based around OpenAI) would mean he would need to address the issue of egress fees as they could be very substantial given all his analytics capabilities are hosted in AWS and would remain so. In addition, he would have to evaluate the perceived risk in large volumes of data transfer between cloud ecosystems – again, something that does not concern him with his current set up.
- 8.1.4. The implications of the adoption of an AI solution outside of the business’s main cloud provider are therefore not as simple as deciding who has the best AI offer or machine learning capabilities. These concerns were voiced by many participants when discussing AI – not in ways that meant they would not consider a provider away from their main cloud provider, but in ways which highlighted that the decision may force them to address factors

within the public cloud market that, in their current set up, they considered trivial or irrelevant to them and their business.

- 8.1.5. It should be noted that when participants began to think about these future challenges that AI might force them to address, they did not always see the solution as one that should be solved by the providers, or a third party, but rather the challenge being one that the business itself would need to solve internally, as the quote below highlights:

"I wouldn't see any of the technical restrictions for cloud providers being a barrier to adopting a new service. I think all the challenges are on the business side. They're not on the cloud side."

Transport, AWS and Microsoft Azure, High

- 8.1.6. Generally, participants described AI as an area that is already being fiercely competed for by the main cloud providers, with a perception that both Microsoft and Google are ahead of others at the moment, but that this is not set in stone. In that sense, AI was seen as an area where competition is driving innovation, and this was considered to be a positive feature of the cloud market, as the quote below illustrates:

"If in the future, let's say if I want to write an application to do the following, I don't actually write the code, I just instruct the thing to do it for me. Then that might mean a change. For example, if Google do that before anybody else or Microsoft does it before AWS and I can create a product with half the engineers or a quarter of the engineers, that would be a kind of tectonic shift I would consider [another provider for]."

Business Services, AWS and Microsoft Azure, Low

- 8.1.7. Alongside this positive view of competition driving innovation in AI there was a recognition that there is no real endpoint to this, with providers constantly updating their large language models (LLMs). This links to a key concern in relation to AI. Participants recognised that although there are differences between providers, and they are all innovating in this area, at some point they (the participants) will need to make a decision as to which AI service they want to commit to. The prevailing view among participants is that from this point onwards, they will largely be locked into the AI service they have chosen, as this quote illustrates:

"One of the things that is a concern currently is lock in. So for our analysis work, we've used AWS, their tooling, their modelling and the lock in in terms of AI feels a lot stronger than it does in other services. For example, the optimisation of certain models on certain clouds would make it very difficult from my understanding to move elsewhere. But it's definitely something that we're looking more into. I don't think we understand what the answer is currently. But it is a concern of ours, and the lock in is a big concern because I think it takes us down a certain way of using AI with certain models."

Transport, AWS, Microsoft Azure and Google Cloud High

- 8.1.8. This concern around AI and machine learning in particular is inextricably linked in participants' minds. Essentially, the value of AI in this context is the AI tool learning based on the unique business data that exists in the cloud. Once this machine learning feedback loop is in place, participants could foresee it being very difficult to justify backing out of, as this quote also highlights:

"I wouldn't say that it would make it easier to migrate, though, because, although some of the AI libraries are very similar, in order to roll out AI as quickly as possible, you almost have to use the machine learning services from AWS to get you there very quickly. That means that you are locked into AWS."

Transport, AWS, Microsoft Azure and Google Cloud High

- 8.1.9. At the most extreme end, one or two participants took the view that AI will be so important, that its development and distribution should not be left in the hands of the cloud providers, but instead should be taken into the control of governments, much like a utility service. The quote below expresses this view

"I would argue they're using the position now that they've got the cloud, to become the monopoly providers [of AI] of the future unless we say to them you have to disentangle your AI, but I don't know whether governments have got the courage anymore to do that."

Healthcare, Google, Medium

- 8.1.10. This view was not commonly held but does illustrate the importance attached to AI by many participants when they are asked to consider future trends.

8.2. Future changes to the cloud market

At the end of the interviews, participants were asked for any suggested areas where they would like to see changes to the public cloud market. It must be stressed, that rarely were these ideas top of mind or strongly expressed – they were in response to direct questions and came at the end of a 1-hour long interview where participants had been given the chance to organise their thoughts in a way that they would not normally do in their day-to-day roles.

- 8.2.1. The most widespread complaint expressed by participants relates to the lack of transparency regarding public cloud charges. While many participants were able to state their total overall spend with their cloud providers – either on an annual or monthly basis – they found it difficult to understand how these charges are arrived at. Examples given in the research describe PDFs of over 100 pages detailing how charges are calculated, with some businesses employing people in their Finance teams dedicated to managing their cloud spend, as the complexity warrants full time members of staff to manage this side of the business.

"The number of different line items on the invoices actually make this even harder, because it's such an overwhelming amount of line items that it's really difficult to even understand what you're being charged for. It's ridiculous that you have to go to a broker to get a tool to understand the billing of Google Cloud – they should be able to make something where it's easy to understand."

E-commerce, Google, Low

- 8.2.2. As outlined earlier in the report, egress fees were not a great concern to participants. However, when asked what changes or improvements they would like to see to the public cloud market, egress fees were mentioned quite extensively. Some favoured an outright ban, others felt they should be reduced, while others said they should be regulated so that they are the same regardless of provider. These views themselves evidence the fact that not all participants are automatically against egress fees – describing them as ‘the cost of doing business’ or ‘like a tax’ – but this view does also link to participant’s desire for fairness in charging data transfer fees, something which regulation around egress fees could provide.
- 8.2.3. The idea of regulation of the public cloud market was not widespread, but is mentioned by a number of participants in different contexts. The research did not pick up a strong sense that public cloud providers are seen as currently abusing their position due to a relatively small number of providers in the market. But some participants did feel that this possibility existed in the market, particularly if consolidation takes place in the future, or one of the main providers decides to leave the market for whatever reason. The quotes below highlight this sentiment, while also showing the desire for this regulation to be light touch:

"You know capital marketplaces are regulated so the infrastructure itself, where people are trading stocks, there is an organisation regulating those infrastructures. I do believe that to some extent there should be some kind of regulation or some guidelines for public cloud providers, because basically they provide technology infrastructure to the real economy so it would be a good thing to have a kind of overseeing body to see what they are doing."

IT, IBM, Low

"And then an ability for us to have a backstop to have support in that negotiation. So, um, something that is being considered in other regulatory areas is a kind of final offer arbitration. So principles where we are unable to reach an equitable outcome with our providers that we have somewhere we can go for support. I think that's something that the CMA should be very seriously considering. This is a regime based on fairness and transparency, and I think that needs to be supported by some kind of regulatory backstop."

Media, AWS, Microsoft Azure and Google Cloud High

"I would recommend the regulator to have a data ombudsman. So, someone who is impartial that has the technical capability and has the regulatory authority to say OK, stop, – this guy needs his data, give it to him."

IT, AWS and Google Cloud, Low

- 8.2.4. This last quote refers to a possible dispute where the cloud provider is refusing access to a client's data. The broader point expresses a desire for an ombudsman for all disputes between providers and clients – not only over data access disputes.
- 8.2.5. Other improvements tend to be linked to participants' desire to retain freedom of choice for as long as possible. One participant referred to Google attempting to link its Google Cloud offer to Google marketing activity. Their view was that these services should be decoupled.

"If you want to have the best of marketing abilities with Google, they will say, 'Well, we can only do that if we have your information to do so'. All your player data we need to track your players and all the information needs to be on a Google cloud for us to understand who your players are such that we can market to new players"

Gaming, AWS and Microsoft Azure, Medium

- 8.2.6. A few participants referred to a desire to see more technical improvements in services that allow and improve the operation of workloads across clouds. Specifically, participants wanted to see improvements in Kubernetes, containerisation, abstraction tools so that the possibility of an integrated multi-cloud model would be less technically daunting than is the case currently. This did not necessarily portray a strong desire for interoperable multi-cloud solutions in the near future, but instead speaks to businesses' desire to retain flexibility and freedom of choice for as long as possible, so that they can respond to unforeseen challenges in the future.

"The API gateway in AWS and the equivalent in Azure are very different. They are very similar in the core, but they are very different in terms of how the APIs are exposed. They are very different across two different platforms simply because there aren't that many open standards and that means that you are locked into one platform."

Transport, AWS and Microsoft Azure, High

"I think that Kubernetes was a game changer because it's platform agnostic, and all the cloud providers agree to contribute to that platform to make it open in a way that you can run it in your own data centre or in other cloud providers. Of course, they have some small features that are specific to each cloud, but it's an extension. But if we use 90% of the product, it will make it easy to migrate between providers. So I think it's more about transforming the cloud to be more open to open standards to share and create these consortiums – this cloud native approach where you can run anywhere."

Mobile gaming, AWS, Medium

9. Appendix – Field materials

9.1. Letter of authority

January 2024

Letter of Authority

The **Competition and Markets Authority (CMA)** is working with Serrula Research and Jigsaw Research – independent opinion research companies – to understand how senior experts within organisations using public cloud infrastructure services, feel about those services and the decisions that led to their current public cloud provider(s) being chosen.

You may be able to help us with our research

Why get involved?

The CMA is an independent government body and the UK's lead competition and consumer authority. As part of our role, we help people, businesses and the UK economy by promoting competitive markets and tackling unfair behaviour. Currently, the CMA is investigating the cloud infrastructure services market: how it functions, how easy it is to switch between providers, whether there are any barriers to switching/working with more than one cloud provider, and whether the market might work better. By getting involved in our research, you will make a key contribution to the evidence we assess as part of this investigation, ensuring that the views and experiences of people like you are fully and accurately reflected in our conclusions about how well the cloud services market is functioning for businesses like yours.

How did we get your details?

The CMA has statutory powers to request customer details from providers in a market which is under investigation. This is what it has done in this case – the CMA requested contact details from businesses who are buying public cloud infrastructure services from the main providers in the market. From that list, businesses are being selected at random to ask if they are willing to take part in the research.

What happens next?

Serrula Research interviewers will be contacting you initially. They will ask you a few questions to ensure eligibility for the research, as it is important the research includes participants from a broad range of different types of business. Assuming your business qualifies and you are able to participate, we would be very grateful if you would take the time to agree to be interviewed.

Participation in the research is voluntary, and everyone who takes part in a follow-up interview with a Jigsaw researcher will receive an incentive as a thank you for their time. Your personal data and interview answers will be used and stored (processed) fully in line with the UK GDPR (please see below for more details).

If you would like to confirm the credentials of Serrula Research and Jigsaw Research, please call the Market Research Society's verification service free on 0800 975 9596. If you have any queries about the research itself, or would rather not be approached about taking part in the research, please contact X.

We very much hope that you will want to help us with this important study.

Yours sincerely,

Senior Market Research Adviser , The Competition and Markets Authority

More about the CMA, your personal data, and data protection

As part of its work, the CMA often asks customers in a particular market to take part in research. To inform the evidence-gathering for its Cloud services market investigation, the CMA has asked Serrula Research and Jigsaw Research [the research companies] to contact and interview a sample of employees of customers who buy public cloud infrastructure services.

The CMA understands that your business buys cloud services from one of the following providers: Amazon (AWS), Microsoft, Google, Oracle or IBM.

Under s174 of the Enterprise Act 2002, the CMA has statutory (legal) powers to obtain customer lists from parties to a CMA investigation, and to share these lists with a research supplier for use as a sample. On this occasion, we have used our powers to require each of the named cloud providers to supply business names and at least one contact email per business from their customer lists. .

Personal data is data that allows an individual to be identified (either directly or indirectly). Your company name and email may be disclosive of an individual and therefore the CMA, and the research companies (as the CMA's agents), may be processing your personal data. The CMA's legal basis for processing customers' personal data is that the processing is necessary for the performance of a task carried out in the public interest (public task) or in the exercise of official authority vested in the controller (GDPR Article 6(1)(e)). Because the CMA has statutory powers to process your personal data, the CMA does not need your consent to process your personal data under data protection law.

You can read more about how the CMA, as a Data Controller, processes personal data (including the contact details for the CMA's Data Protection Officer, information about your rights in relation to your personal data, and details of how to exercise those rights, including how to complain to the Information Commissioner's Officer) in the CMA's Privacy Notice. You can find a link to this document (titled 'Personal information charter') on the CMA's homepage at:

www.gov.uk/government/organisations/competition-and-markets-authority. You can read more about Jigsaw Research's privacy policy here: <https://www.jigsaw-research.co.uk/privacy-policy/> and about Serrula Research's privacy policy here: <https://www.serrula-research.com/privacy-policy>.

While the CMA's market investigation is ongoing, the CMA, and the research companies, will process your personal information securely at all times. Your personal data will be processed for the purposes of this research project and market investigation only. All your personal information will be securely deleted on conclusion of the CMA's Cloud Market Investigation.

If you agree to be interviewed by the research companies, you agree to the processing of any additional personal data you share with them during your recruitment to the research and/or during your interview. This personal data will not be shared with the CMA, or any third party, in a way that would allow you to be individually identified, except in the event that we are required to do so by law. Please note that by agreeing to be interviewed, you will also be agreeing to the CMA processing aggregated data from which you cannot be individually identified, in the form of Jigsaw Research's report on the findings from the research for the CMA. Again, this personal data will be processed for the purposes of this research project and market investigation only and securely deleted on conclusion of the market study.

9.2. Recruitment screener

Introduction

Good morning/afternoon/evening. My name is _____, and I work for Serrula Research. We have been commissioned by Jigsaw Research to organise some research on behalf of their client, the Competition and Markets Authority (the CMA). The CMA is the UK's principal competition and consumer protection body responsible for promoting competitive markets and tackling unfair behaviour.

We are recruiting IT decision-makers to take part in an important research project which aims to understand customers decision-making on their choice of public cloud infrastructure services. This research will inform the CMA's cloud services market investigation.

The study is not a solicitation for business now or later. Neither you nor your company will be identifiable in the findings of the research shared with the CMA.

I'd like to ask you some questions to determine if you qualify for this particular study. The answers that you give to me today will be shared with **Serrula Research** and **Jigsaw Research**.

The research activity and this interview will be conducted in accordance with the Market Research Society Code of Conduct, and the information you provide will be treated in accordance with data protection law. You have the right to withdraw your consent to our processing the information you provide or object to our processing of your information.

This interview is just to establish eligibility for the research project and will take around 5 minutes.

If you are eligible, you may be invited to attend an online interview which will last approximately 1 hour. You will be provided with a list of the topics to be covered in the interview ahead of time. Most of the questions will relate to your experience as a decision-maker regarding your organisation's purchase and use of cloud services, but some may draw on your wider expertise and knowledge.

You will receive a £X incentive for participating in the interview. This payment will be paid via BACs by Serrula Research if you qualify and attend the interview.

Q1. Are you interested in proceeding further?

Yes	CONTINUE TO Q1a
No	THANK and CLOSE

Q2. What business sector do you work in?

DO NOT READ OUT

Business and professional services	1	CODES 1 TO 7 – RECRUIT MIN 6 PER BUSINESS SECTOR
IT	2	
Retail, consumer, hospitality	3	
Telecoms, media and entertainment	4	
Banking, insurance and financial services	5	
Manufacturing, utilities and primary industry	6	
Healthcare and life sciences	7	
Public sector	8	CLOSE
Other (write in)	9	OTHER – AIM TO RECRUIT MINIMUM 2

RECRUIT MINIMUM 6 PER BUSINESS SECTOR (CODE 1-7)

RECRUIT MINIMUM 2 OTHER (CODE 9)

Q3. Which, if any, of these public cloud computing services does your company currently use? By public cloud we mean a cloud deployment model where cloud services are open to all customers willing to pay and computing resources are shared between them.

SaaS – Software as a service (i.e. software which is licensed on a subscription basis and is centrally hosted by a cloud provider. Sometimes referred to as ‘on-demand’ or ‘web-based’ software e.g. Salesforce.com or Microsoft 365)	1	CLOSE IF <u>ONLY</u> USE SaaS MUST USE IaaS/PaaS TO CONTINUE
IaaS – Infrastructure as a service (i.e. delivers computing, network, and storage resources to customers on-demand, over the internet (or over a private connection) and on a pay-as-you-go basis. E.g. Amazon Web Services (AWS) EC2, Microsoft Azure Virtual Machines, Google Compute Engine (GCE)	2	CONTINUE IF IN IT SECTOR GO TO Q3b, IF IN ANY OTHER SECTOR GO TO Q4
PaaS – Platform as a service (i.e. A platform for software creation delivered via the web allowing developers to design and create customised applications without having to worry about operating systems, software updates, storage, or infrastructure. e.g. AWS Elastic Beanstalk, Microsoft Azure App Service, Google App Engine.)	3	CONTINUE IF IN IT SECTOR GO TO Q3b, IF IN ANY OTHER SECTOR GO TO Q4
None of the above – we only use on-premises/private cloud services	4	CLOSE IF <u>ONLY</u> ON-PREMISES/PRIVATE CLOUD MUST USE IaaS/PaaS TO CONTINUE
Don't know	5	SEEK REFERRAL TO RIGHT PERSON TO ANSWER THIS, THEN CLOSE

MUST USE CODE 2 OR 3 – IaaS OR PaaS – TO CONTINUE

ADDITIONAL QUESTION FOR THOSE IN IT SECTOR (Q2 = 2)

Q3b. Does your company sell PaaS and/or SaaS services to customers where you do not own any of the underlying raw computing resources?

Yes	CONTINUE TO Q3c
No	GO TO Q4

Q3c. What proportion of your spend on IaaS/PaaS is to support your customers who are buying PaaS/SaaS services from you rather than for your own internal business use?

50% or more of our spend on public cloud infrastructure is for supporting customers rather than internal use.	1	THANK AND CLOSE
Less than 50% of spend on public cloud infrastructure for supporting customers rather than internal use.	2	CONTINUE

THOSE IN IT SECTOR MUST NOT MAINLY USE PUBLIC CLOUD FOR CUSTOMERS RATHER THAN THEIR OWN INTERNAL USE (IE. WE DO NOT WANT TO INTERVIEW IT BUSINESSES WHERE THEIR MAIN USE IS AS INDEPENDENT SOFTWARE VENDORS/MANAGED SERVICE PROVIDERS RATHER THAN END USERS)

Q4. How long has your company used public IaaS/PaaS service(s)?

Less than 5 years	1	CONTINUE
5+ years	2	CONTINUE
Our organisation has always used cloud	3	CONTINUE

AIM FOR A SPREAD OF CODES 1,2 AND 3

Q5. Which of the following best describes your responsibilities with regards to decision making about public IaaS/PaaS service(s) in your organisation?

I am solely/jointly responsible for decision-making about IaaS/PaaS (e.g. suppliers to use, purchase decisions etc) <u>for the whole company.</u>	1	CONTINUE – CHECK POTENTIAL FOR REFERRAL
I influence the decision-making process for IaaS/PaaS (e.g. suppliers to use, purchase decisions etc) <u>for the whole company</u> – but do not have the final say.	2	
I am solely/jointly responsible for decision-making about IaaS/PaaS (e.g. suppliers to use, purchase decisions etc) <u>for my division or department only.</u>	3	CONTINUE – CHECK POTENTIAL FOR REFERRAL
I influence the decision-making process for IaaS/PaaS (e.g. suppliers to use, purchase decisions etc) <u>for my division/department</u> – but do not have the final say.	4	
I have the final sole/joint say about all IaaS/PaaS for my company but do <u>not</u> take an active role in the decision-making process.	5	ASK FOR REFERRAL TO DECISION-MAKER(S) THEN CLOSE
I have minimal involvement in decision-making about IaaS/PaaS.	6	
I have no involvement in decision-making about IaaS/PaaS.	7	

ALL TO CODE 1 OR 2 ONLY – GO TO Q7

IF CODE 1-2, CONTINUE PLUS SEEK POTENTIAL REFERRAL FOR DIVISIONAL/DEPARTMENT/MORE TACTICAL DECISION MAKERS

IF CODE 3-4, CONTINUE PLUS SEEK POTENTIAL REFERRAL FOR WHOLE COMPANY/MORE STRATEGIC DECISION MAKERS

IF CODE 5-7, SEEK REFERRAL TO THE DECISION MAKER(S) THEN CLOSE

Q6a. Please could you tell me what your job title is?

Write In _____

Q6b. Which of the following best describes your role?

Senior management e.g. MD, CEO, Board Director	1	CHECK WITH JIGSAW IF ANY UNCERTAINTY ABOUT DECISION MAKING ROLE, THEN GO TO Q7a
Commercial/Sales/Marketing	2	
Financial	3	
Operations/logistics	4	
Purchasing/procurement	5	
Production/manufacturing	6	
IT	7	GO TO Q7a
Other (WRITE IN)	8	CHECK WITH JIGSAW IF ANY UNCERTAINTY ABOUT DECISION MAKING ROLE, THEN GO TO Q7a

Q7a. Which of these best describes your organisation?

Headquarters in the UK	1	CONTINUE (ASK Q8)
Offices in the UK, but headquarters are in another country	2	CONTINUE (ASK Q8)
Don't know	3	CLOSE

Q8. Do you have a degree of autonomy in terms of buying IaaS/PaaS for your organisation's use in the UK?

Yes	1	CONTINUE
No – please explain any limitations/restrictions	2	CHECK WITH JIGSAW
Don't know		CLOSE

Q9. What is the name of the provider of IaaS/PaaS that your organisation is currently using?

Amazon Web Services (AWS)	1	CONTINUE (SEE QUOTAS)
Microsoft Azure	2	CONTINUE (SEE QUOTAS)
Google Cloud	3	CONTINUE (SEE QUOTAS)
IBM Cloud	4	CONTINUE (SEE QUOTAS)
Oracle Cloud	5	CONTINUE (SEE QUOTAS)
Other (please write in)	6	CLOSE IF <u>ONLY</u> USE OTHER PROVIDER

RECRUIT AT LEAST 10 WHO USE AWS (CODE 1)

RECRUIT AT LEAST 10 WHO USE MICROSOFT (CODE 2)

RECRUIT AT LEAST 10 WHO USE GOOGLE (CODE 3)

RECRUIT AT LEAST 7 WHO USE IBM (CODE 4)

RECRUIT AT LEAST 7 WHO USE ORACLE (CODE 5)

RECRUIT MIN 2 LISTED IN SAMPLE AS HIGH SPEND PER PROVIDER

RECRUIT MIN 2 LISTED IN SAMPLE AS MEDIUM SPEND PER PROVIDER

RECRUIT MIN 2 LISTED IN SAMPLE AS LOW SPEND PER PROVIDER

Q11. What type of contract does your organisation have with your current provider(s) of IaaS/PaaS?

Contract	1	CONTINUE (ASK Q12)
PAYG (charges based on usage)	2	CONTINUE (ASK Q13)

Q12. Please tell us which of the following apply to the contract(s) from your IaaS/PaaS provider(s) in 2022 or since then?

No discount – Are paying the price as it was quoted to us (by provider directly, intermediary, on a marketplace etc) without any discount	1	CONTINUE
Have negotiated a general discount on all or some services we buy	2	CONTINUE
Are receiving a discount because we are buying a range of cloud services from the same provider	3	CONTINUE
Are receiving a discount because we are also buying some non-cloud services from the same provider	4	CONTINUE
Are receiving a discount because we have committed to an agreed minimum level of spend over the length of contract – (CSD)	5	CONTINUE

Q13. Please tell us how much your organisation spent on Egress fees in the last twelve months?

Less than £50,000	1	RECRUIT A MIX OF EGRESS FEE SPEND LEVELS
£50,000 – £100,000	2	
£100,000 – £250,000	3	
£250,000 – £500,000	4	
Over £500,000	5	
Don't know	6	

IF THE RESPONDENT QUALIFIES, EXPLAIN THE REQUIREMENTS FOR PARTICIPATING:

S1. We would like to interview you online via Zoom. The interview will last up to an hour. We will send a list of the topics to be covered in advance of the interview. You will receive an incentive of £X as a thank you for your participation. Are you willing to take part?

Yes	1	CONTINUE
No	2	THANK and CLOSE

CONSENT: Thank you for agreeing to take part in our market research interview. In order to help protect your privacy, I will read out some terms and conditions that we will ask you to agree to and sign when you attend the interview.

S2. This study is being undertaken for market research purposes by Jigsaw Research who adheres to the MRS code of conduct. The responses you give during the interview/discussion will be reported anonymously and in aggregate with responses from other research participants. The recordings may also be sent by secure transfer for transcription. Jigsaw Research will not disclose audio files/transcripts/interview notes gathered as part of a research exercise to CMA. Only anonymised quotes are used in our presentation and report to CMA. The information obtained about you during this project will be used solely for the purpose of this project.

Are you happy to proceed on this basis?

Yes	1	CONTINUE
No	2	THANK and CLOSE

S3. As part of our research process, and only with your express consent on the day, a member of the CMA team may wish to observe our interview with you, but we would take efforts to protect your confidentiality in this instance.

Are you happy to proceed on this basis?

Yes	1	CONTINUE
No	2	CONTINUE

DO NOT EXCLUDE IF UNWILLING TO BE OBSERVED – RECRUIT IF ELIGIBLE AND FLAG AS NO OBSERVATION OF INTERVIEW.

S4. Personal data is data that allows a living individual to be identified, either directly or indirectly. As part of this research, Jigsaw would like to make a digital recording of our interview with you for analysis purposes. We will securely delete the recording from our systems on the final determination of the investigation or any legal challenge to it.

Are you happy to proceed on this basis?

Yes	1	CONTINUE
No	2	THANK and CLOSE

S5. Please note that by giving your consent to participate in an interview, you consent to Serrula Research Ltd and Jigsaw Research Ltd using and storing (processing) the personal data we have collected from you during your recruitment to the research and any further personal data you share with us during your interview. Serrula Research Ltd and Jigsaw Research Ltd will use and store (process) your personal data for the purposes of this research project only.

Your personal data will not be shared with our client, the Competition and Markets Authority, or with any third party, in a way that would allow you or your company to be individually identified, except in the event that we are required to do so by law.

Please note that you are also consenting to the CMA processing aggregated data from which you cannot be individually identified, in the form of our report for the CMA on the findings from the research.

You have the right to access and rectify any data held on you by Serrula Research or Jigsaw Research and to withdraw consent at any time or to object/restrict any processing of your data.

We will securely delete all your personal data from our systems on the final determination of the investigation or any legal challenge to it.

If you have any questions about how your personal information will be used, you can contact Jigsaw Research Ltd on 020 7291 0810 or by emailing datasecurity@jigsaw-research.co.uk

If you have any concerns about Jigsaw's information rights practices, you can report it to the ICO via their website by calling 0303 123 1113.

Are you happy to proceed on this basis?

Yes	1	CONTINUE
No	2	THANK and CLOSE

S6. If we needed to contact you within 6 months of this study for any follow-up questions relating to this study specifically, is this OK?

Yes	1
No	2

RESPONDENT DETAILS

Name	
Address	
Phone	
Email	
Date of depth attending	

RECONFIRM DAY/DATE/TIME OF THE DEPTH WITH RESPONDENT

INTERVIEWER DETAILS

Name	
-------------	--

INTERVIEWER DECLARATION

I declare that this interview has been conducted strictly in accordance with your specifications within the MRS Code of Conduct and with a person totally unknown to me.

Signed (**INTERVIEWER**) : _____ Date _____

9.3. Discussion guide

Introduction (5 minutes)

- Thank respondent for participation in the research, which is being conducted on behalf of CMA and re-iterate value of their experiences and perspective to this investigation.
- Provide background and set the context for the rest of the discussion.
- Confirm position on anonymity and recording/transcription of interviewing.
- For pilot interviewing: confirm position on consent for observation by CMA staff.
- State that the interview will last around an hour

Explain: The intention is to learn about your company and thoughts, experiences and plans on the use of public cloud infrastructure services now and in the future. We are interested in your personal views in your role as an IT professional making decisions about public cloud infrastructure services for your organisation. We are also interested in your wider expertise and knowledge.

Before we start, just a couple of points to clarify:

- A. When we talk about ‘public cloud infrastructure services’ in the interview we are focussing on Infrastructure as a Service and Platform as a Service (IaaS and PaaS services) but not Software as a Service (SaaS) *SHOW EXPLAINER SLIDE*
- B. We are also focussing on public cloud services (i.e. a cloud deployment model where cloud services are open to all customers willing to pay and computing resources are shared between them), that means we’re less interested in issues around private cloud services.
 - Firstly, I’d like to understand a little about your company and your role. Tell me about your company. Can you briefly introduce and describe your business?
 - e.g. number of full-time staff, types of products/services sold and customers they have, how long they have been trading
 - Can you tell us a bit about your own role and the IT team in your business?
 - cover size and structure of the IT team, their role technical/commercial/strategic, how long they have been working in the sector etc...
 - Can you also give me a brief overview of how decisions are typically made about use of public cloud in your business?
 - Do you have a central public cloud strategy?
 - How do you decide whether to use public cloud infrastructure services vs private cloud/on-premises?
 - How are potential cloud infrastructure service options identified and evaluated?
 - How formal is this process e.g. use of formal tender/evaluation/procurement processes?
 - Does the process vary at all by difference use cases/workloads?

Journey towards public cloud infrastructure (10 mins)

Explore their journey of adopting public cloud infrastructure services, which public clouds they use and how they use them.

- Can you give me an overview of all the various public cloud services that your company uses for your infrastructure and platform needs and use cases?
 - Which cloud providers are you using?
 - For infrastructure (IaaS) and platforms (PaaS)
 - (MODERATOR CHECK AGAINST SAMPLE INFO)
 - Do you buy public infrastructure services directly from your cloud provider? If not, do you use any advisers or intermediaries? How does this work?
 - IF USES MULTIPLE CLOUD PROVIDERS: *Thanks, I will come back with more detailed questions on multi-cloud later in this interview.*
 - What are the main use cases on each of the cloud services that you use?
 - What categories of public cloud infrastructure services are you using them for? *SHOW CLOUD SERVICES EXPLAINER SLIDE AS NEEDED*
 - PROBE AROUND IAAS CATEGORIES: e.g. compute, storage, networking
 - PROBE AROUND PAAS CATEGORIES: e.g. databases, data analytics, containers, functions, ML, developer tools, hybrid and multi-cloud
- How much approximately does your company spend on IaaS/PaaS public cloud services annually?
 - And as an estimate, what proportion of your total IT spend is now spent on these public clouds?
 - If this is hard to estimate: What proportion of your workloads is now on these public clouds?
 - And how much data do you store in the public cloud?
- Where are you in your journey of adopting and using public cloud infrastructure services?
 - When did you start using public cloud infrastructure for your use cases/workloads?
 - Have you migrated any use cases/workloads in your company to public cloud from traditional/on-premises IT and/or private cloud?
 - If yes: why? what prompted the move to public cloud?
 - Are there any use cases/workloads in your company that are cloud-native/cloud-first? (e.g. Workloads that are designed and built to maximise the benefits of cloud, as opposed to workloads directly migrated with minimal redesign from on-premises)
 - If yes: why?
 - Have you moved any use cases from the public cloud back to on-premises/private cloud?
 - If yes: why?

Choice of public clouds, deployment models and multi- vs single cloud (25 mins)

Discover how they chose their public clouds, the key purchase criteria involved and explore any potential barriers to multi-cloud models.

- What are the key reasons for using [CLOUD PROVIDERS CITED ABOVE] for your public cloud infrastructure services?
- Which factors were the most important ones when the decision was made to use [INSERT PROVIDERS]?

PROBE ON MOST RELEVANT IF NOT MENTIONED:

- Reputation of the service/service provider
- Service quality e.g. reliability, resilience, security, support from provider
- Price including discounts/credits
- Range of services
- Ease of integration with existing IT
- Ease of integration with other public clouds
- Ability to use existing software licenses
- Cloud-specific skills of employees
- AI capabilities
- Number/location of data centres and the extent of geographic coverage
- Existing relationships with cloud provider/ISV
- Data sovereignty requirements

MAKE SURE TO PROBE ON ALL ACROSS INTERVIEWS

- Did you consider any other public clouds when purchasing these IaaS or PaaS services?
 - If no: Why didn't you consider other providers?
 - If yes: Which other public clouds did you consider? Why those? What was ultimately the reason why you decided against this provider?

MODERATOR MAKE SURE TO GET VIEWS ON LARGE PROVIDERS (AWS, MICROSOFT, GOOGLE) AND SMALL PROVIDERS (IBM, ORACLE, OVHCLOUD) ACROSS ALL INTERVIEWS

I'd like to spend a bit more time on the costs associated with the public cloud services you use...

- What type of billing arrangement/type of contract do you have with this/these public cloud providers?
 - For example, do you have a pay as you go (PAYG) arrangement or a contract? PROBE FOR ALL IaaS/PAAS PROVIDERS CITED ABOVE
 - Were there any negotiations at all on the price?

- If yes: Tell me about the price negotiation process
- Did you negotiate a discount with your public cloud provider/s based on how much you commit to spend with them?
 - EXPLAIN IF NEEDED: Some providers offer so-called Committed Spend Discounts (CSD), i.e. whereby the customer gets a discount if they commit to spend a minimum amount with one particular provider
- How important was the Committed Spend Discount in your decision to use this particular provider?
- How important do you think the Committed Spend Discount will be in your choice of where to allocate/which provider to use, for your next use case for the public cloud?
 - PROMPT IF NEEDED: For example, will you be more likely to choose your current cloud provider for your next use case as a direct result of this discount?
- Were any other commitments to your public provider(s) a factor in any discount you received?
 - PROMPT IF NEEDED: For example, commitments to a certain type or length of contract, or the proportion of your total spend on public cloud given to that provider?
- How relevant were 'egress fees' incurred by your public cloud infrastructure in your choice of cloud provider/Independent Software Vendor (ISV)? (*PROMPT IF NEEDED: By egress fees, we mean the fees that occur when data is transferred out of a cloud system*) **SHOW EXPLANATORY SLIDE IF NEEDED**
 - If relevant: why? PROBE around one-off switching costs vs ongoing data fees
 - If not relevant: why not? PROBE around workaround e.g. designing for minimum egress fees
- And how relevant were **software license agreements** in your choice of cloud provider/ISV?
 - PROBE: how much of your IT is affected by these agreements
 - how important they are to your public cloud infrastructure
 - what are the potential costs or additional work involved of using a different cloud provider when considering these agreements?

Single-cloud customers only

- From your answers so far, I understand that you use a single-cloud approach for your business, is this correct? (INTERVIEWER NOTE: This means IaaS/PaaS from one provider. They may still use some on-premises/private cloud in conjunction with one public cloud provider, or use other providers for SaaS)
- Why have you decided to use a single public cloud provider for your business?
- What do you think more broadly about the multi-cloud model?
 - What do you see as advantages and what do you see as disadvantages of a multi-cloud approach?
 - Does this vary for different types of multi-cloud?
 - **SHOW EXPLANATORY SLIDE AS NEEDED ON HYBRID/MULTI-CLOUD**
 - PROMPT e.g. hybrid multi-cloud where you use a mix of public clouds, private clouds or on-premises?

- Or multi-cloud from more than one public cloud provider, where you may have different degrees of integration? (see technical note for more information)
- Have you considered using a multi-cloud approach for your public cloud infrastructure?
- What made you decide against using multiple cloud providers?
- What role, if any, did **technical barriers** play in your decision against multiple cloud providers? For example...

PROBE ALL:

- Difficult integration of ancillary cloud services across public clouds, such as observability, orchestration, (Identity and Access Management) IAM and billing
 - PROBE: Would you like to be able to integrate ancillary services more easily across different public cloud providers? Please provide examples of how you would like to integrate if it was easier to do so
 - IF TERRAFORM (An open-source Orchestration solution) MENTIONED: How and why do you use Terraform?
 - How much resource (time/effort) has it cost the organisation to use Terraform? What is the benefit in doing this?
- Difficult integration of data stored on one cloud with PaaS services like analytics on another cloud
 - What makes the integration difficult? Is it difference in functionality, different protocols or proprietary APIs
- Uncertainty around potential complications of using multiple cloud providers
- Lack of technical skills required to work with different clouds
- Effort required to integrate and operationalise cloud services from different providers
- Data latency
- IF ABSTRACTION MENTIONED: Please explain how, if at all, this has reduced some of the technical barriers associated with running multiple cloud providers?
 - What costs (time/effort/financial) were incurred by the business in undertaking this abstraction process?

MODERATOR ASK TO EXPLAIN:

- If and how each these pose a challenge for multi-cloud
- Which are most difficult to overcome? Why?
- How might they be overcome PROMPT IF NEEDED; e.g. duplicating effort, acquiring skills etc
- Earlier I asked about **egress fees**... How important were egress fees in your decision in favour of a single cloud provider (as opposed to multi-cloud)? Please explain in as much detail as you can the role egress fees played, if any, in your considerations.
- Similarly, I touched briefly on **software license agreements** earlier in the context of choosing a cloud provider... did your decision in favour of a single IaaS/PaaS cloud approach have anything to do with the software license agreements your business has with the same provider? Please explain

in as much detail as you can the role software license agreements played (if any) in your considerations.

- IF PLAYED A ROLE, PROBE: around ability to keep using these licenses as you migrated into the cloud?
- any disadvantages/restrictions on continued use of these software licenses with other cloud providers
- OR any advantages/cost incentives to obtain new licences of the software with the same cloud provider (see technical note for more information)
- And to what extent did any **committed spend discounts** (CSD) you were offered by any of the providers matter to your decision to go with one single cloud provider as opposed to multiple cloud providers?
- Of all these technical and financial factors, which were the most important, if any, in your decision to use a single cloud approach?
- How likely are you to consider using a different public cloud provider for new use cases in your business?
 - If likely: What role do the barriers we just discussed play for new use cases?

Multi-cloud customers only

- From your answers so far, I understand that you use a multi-cloud approach for your business, is this correct?
- Why have you decided to use more than one public cloud provider?

PROBE:

- Unintentional (e.g. due to a merger)
- Making use of best service offers across different providers
- Value for money
- Resiliency/Disaster Recovery
- Specific needs that are better met by different providers
- To encourage competition between providers for our business

How integrated are the cloud infrastructure services you are using across different public cloud providers? By integrated we mean operations or functions of the public clouds are interlinked or dependent on each other. **SHOW EXPLANATORY SLIDE AS NEEDED** on duplicate fail-safe backup/siloed and integrated multi-cloud:

IF INTEGRATED:/MORE INTEGRATED

- Why have you chosen this approach?
- What made it easier and what made it more difficult to integrate your different cloud functions?

PROBE ON BARRIERS:

- Difficult integration of ancillary cloud services across public clouds, such as observability, orchestration, IAM and billing

- PROBE: Would you like to be able to integrate ancillary services more easily across different public cloud providers? Please provide examples of how you would like to integrate if it was easier to do so
- IF TERRAFORM (An open-source Orchestration solution) MENTIONED: How and why do you use Terraform?
- How much resource (time/effort) has it cost the organisation to use Terraform? What is the benefit in doing this?
- Difficult integration of data stored on one cloud with PaaS services like analytics on another cloud
- Uncertainty around potential complications of using multiple cloud providers
- Lack of technical skills required to work with different clouds
- Effort required to integrate and operationalise cloud services from different providers
- Data latency
- IF ABSTRACTION MENTIONED: Please explain how, if at all, this has reduced some of the technical barriers associated with running multiple cloud providers?
 - What costs (time/effort/financial) were incurred by the business in undertaking this abstraction process?

PROBE ON ENABLERS:

- Support provided by my public cloud provider/s
- Actions our business has taken e.g. on staff training, buying-in technical support services/investment in our IT function
- How important are these factors in your decision-making? For example, would you still have opted for a multi-cloud approach if you hadn't been able to integrate them in the way you did?

IF NOT INTEGRATED/LESS INTEGRATED:

- Is this a deliberate choice or the result of certain barriers that stop you from integrating the different cloud services?
- What barriers are most relevant in this context?

PROBE ON BARRIERS:

- Difficult integration of ancillary cloud services across public clouds, such as observability, orchestration, IAM and billing
- Difficult integration of data stored on one cloud with PaaS services like analytics on another cloud
- Uncertainty around potential complications of using multiple cloud providers
- Lack of technical skills required to work with different clouds
- Effort required to integrate and operationalise cloud services from different providers
- Data latency

- How far, if at all, do you think there are any barriers or challenges to you using more than one public cloud provider related to the level or way in which **egress fees** are charged by your provider? How/why are these a barrier/challenge?
- How far, if at all, do you think there are any barriers or challenges to you using more than one public cloud provider, related to any **software licence agreements** your business has entered? What are they/How/Why are they are a barrier/challenge?
- And to what extent did any **committed spend discounts** (CSD) you were offered by any of the providers influence your choice of where to allocate/which provider to use? or might influence your next use case for the public cloud? Probe incentives to use existing providers and spend levels across them

Customer Experience (5 mins)

Cover satisfaction with current provider/s, consideration, and experience of switching between providers.

- Overall, how satisfied are you with your current public cloud provider(s)?
 - Why? What are their strengths and weaknesses?
 - PROBE: quality, price/price inflation, range of services, innovation, support/service
 - How did you experience the onboarding process when you first switched/migrated to this cloud provider? PROBE: What support, if any, was offered by your cloud provider(s)?
 - Where do you think they could improve their services?
- IF MULTI-CLOUD CUSTOMER:
 - Are there any differences between your cloud providers in terms of how satisfied you are with their services? Why?
- Have you renegotiated your contract with your current public cloud provider(s)?
 - If so: How did you find this experience?
 - What was the outcome of the renegotiation? PROBE: spend, discounts, nature of contract
 - PROBE: Was there any pressure to increase your \$ spend commitment to maintain your % discount on renegotiation?
- How easy or difficult have you found it to predict/control your level of spend on your public cloud/s? Why? How far is your spend in line with your expectations?
 - Probe: any use of provider/other monitoring tools to help track usage and spend

Switching Cloud Provider (10 mins)

I now want to ask about your thoughts and experiences of switching between cloud providers/between public clouds. This could be switching a use case (or part of a use case) from one public cloud to another (e.g. migrating your application from AWS to GCP), and/or switching from using a given cloud infrastructure service hosted on one public cloud to another cloud infrastructure service with similar functionalities hosted on another public cloud (e.g. switching from using AWS RedShift on AWS to Google BigQuery on GCP). This would not include moving from a non-public cloud IT environment onto the public cloud or vice versa.

- How frequently do you review your public cloud provider(s) and the contract/arrangements you have with them?
 - PROBE on rationale/triggers to review
 - Formality e.g. of review process, procurement process and use of tenders etc.
- Have you ever considered switching between cloud providers/between public clouds?
 - Why? Why not?
 - PROBE FOR BARRIERS TO SWITCHING:
 - Technical barriers
 - Egress fees
 - Software license agreements that would be impacted by switching to another cloud provider
 - Structure of contract/discounts offered
 - Any other reasons not to switch provider?

MODERATOR NOTE: DETAILS OF THESE ARE INCLUDED IN EARLIER SECTIONS OF THIS GUIDE

- Have you ever switched between cloud providers/between public clouds?
 - Why/why not?
 - SEE PROBES FOR BARRIERS ABOVE AND
 - Which of these factors are most important when considering whether to switch public cloud provider?

IF THEY SWITCHED:

- What use cases, cloud infrastructure services and public clouds did you switch?
 - Mostly IaaS or use cases using a lot of PaaS services
 - What public cloud providers were involved in the switch?
 - How satisfied were you with the experience of switching? Why?
 - To what extent did you encounter any challenges or barriers to switching? (reference barriers previously raised by respondents in the multi-cloud section if relevant)
 - If so, how did these affect your experience?

- To what extent did you encounter any enablers or aspects of the process that made the switching experience easier or more efficient and effective?
 - If so, how did these affect the process and your experience?
 - PROBE: To what extent did your existing cloud provider facilitate the switch or make the switch more difficult? Tell me about how.
- What do you think about switching between cloud providers/between public clouds in the future?
 - How seriously are you considering this as an option for your business? What, if any, benefits might you expect from switching in the future?
 - What, if any, costs might you expect from switching? Why?
 - How attractive are the current set of public cloud providers in the UK market to you as a cloud customer?
 - What would prompt you to consider switching between cloud providers in the future? Why?
- For the use cases you run on your public cloud, have you ever switched these to other IT environments such as a private cloud (e.g. hosted or on-premises) or to traditional IT?
 - If so Why/Why not?
 - Would you consider doing this in the future? Why/Why not?

Attitudes to future developments including AI (5 mins)

Explore attitudes to AI-enabled services and future cloud developments

- How do you see your use of public cloud infrastructure services changing over the next five years?
 - Why?
 - listen for/probe if raised new cloud suppliers, multi-cloud models and mentions of AI-enabled services.
- Thinking now about developments in the cloud market and specifically about AI-enabled services:
 - Do you currently use any AI-enabled cloud services from your (insert provider/s)?
 - If so, what use cases are involved?
 - Do you have any AI-enabled services which are dependent on cloud infrastructure?
 - If so, what are the implications, if any, of this for your use of public cloud infrastructure services? How, if at all, is this affecting your choice of public cloud provider/s?
- Given the discussion we have just had, are there any other issues or points you would like to mention about your experience of using public cloud infrastructure services and how the market is developing in the UK?
- Are there any aspects of the way individual providers or the overall public cloud market currently works that you would like to see change?

Thank you and close



THANK YOU

Jigsaw Research Ltd
3rd Floor, 53 Parker Street
London WC2B 5PT
Tel: +44 (0)20 7291 0810
Web: www.jigsaw-research.co.uk
Email: info@jigsaw-research.co.uk

Registered in England no.: 3650224
Registered Address: 11 Soho Street
London W1D 3AD
VAT no.: 326 011 940

