Part D

Annexes

Table of Contents

Contents **Page**

Annex 1

Key competitors in the cloud market & evolution of rivalry through innovation and strategy

1 Microsoft

(1) AWS was the first to offer a modern public cloud service in 2006.[1] Former Senior Vice-President of Research and Development for the Online Services Division[2] and current CEO, Satya Nadella, later reflected:

> *We knew by looking at what Amazon was doing that we needed to reinvent ourselves, and by the way, the margins were going to be very, very different – tough challenge.*[3]

(2) Microsoft announced its intention to build Windows Azure and Azure Services Platform (PaaS products) in 2008.[4] Eventually released in 2010, Microsoft was observed to be slow in delivering features.

(3) An industry analyst (in 2012) and the CTO for Microsoft Azure (in 2014) both observed that Microsoft's PaaS offering failed to meet the needs of large and small organisations seeking to migrate to the cloud.[5] Microsoft CEO Steve Ballmer later acknowledged:

> *We started actually with platform as a service instead of infrastructure as a service. Probably we would do that a little bit differently. It cost us a little bit of time in the eventual battle, if you will, with AWS.*[6]

(4) In response, Microsoft invested heavily in R&D to improve Azure, such as launching an IaaS offering in 2013 to support customer transitions to the cloud.[7] By 2014, Gartner rated Microsoft as a "leader" in IaaS,[8] and industry reports estimated that Microsoft had grown its cloud business by 154% year over year.[9] Industry analysts noted that Microsoft's products and services now offered an alternative to AWS.[10] By 2015, Microsoft had become the second-largest provider of cloud services after Amazon.[11] AWS remained larger than its three closest competitors (Azure, Google and IBM) combined in 2017[12] and it continued to lead competitors in terms of features and scalability.[13]

---

[1] History of AWS - W3schools.

[2] Microsoft Welcomes new CEO Satya Nadella | by Brad Groux | Microsoft Expert (msft.expert).

[3] Satya Nadella wishes Microsoft got into public cloud earlier (cnbc.com).

[4] Microsoft, Windows Azure and the Azure Services Platform: Making Microsoft's Software-plus-Services Vision a Reality - Stories.

[5] Wired, Time for a Second Look at Windows Azure?.

[6] Steve Ballmer wishes Microsoft entered cloud computing business sooner (cnbc.com).

[7] Microsoft Offers New IaaS Services To Compete With Amazon -- Redmondmag.com.

[8] Amazon and Microsoft top Gartner's IaaS Magic Quadrant | ZDNET.

[9] "AWS revenues exceed $1 billion per quarter, but Microsoft growing faster," Computer Weekly, May 19, 2014, ComputerWeekly.com.

[10] Microsoft Preps for Public Embrace of Linux | WIRED.

[11] "Amazon Dominates Public IaaS and Ahead in PaaS; IBM Leads in Private Cloud," Synergy Research Group, October 30, 2016.

[12] "According to Canalys, AWS dominates the market with 33.8 percent global market share while its closest competitors—Microsoft, Google, and IBM—together account for 30.8 percent of the total market. Other smaller competitors include Oracle and Alibaba Cloud." (see Here is Why AWS Still Dominates the Cloud Market (comparethecloud.net)).

[13] Here is Why AWS Still Dominates the Cloud Market (comparethecloud.net).

(5)   Throughout its history as a challenger to AWS, Azure has emphasised portability and ease of migration to gain customers who may have already adopted AWS. With the trend towards verticalisation, Microsoft launched an initiative in 2017 to organise its account teams along industry lines and focus on vertical markets. Partner ISVs have contributed to this push in the form of vertical- or industry-specific solutions built on top of Microsoft's cloud offerings, and integrators providing industry experience for customers and ISVs.

(6)   Price competition has driven cloud providers, including Azure, to invest in innovation to reduce costs. For example, Project Silica was introduced in 2017 to develop the world's first storage technology designed to store data in quartz glass: a low-cost, durable solution offering a lifetime of thousands of years and eliminating the costly cycle of periodically copying data to a new media generation.[14]

2   AWS

(7)   AWS was the first to offer a modern public cloud service when it launched the Amazon Elastic Compute Cloud (EC2) and Simple Storage Service (S3) in 2006.[15] While competitors have gained ground recently, AWS remains the de facto industry leader.

(8)   In recent years, AWS has shifted its sales and solution development strategy towards industry verticalisation. As the leader in IaaS, it has targeted broad fields like automotive and healthcare and work with large organisations to build "industry-configured" products based on market needs. Through its work with large enterprises (e.g. Toyota and Volkswagen in the automotive industry),[16,17] it can develop tailored solutions well-suited for large customers. Once these solutions are developed, they are marketed and sold to the broader industry.

(9)   AWS still retains a global market share of nearly double the next largest competitor, Azure. It has also been actively investing in its cloud infrastructure. More recently, AWS developed Graviton instances on its Arm-based chip design, offered at optimised price-performance ratios compared with x86 instances.[18]

3   Google

(10)   First launched in 2008 as a PaaS in the form of Google App Engine, Google Cloud has evolved into a global public cloud platform. As a challenger to AWS, just like Microsoft, it has emphasised portability to gain customers. For instance, it has developed Kubernetes, the most popular container solution, and container orchestration. Google has also leveraged its reputation for AI/ML capabilities in its cloud products, offering services based on its Tensor Processing Unit chips, which work with its widely used TensorFlow ML library.[19]

(11)   GCP has focused on their strengths in AI/ML and advertising sale verticals. It has partnered with strategic customers, such as the Mayo Clinic and Ascension in the healthcare space,[20,21] to build solutions that leverage these strengths.

[15]   History of AWS - W3schools.
[15]   History of AWS - W3schools.
[16]   Toyota and Amazon Web Services Collaborate on Toyota's Mobility Services Platform - Toyota USA Newsroom.
[17]   Volkswagen Industrial Cloud Case Study (amazon.com).
[18]   "IDC MarketScape: Worldwide Public Cloud Infrastructure as a Service 2022 Vendor Assessment", Nov 2022, p 6.
[19]   "IDC MarketScape: Worldwide Public Cloud Infrastructure as a Service 2022 Vendor Assessment", Nov 2022, p 8.
[20]   Our partnership with Ascension | Google Cloud Blog.
[21]   How Google and Mayo Clinic will transform the future of healthcare | Google Cloud Blog.

(12) While Google Cloud has consistently garnered a global revenue share somewhat smaller than that of AWS, the service is expected to generate significant profits for the company in the near future. In a 2022 report, analysts noted a renewed emphasis on sales, with headcount in that department more than tripling between 2019 and 2022, and an increase in product release cadence that led to a modest growth in market share.[22] They found margins improving meaningfully as the business scaled and stated that "*we see no structural reason why Google Cloud operating margins should not mimic that of its larger peers such as Azure and AWS*".[23] In 2023, they commented that Google Cloud revenue had grown 32% year-over-year, and operating margins had improved even more than anticipated.[24]

4   Oracle

(13)   Prior to the widespread adoption of cloud computing, Oracle dominated the market for high-end databases[25]. On Oracle's dismissal of cloud computing, analysts noted that "*there is little incentive for Oracle to embrace and cannibalize over $35bn in revenues by investing in generally more efficient multi-tenant SaaS.*"[26]

(14)   As noted in the Response above, Oracle's lack of an IaaS offering prior to 2016 was a source of weakness - as one AWS customer noted, "*customers would have no interest in hosting their database in an Oracle cloud and the application somewhere else, as this introduces unnecessary latency.*"[27] Eventually, Oracle launched its public cloud offering, Oracle Cloud Infrastructure (OCI), in 2016. To steer customers to OCI, Oracle focused on ensuring that Oracle workloads would run faster and cheaper on OCI than AWS.[28] In 2018, Oracle unveiled a key technical innovation, Autonomous Database, to further strengthen its IaaS position.[29] Autonomous Database is a fully automated database service that makes it easy for all organisations to develop and deploy application workloads regardless of complexity, scale, or criticality.[30] Between 2018 and 2020, Oracle also focused on technically differentiating OCI's infrastructure from AWS, with a particular focus on security.[31] 2020 saw several customers moving workloads from AWS to OCI, with 8x8[32] and Sky.One[33] moving all of their workloads, and Zoom,[34] Jeffries,[35] and DTCC[36] moving some critical workloads to OCI.

(15)   OCI has quickly built up a broad range of VM instance types, support for containers via the Container Engine for Kubernetes, and storage solutions. Oracle claims its OCI block storage

---

[22] Benjamin Black et al, "One Alphabet Avenue – The Unorthodox Real Estate Play," *Deutsche Bank Research*, Mar 10, 2022, p 14.

[23] Benjamin Black et al, "One Alphabet Avenue – The Unorthodox Real Estate Play," *Deutsche Bank Research*, Mar 10, 2022, p 16.

[24] Benjamin Black et al, "Striking A New Tone," *Deutsche Bank Research*, Feb 3, 2023, p 3.

[25] Oracle had 44.4% of the database market in 2006, second highest was IBM with 21.2%. Source: IDC: Oracle extended lead over IBM in 2006 database market | Computerworld.

[26] Deutsche Bank, "Oracle World brings few surprises," 26 September 2008.

[27] Deutsche Bank, Oracle Corporation, 6 June 2016.

[28] Oracle 2017 Q3 Earnings Call, Capital IQ.

[29] Oracle 2018 Q3 Earnings Call, Capital IQ.

[30] Ibid.

[31] Oracle 2019 Q4 Earnings Call, Capital IQ.

[32] Oracle 2020 Q4 Earnings Call, Capital IQ.

[33] Ibid.

[34] Oracle 2021 Q1 Earnings Call, Capital IQ.

[35] Oracle 2020 Q4 Earnings Call, Capital IQ.

[36] Oracle 2020 Q3 Earnings Call, Capital IQ.

costs 57% less than other services without sacrificing performance.[37] OCI Dedicated Region capabilities are unique among public IaaS providers, offering the same 100-plus cloud services available through OCI's public cloud regions on premises in the physical location of the customer's choosing.[38]

(16)   2023 also represented several wins for OCI over AWS, with One-One Intelligent Information Technology,[39] Oxford Nanopore,[40] and Entel[41] moving critical workloads from OCI to AWS. OCI has been recognised as a Visionary in the 2022 Gartner Magic Quadrant for Cloud Infrastructure and Platform Services (which includes IaaS and PaaS).[42]

(17)   OCI is gaining market momentum rapidly, with 53% year-over-year revenue growth in 2022 for the IaaS/PaaS business. OCI touts what it calls its "*second mover advantage*", or its ability to build upon the first-generation public clouds like AWS, Azure, and GCP to offer a compelling combination of performance, security, availability, and price.

(18)   Finally, Oracle is reported to be priced very competitively relative to IaaS competitors. It is reported to also maintain a sizeable R&D budget, spurring key innovations including OCI Gen 2, Autonomous Database, and MySQL Heatwave. Going forward, analysts expect R&D to remain a priority with a particular focus on OCI, Fusion Apps, the migration of Cerner's software portfolio to the Cloud and the core Oracle database business.

5   IBM

(19)   IBM first investigated a cloud strategy in 2007 via a research partnership with Google.[43] In 2007, IBM also announced Blue Cloud, an offering that would allow customers to maintain distributed computer architectures in their own data centers. However, IBM did not yet support the hosting and infrastructure services like AWS.[44] In March of 2010, IBM finally announced SmartCloud, an IaaS product.[45] However, IBM's cloud focus remained on customer data centres, and the IBM-hosted offering was limited – targeting application development but not end production (i.e., running the applications).[46][47] In 2011, IBM offered an improved SmartCloud with expanded hosting capabilities[48] but still struggled to gain customer traction.[49] IBM subsequently acquired IaaS provider Softlayer in 2013.[50] Industry

---

[37]   "IDC MarketScape: Worldwide Public Cloud Infrastructure as a Service 2022 Vendor Assessment", November 2022, p 12; and Brad Zelnick et al, "Charting the path to FY26 targets; all roads run through the Cloud," *Deutsche Bank Research*, Feb 3, 2023.

[38]   OCI Dedicated Region (oracle.com).

[39]   Oracle 2023 Q2 Earnings Call, Capital IQ.

[40]   Oracle 2023 Q2 Earnings Call, Capital IQ.

[41]   Oracle 2023 Q1 Earnings Call, Capital IQ.

[42]   Oracle, Oracle recognized as a Visionary in 2022 Gartner® Magic Quadrant™ for Cloud Infrastructure and Platform Services; Gartner, Magic Quadrant for Cloud Infrastructure and Platform Services.

[43]   Google and the wisdom of clouds (nbcnews.com).

[44]   IBM's Blue Cloud is Web Computing By Another Name | TechCrunch.

[45]   Internet Archive-WaybackMachine.

[46]   Red Hat's KVM virtualization proves itself in IBM's cloud | Network World.

[47]   What is each server for: Dev, Test, UAT, Staging, Demo and Production – dblSec (itsec.rs).

[48]   Internet Archive-WaybackMachine.

[49]   IBM Faces a Crisis in the Cloud | Mashable.

[50]   Cloud computing and AI: Can IBM finally catch the wave this time around? | ZDNET.

analysts reported that "*SoftLayer was only founded in 2005, and IBM could have saved itself $2 billion by the timely development of its own software.*" [51,52, 53]

(20) In recent years, IBM has increased its focus on hybrid cloud solutions to help customers run heterogeneous infrastructure and have partnered with public cloud providers.[54] In 2018, IBM announced their acquisition of RedHat, a crucial move in their overarching hybrid cloud strategy.[55]

(21) IBM has continued to invest in computing. It was one of the first organisations to release a roadmap to quantum computing. IBM currently sees the main consumer of quantum computing to be scientific organisations and governments similar to the way they use supercomputing for their models and calculations.

6 OVHCloud

(22) OVHCloud began offering public cloud IaaS in 2011. OVHCloud has developed a range of offerings at the IaaS level, with support for AMD and Intel-based VMs, block, file, and object storage. In recent years, it has invested in building out higher-layer software services such as AI/ML, databases, and container orchestration.[56]

(23) OVHCloud provides valuable service adjacency, given that it also offers dedicated private servers, web hosting, and hosted private clouds, which makes it a helpful partner for companies still early in their transition to the public cloud.[57] It also strongly emphasises data security and sovereignty, positioning itself well to tap into demand for data-sovereign cloud solutions, especially in regulated sectors and industries that manage sensitive data. OVHCloud is also expected to benefit from a greater focus on integrated PaaS solutions over the coming years.[58] As with larger public cloud providers, OVHCloud has benefitted from the broader growth in cloud services, with its recent Public Cloud 1H revenue growing 20.5% year-over-year.[59]

(24) EU-based cloud players (such as OVHCloud) will continue to be relevant competitive constraint in the cloud market as a result of (amongst other factors) EU digital sovereignty and the strategic policy and regulatory environment of the EU (including EU Member States).

---

[51] Cloud computing and AI: Can IBM finally catch the wave this time around? | ZDNET.

[52] Sources observed that "IBM has acquired cloud-computing infrastructure provider SoftLayer Technologies in an effort to better compete with Amazon." See: IBM Acquires Cloud-Computing Firm SoftLayer (pcmag.com).

[53] "Instead of investing in cloud data centers, IBM has managed by press release for years, with current CEO Virginia Rometty continuing the policy of her predecessors. IBM hypes blockchain, it hypes Artificial Intelligence, and it hypes its Watson front-end. But it lacks the capital firepower to capitalize." (A Tech Giant No More: IBM Is Too Small to Compete in the Cloud Era | InvestorPlace; https://investorplace.com/2018/04/international-business-machines-corp-ibm-too-small-to-compete-in-the-cloud-era/; Cloud computing and AI: Can IBM finally catch the wave this time around? | ZDNET).

[54] IBM 2021 Q4 Earnings Call, Capital IQ.

[55] IBM 2018 Q4 Earnings Call, Capital IQ.

[56] "IDC MarketScape: Worldwide Public Cloud Infrastructure as a Service 2022 Vendor Assessment", November 2022, p 14.

[57] "IDC MarketScape: Worldwide Public Cloud Infrastructure as a Service 2022 Vendor Assessment", November 2022, p 14.

[58] Toby Ogg et al, "Macro headwinds though cost reduction plan and pricing support in H2," *J.P. Morgan; Europe Equity Research*, Apr 20, 2023, p 2.

[59] Toby Ogg et all, "Q2 Private/Public ahead though lowering revenue & EBITDA guidance," *J.P. Morgan; Europe Equity Research,* Apr 18, 2023, p 1.

7 C3 AI

(25) C3 AI markets itself as an Enterprise AI software provider for accelerating digital transformation. It offers vertical-focused services in higher layers of the stack, while operating across multiple public cloud layers. They have a multi-cloud focus, with their key product running on Azure, AWS, and GCP, private clouds and customers' data centres. Their vertical focus, with industry-specific apps, allows them to start with one application at a business, gain expertise in that domain, and expand to adjacent areas.

(26) C3 AI has jumped to the forefront in the industry, with Forrester naming them a Leader in the 2022 Forester Wave for AI and Machine Learning Platforms, ahead of Microsoft, Google, and Amazon Web Services.[60] They have also been named a Google Cloud Technology Partner of the Year award winner[61] and a Forbes Top 100 Cloud company multiple times.[62]

8 Evidence of numerous recent entrants

(27) Evidence suggests that, in recent years, many providers of cloud services have emerged by focusing on certain customer groups or functionalities. Some new entrants even specialise in providing access to basic infrastructure at a low cost. Below is an illustrative list of entrants. While not exhaustive, it illustrates the continuing innovation in this sector globally and in the UK.

8.1 CoreWeave

(28) CoreWeave was founded in 2017 to address the need for Graphics Processing Unit (GPU) computing,[63] especially for generative AI technologies. It was initially focused on cryptocurrency applications but pivoted to general-purpose computing and generative AI technologies. CoreWeave claims its hardware for inference (i.e. serving AI models) is industry-leading, and able to "autoscale" within three seconds. CoreWeave also recently launched an accelerator program for customers, which provides companies compute credits in addition to discounts and other hardware resources. Recently, in April 2023, CoreWeave secured $221 million in funding, including from NVIDIA, valuing CoreWeave at $2 billion.

8.2 Paperspace

(29) Paperspace (which focuses on generative AI workloads and a serverless tool that abstracts from underlying hardware) started in 2014. It helps developers build AI and machine learning applications based on its proprietary GPU-based hardware. In addition to providing servers, it also released a serverless tool to make it easier to deploy and manage AI and machine learning workloads.[64]

8.3 DigitalOcean

(30) DigitalOcean is a self-serve SaaS business founded in 2012 with revenues of $250 million at the end of 2019.[65] DigitalOcean has focused on targeting developers as its core

---

[60] The Forrester Wave™: AI/ML Platforms, Q3 2022.

[61] C3 AI Named Google Cloud Technology Partner of the Year for AI and Machine Learning.

[62] C3.ai Awarded Forbes Cloud 100 and Datanami Honors.

[63] AI workloads require GPU computing capacity.

[64] Paperspace scores $13M investment for AI-fueled application development platform | TechCrunch.

[65] DigitalOcean raises $100M in debt as it scales toward revenue of $300M, profitability | TechCrunch.

customers, wooing them with easy-to-understand platform tutorials and simple and affordably priced services.[66]

8.4 Scaleway

(31) Scaleway, launched in 2014, focuses on providing cloud services to small businesses by offering very cheap cloud instances.[67] It also offers multi-cloud management tools and serverless architecture. It has been expanding its data centres in Europe, and in 2022 it built further data centres in Amsterdam and Warsaw.[68]

8.5 Clever Cloud

(32) Clever Cloud, a French start-up, is a cloud-hosting company that operates a Platform-as-a-Service (or PaaS) for developers and handles all code deployment so that developers can focus on coding.[69] It builds its GPU-based instances for machine learning purposes and builds its software.[70] It provides cloud computing resources by abstracting away from infrastructure management by running code on a web interface. Clever Cloud clients include Airbus, MAIF, Compte Nickel, Sogeti and the South African Ministry of Health.

9 Expansion of data centres over time

(33) Figure 1 illustrates the geographic expansion of data centre regions[71] by cloud providers over time. Notably, Oracle's and GCP's regional coverage has surpassed AWS and Azure's since 2022. Following a strategy to focus on its hybrid cloud offering contrast, IBM's recent investments in cloud Infrastructure have been to add new data centers and availability zones to *existing* regions, rather than expanding to new regions.[72]

Figure 1: New regions by cloud provider over time

---

66 Amid the cloud giants, small providers find their niche (fastcompany.com).

67 Scaleway launches cloud instances that cost $2.10 per month | TechCrunch.

68 About us | Scaleway.

69 Clever Cloud launches GPU-based instances | TechCrunch; Home – Clever Cloud (clever-cloud.com).

70 Clever Grid On-Demand GPU for ML & AIClever Grid (·) On-Demand GPU for ML & AI.

71 A region is a localized geographic area in which public cloud resources are available. It may consist of one or more data centres and availability zones. Availability zones are isolated data centers located within specific regions in which public cloud services operate. See: The Big Deal About Cloud Regions, LucidChart; Availability Zones, TechTarget.

72 IBM, IBM Cloud Global Data Centers.

Total Cloud Regions
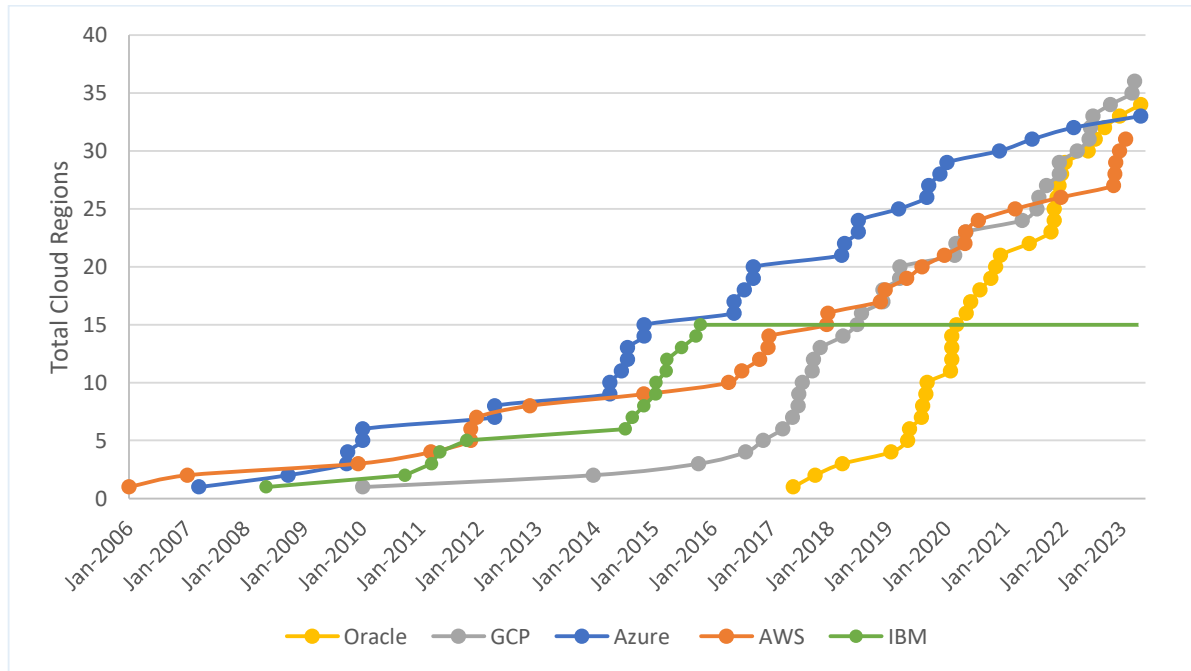
(34)   Oracle was the last to establish a data centre in late 2016[73], and since 2019 has invested heavily in data centres as seen in their steep footprint growth. Oracle's relatively late and aggressive expansion is supported by the pivot in its cloud strategy described above. IBM specifically invests substantially in creating Multi-Zone Regions (MZRs), which are regions containing three data centres,[74] to improve the availability and redundancy of services within a particular region.[75] Such data centre architecture is particularly suitable for multi-cloud deployments.[76]

---

[73]   The Motley Fool, Oracle is Going to War with Amazon.

[74]   IBM, Multi-Zone Region (MZR) overview.

[75]   IBM, Data Center Consolidations.

[76]   IBM launches new Cloud Multizone Region in Spain | ZDNET.

Annex 2
Key innovations in cloud infrastructure

(1)   Section 1 below summarises the breadth and scale of innovations made each year in cloud infrastructure by AWS, GCP and Azure. Investments and innovations are prevalent in every single layer from infrastructure to analytics, management tools, and other services. AWS, Azure, and GCP alone have each introduced hundreds of product improvements and new products in just the last five years. A list of key innovations specifically for infrastructure-related products is presented in Section 2.

1   Pace of innovation by category and type

(2)   Figures 1-3 below show the breadth as well as the scale of cloud-related innovations launched by AWS, Azure, and GCP in the last five years. It depicts an environment of continuous and robust innovation, with each provider making several releases in a wide variety of categories each year.

Figure 1: AWS # innovations by category, 2017 - 2022
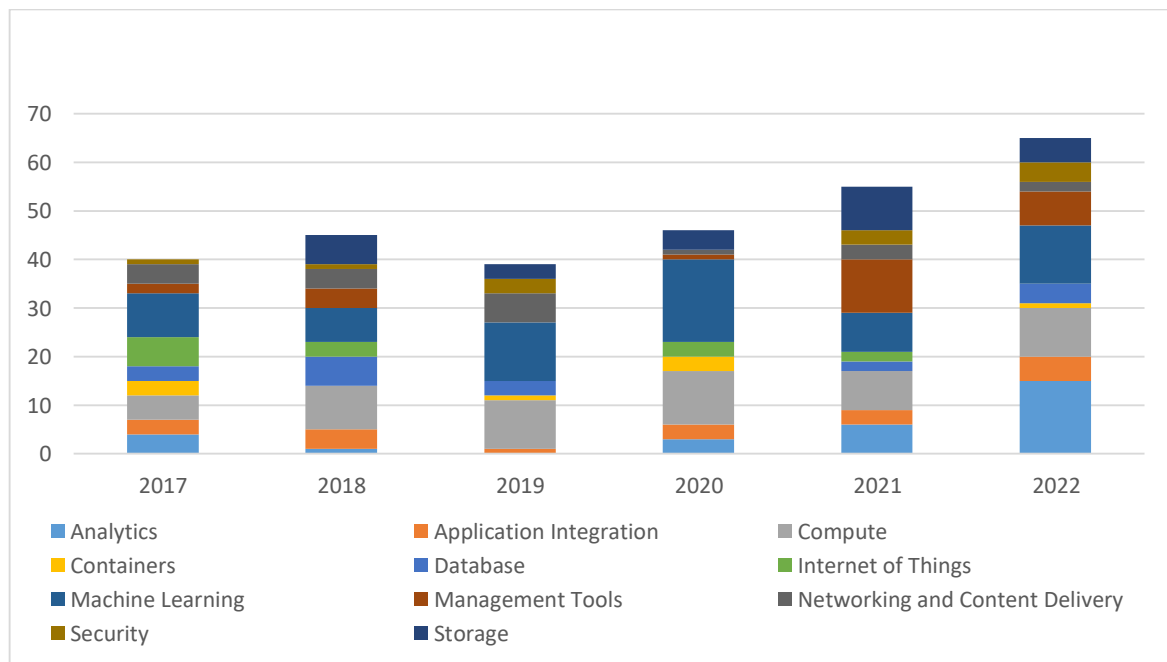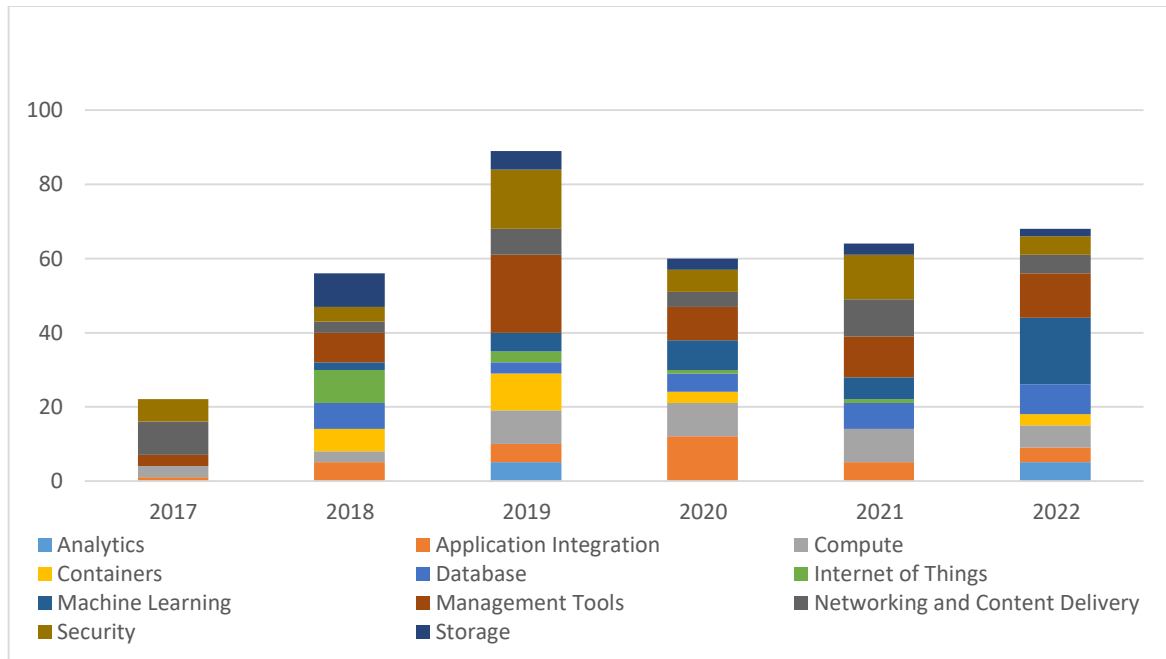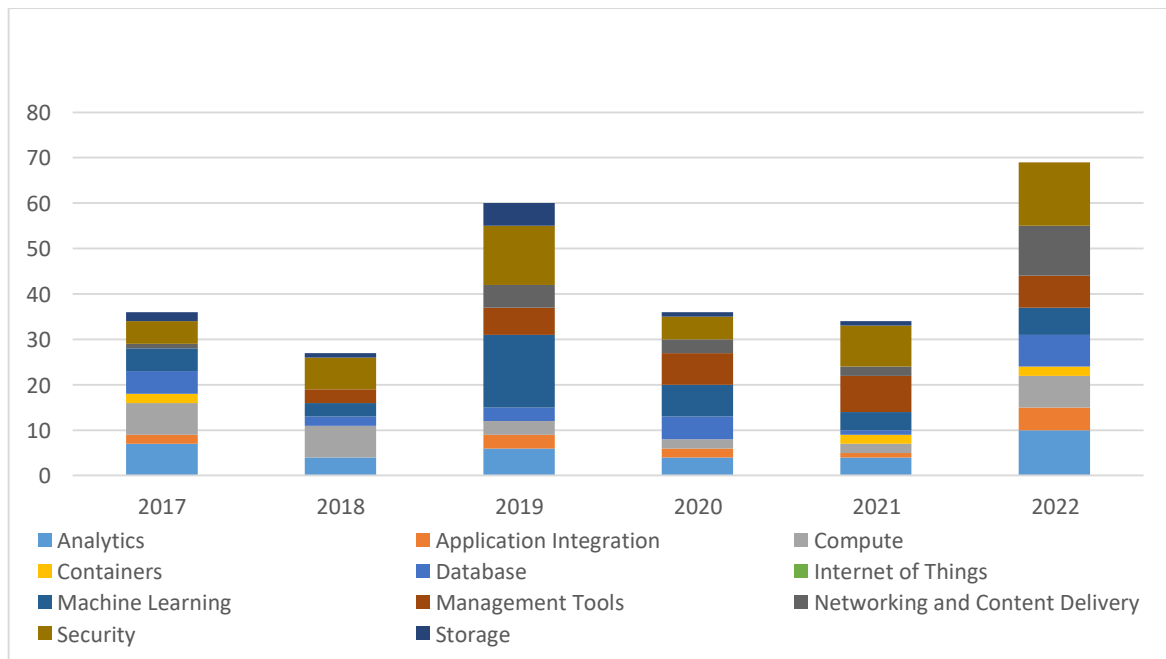
Figure 2: Azure # innovations by category, 2017 - 2022

Figure 3: GCP # innovations by category, 2017 - 2022

1 Key innovations by year by each hyperscaler

1.1 Compute

	aws	Microsoft Azure	Google Cloud
(2017)	• Several instances for data-intensive workloads	• Several instances optimized for low-latency	• Expansion of vCPUs accessible by instances • Dynamic load balancing for instances
(2018)	• Arm-based instances, new high-performance instances • Instance hibernation to facilitate load balancing	• Instances optimized for graphic-intensive workloads, HPC workloads • Features that facilitate VM management	• Instances for SAP HANA workloads • TPU chip for TensorFlow Lite ML models
(2019)	• Instances optimized for ML Inferencing • Instances for Redshift • Instances optimized for EBS	• Instances for RDBMS, HPC workloads • Instances for graphics-intensive workloads	• Instances for general-purpose workloads • OpenTitan chip to enable transparency for VMs
(2020)	• Instances for building Xcloud applications • Instances for graphics-intensive workloads	• Instances optimized for lower latency, RDBMS workloads • Features that facilitate VM load management	• Instances that feature dynamic load management • Instances with expanded GPU access for HPC, ML workloads
(2021)	• Instances optimized for high disk I/O performance • Instances integrated with Gravitron2 price-performance benefits	• Instances optimized for general CPU performance • Instances optimized for memory-intensive enterprise applications	• Instances optimized for HPC, ML workloads • Instances optimized for cost-effective, scaled out workloads
(2022)	• General purpose instances for EC2 with expanded network bandwidth • Instances optimized for HPC, deep learning inference	• Instances optimized for HPC workloads • Features that facilitate patching of VMs without the need to reboot	• Instances optimized for HPC workloads • Instances optimized for DBMS workloads

1.2 Storage

	aws	Microsoft Azure	Google Cloud
(2017)	• De-duplication software that reduces the cost of storage for unstructured data	• Archive Blob storage for rarely accessed data in the lowest price tier	• Expanded memory for instances, doubling vCPUs customers can run per instance
(2018)	• Intelligent tiering, which optimizes storage when access patterns change • Simplified S3 storage management	• DataBox physical storage to facilitate transfer of data to Azure data center • Immutable data storage for regulated data	• Immutable data storage for regulated data
(2019)	• Analytics-optimized storage for S3 • Deep Archive storage, for long-term, infrequently accessed data	• Upgraded storage performance for Azure Disks • Automatic caching of data held both on-premise and in Azure, reducing latency	• Archive storage for long-term, infrequently accessed data • Managed file storage for HPC • Active disk-duplication across two zones in the same region
(2020)	• Duplication of multiple copies of data in different storage classes • New SSD EBS volumes for enhanced storage	• Private storage of data over a VPN • Shared disks, which allow a single disk to be used across multiple VMs	• Scaled file deployment for high-capacity file stores • Upgraded SSDs that reduce instances required for workloads
(2021)	• Nitro SSDs, optimized for high-scale Cloud operations • Service to automate backup and restoration of applications	• Zone-redundant storage for mission-critical workloads • Performance tiers for Azure premium SSD, for resource-intensive workloads	• Service to deliver critical file-based applications with high regional availability • New higher-capacity SSDs for Compute Engine
(2022)	• Shared file storage that automates provision and capacity management • High-speed file cache to reduce latency	• Azure Premium SSD v2, further optimized for resource-intensive workloads • Cloud-native storage area network	• Automatic storage-class provisioning optimized for price-performance • Persistent Disk that enables dynamic performance tuning

1.3 Networking

	aws	Microsoft Azure	Google Cloud
(2017)	• Peering relationships between VPCs • HTTP processing at nearby locations to reduce latency	• Expanded VPN gateway performance • Peering between different Virtual Networks	• Common network for multi-tenant deployments • Dynamic allocation of server resources
(2018)	• AWS Outposts for native AWS services on on-premises facilities • Dynamic routing of traffic to multiple regions	• Physical Network Appliance optimized for AI Edge computing • Increased connectivity between on-prem networks	• High-speed network interfaces for TPU devices • Sole Tenant Nodes optimized for individual customer use cases
(2019)	• Optimized connectivity between AWS and on-prem networks • AWS Compute and Storage embedded in 5G networks	• Network insights added to Azure Monitor for centralized network management • Expanded satellite connectivity for ExpressRoute	• Facilitated traffic control across multiple regions • High availability VPN for secure connections between on-prem and VPC
(2020)	• Fleet manager for managed instances from a single location • Network diagnostics for reachability between VPCs	• Upgraded satellite communication capabilities for data analysis • Physical Edge computing devices for remote areas	• Undersea cable between U.S. and U.K. for network resilience • Centralized network monitoring
(2021)	• Optimized connectivity between on-prem networks • Automated network management features • Automated network configuration risk assessment	• Single-tenant hosting for increased reliability in high traffic • Managed Virtual Network for secure connections between data stores	• Enhanced connectivity between on-prem, GCP, and other Cloud networks • Upgraded traffic management
(2022)	• Enhanced communication between multiple cloud services • Enhanced private connectivity between services	• DNS service to conditionally forward DNS queries from virtual network • Automated traffic distribution for low latency	• Dynamic compression to reduce the latency of transmitted across network • Private access for managed services from inside VPC

1.4 Middleware

	aws	Microsoft Azure	Google Cloud
(2017)	• Managed service for Kubernetes on AWS • Serverless development environment for applications	• Managed service for Kubernetes on Azure • Serverless functions development environment	• Managed tool for building docker containers on GCP • Python SDK for managed data processing service
(2018)	• Central management for code and data shared across functions • Serverless functions development environment	• Cloud-enabled managed app building platform for real-time • Upgrades to serverless functions development environment	• CI/CD platform for building and testing applications in the Cloud • Serverless functions development environment
(2019)	• Environment for ML development • Dynamic initialization of functions for latency-sensitive services	• Extension of Azure management to any infrastructure • Environment for ML development	• IDE plug-ins for easier application integration with Cloud • Auto-migration of VMs to GKE containers
(2020)	• Manual Kubernetes cluster creation for developers • Simplified access to AWS-enabled shell prompt for developers	• Environment for .NET microservices development • Centralized cluster Stop/Start feature for easy cluster management	• Instances that feature dynamic load management • Instances with expanded GPU access for HPC, ML workloads
(2021)	• Kubernetes autoscaler optimized for application load • Managed delivery service for container and serverless applications	• Facilitated migration to Kubernetes containers • Service for automated configuration and management of servers	• Managed control plane for Kubernetes containers • Standardized Kubernetes between workloads located on-prem and in GCP VMs
(2022)	• Upgrades to serverless development environment • Facilitated deployment of Kubernetes software to clusters	• Container optimized for ML training workloads • Managed serverless container service for modern apps	• Fleet management for container clusters across clouds and on-prem • Facilitated deployment of releases to various GCP services.

1.5 Security

|	aws	Microsoft Azure	Google Cloud
(2017)	• VPC Endpoints made optimizable for Privacy	• DDoS protection for virtual networks from targeted attacks	• Centralized key management service • Service for secure application access
(2018)	• Centralized security hub across accounts • Centralized service for management, retrieval, and rotation of credentials	• Cloud-native network security service • Protection for data processed in public clouds	• FIDO Security Key for easy verification of integrity • ML-based automated DDoS protection
(2019)	• Centralized data security visualization service • Isolated compute environments for sensitive data	• Service to modernize any infrastructure with cloud security protection • Centralized traffic flow monitoring for security	• Managed protection for internet-facing applications • Customer-defined security perimeters around GCP resources
(2020)	• ML-based system to detect abnormal behavior in industrial machinery	• Isolated hosting environment for most sensitive web workloads • Managed import and export over VPN for security	• Centralized certificate authority service • Confidential VMs for secure computing
(2021)	• Automated DDoS mitigation service • Automatic detection of sensitive information hardcoded into storage files	• Single-tenant key management service • Continuous monitoring of device security posture	• Additional data protection features for cloud storage • Automated exploit attempt detection to protect customer environments
(2022)	• Remote secure access for corporate applications without VPN • External key store for on-prem or non-AWS cloud	• Managed security solutions for hybrid and multicloud • Automatic detection of exposed secrets and open-source dependencies	• ML-based adaptive protection to automatically deploy customer security rules • Simplified configuration and segmentation for firewall

Annex 3
Flaws in Ofcom's Survey

1 Flaws in survey interpretation and/or design

(1) To inform its assessment in the Interim Report, Ofcom conducted an online survey during September-October 2022 of 1004 ICT decision-makers in the UK, across size bands and sectors. It also conducted 64 interviews with UK decision-makers with a broad range of organisations. To qualify, decision-makers had to work for organisations that were existing users of cloud computing services (IaaS, PaaS or both) or were considering adoption within 12 months.

1.1 Current and forward-looking concerns of Ofcom materially contradicted by Ofcom's own survey

(2) Ofcom expresses significant concerns that competition is being limited in the market for cloud infrastructure services. In particular, purported lock-in of existing customers increases the hurdles to switching providers and lowers competitive pressure on market actors.

(3) These concerns do not seem to be mirrored by the respondents in the IaaS and PaaS market. Figure 1 presents Ofcom's survey results, where IaaS and PaaS customers were asked to assess the competitiveness in the relevant markets. 89% of respondents indicated that the level of competitiveness in IaaS or PaaS was at least average or above. In fact, the number of respondents attributing IaaS and PaaS a high level of competitiveness far exceeds those with concerns about the competitiveness of the market. Figure 1 also stratifies the results by the cloud services provider used by the client. The bigger the cloud provider is, the more customers think that the market is competitive.

Figure 1: Perceived competition in the IaaS and PaaS Market
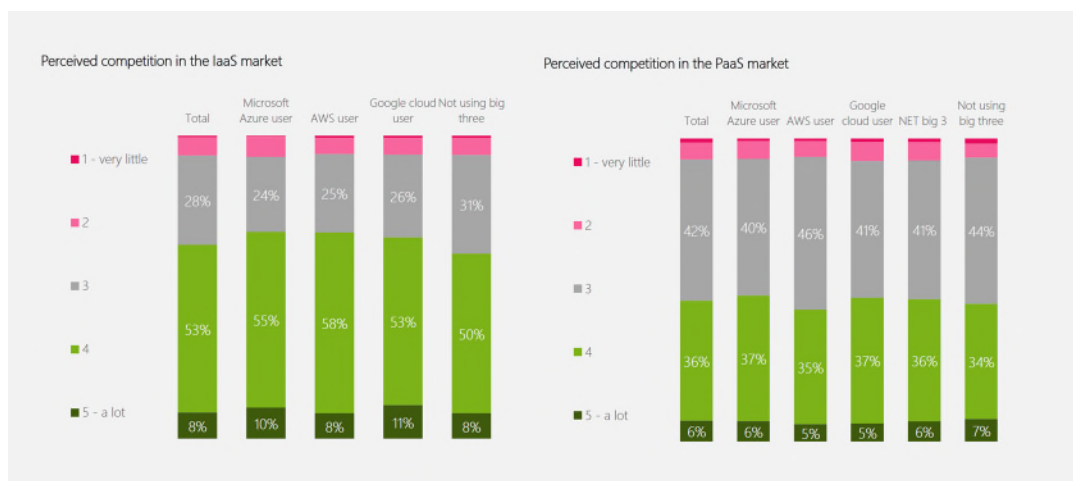
Source: Ofcom and Context Consulting, Cloud Services Market Research, Summary of Findings, March 2023, p.132.

(4) It is not surprising that Microsoft Azure and AWS IaaS users tend to perceive a higher level of competition than the average. The IaaS market especially is characterised by hyperscalers making use of economies of scale and scope to produce IaaS services at the lowest possible costs. Naturally, the two largest – and thus most efficient – competitors offer some of the most competitive products to their customers, rather than locking them in.

(5)   Additionally, Ofcom is concerned about customers' future ability to switch cloud services provider. This concern is again not mirrored by the respondents of the Ofcom survey. As Figure 2 shows, the vast majority of current cloud services customers are expecting or at least considering a multi-cloud solution for their business needs in the future. Currently, more than half of cloud services customers are already employing a multi-cloud solution (see Figure 4 below).

Figure 2: Future consideration of multi-cloud architectures
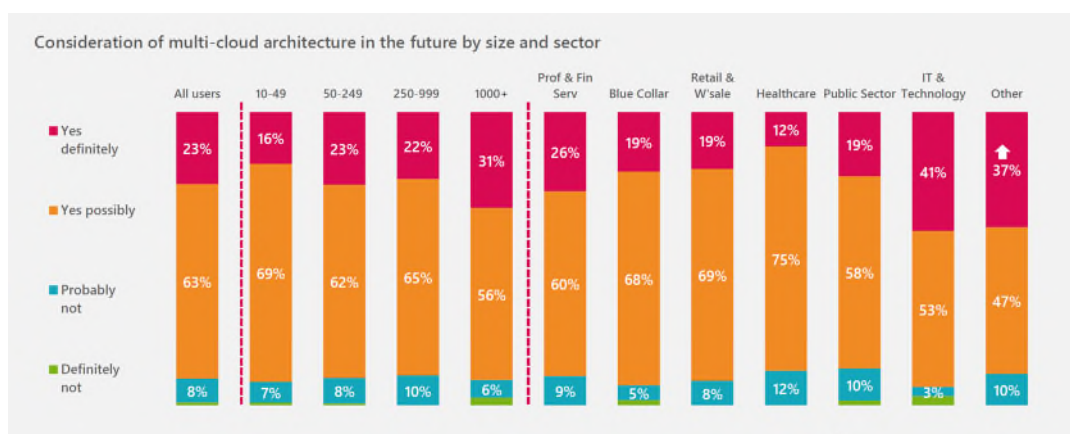
Source: Ofcom and Context Consulting, Cloud Services Market Research, Summary of Findings, March 2023, p.77.

1.1.1   Interim Report conclusions are not fully supported by the underlying data

(6)   As discussed in the next sub-sections, the survey shows that:

(i)   multi-cloud is as common as single-sourcing among the respondents;

(ii)   the estimated price rise does not control for quantity and therefore cannot be distinguished from a spend increase;

(iii)   price rises were higher or similar for hyperscalers and some non-hyperscale providers like Digital Ocean and VMWare and is not consistent with price rises being associated primarily with the business practices of the hyperscalers; and

(iv)   the proportion of customers identifying "existing relationships for other services" as a reason for choosing a supplier is similar across almost all cloud services providers and is not unique to Microsoft.

(7)   Finally, Ofcom's qualitative research raises an important question about switching in the counterfactual. The qualitative research clearly suggested that there is inherent inertia which reduces the incentives to switch, and companies are eager to avoid the disruption this would entail. It found that many decision-makers often perceived any lock-in as a function of internal factors rather than provider-imposed restrictions. Given the inertia, something significant would need to happen to prompt a switch away from a cloud services provider, such as a substantial price hike, a deterioration in technical performance, or significant security concerns.[77] Therefore, it is not clear from the Interim Report to what degree the business practices highlighted impact customer switching and multi-cloud behaviour in

---

[77]   Interim Report Annex, para A7.39.

practice, over and above the inherent costs of switching and multi-sourcing in cloud computing.[78]

Figure 3: Switching behaviour qualitative survey respondents
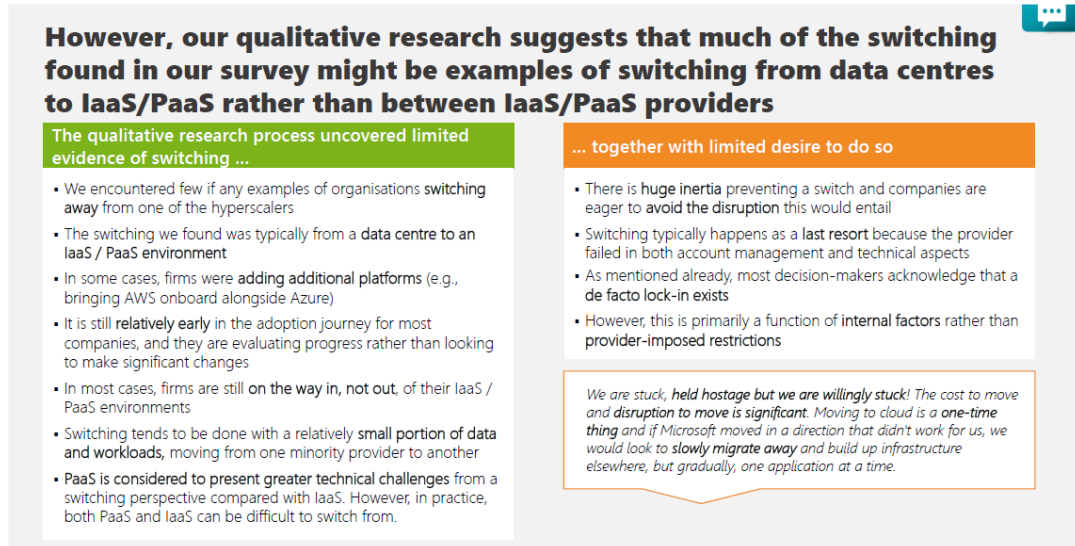
Source: Slide 110, Cloud Services Market Research, Ofcom's interim report.

1.1.2 Prevalence of multi-cloud

(8)   Ofcom concludes that about half (48%) of the respondents to its survey do not multi-home. However, the survey shows that 52% use more than one cloud provider. Ofcom reports that if respondents who were only using private cloud are excluded, the proportion of respondents who were using more than one IaaS/PaaS provider **increases to 61%**.[79] Ofcom purports that this figure is likely to overstate the fraction of customers using multiple public cloud providers because customers may have understood IaaS/PaaS providers to include private cloud providers and ISV providers, as well as public cloud providers. Ofcom does not provide any evidence of the extent of this over-estimation in the Interim Report, so it is not clear how to interpret the results.

(9)   Ofcom's acknowledgement in the Interim Report that customers may have "primary" suppliers (such as AWS or Microsoft due to their existing relationships) as well as "secondary" suppliers (for new workloads that are sufficiently separate to a customer's existing cloud usage) contradicts Ofcom's argument that there are barriers for smaller providers as there is acknowledgement of multi-homing prevalence.[80] In addition, Ofcom does not distinguish between IaaS and PaaS: customers might not switch IaaS but could do so for PaaS. Customers already tend to use software from multiple vendors and will look for the solution most tailored to their needs.

---

[78]   Ofcom acknowledges these inherent costs at para 6.15-6.16 of the Interim Report.

[79]   Interim Report Annex, para A7.5.
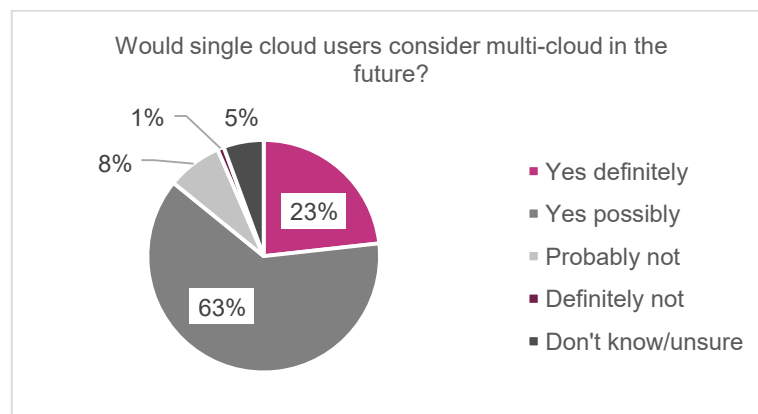
[80]   Interim Report, para 5.278.

Figure 4: Number of cloud providers used by respondents

Number of cloud providers, n=880

- 1 cloud provider only — 48%
- 2 cloud providers — 30%
- 3+ cloud providers — 22%
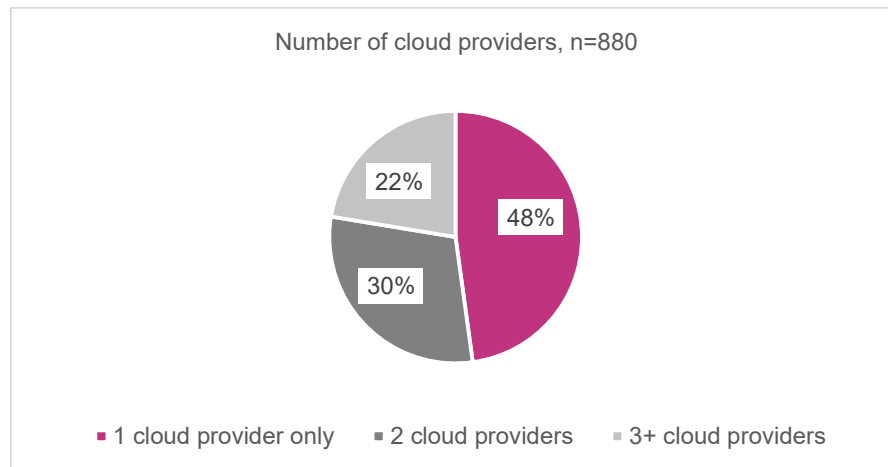
1.1.3 Ofcom's survey shows that multi-cloud is likely to continue to be common

(10) Ofcom's survey shows that customers prefer multi-homing. Almost all single-cloud users of IaaS/PaaS or those considering purchasing IaaS/PaaS would consider a multi-cloud deployment in the future. 23% said they definitely would, and another 63% said they possibly would.[81]

Figure 5: Consideration of multi-cloud use in the future

Would single cloud users consider multi-cloud in the future?

- Yes definitely — 23%
- Yes possibly — 63%
- Probably not — 8%
- Definitely not — 1%
- Don't know/unsure — 5%

(11) Multi-cloud is not restricted to IaaS. The survey suggests that multi-homing customers consist of users of SaaS and PaaS, as well as IaaS.[82] Although it is not clear to what extent a respondent multi-homes for a particular service, at a high level, the below figure shows that multi-homing is likely to be common across use cases.

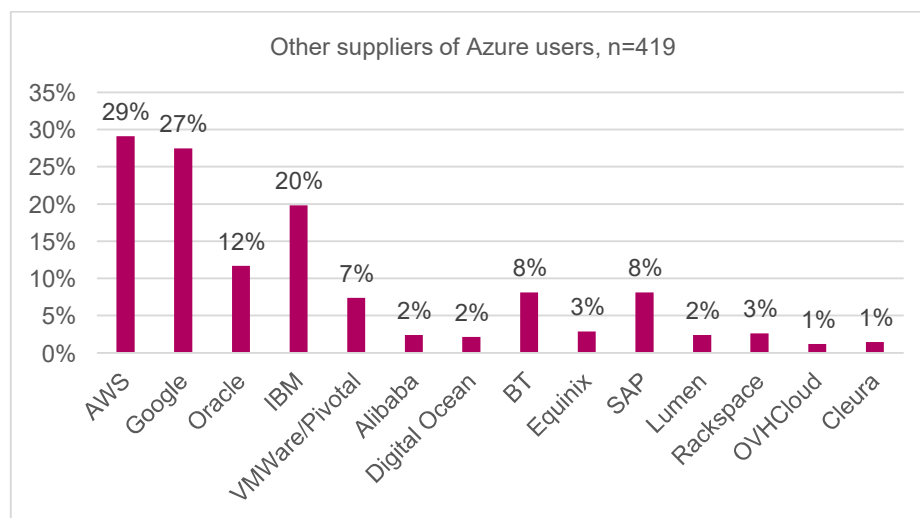---

[81] See responses to Q30 of the Ofcom data table.

[82] See responses to Q30 of the Ofcom data table.

Figure 6: Usage of SaaS, IaaS, and PaaS services by number of cloud providers used

Multi-cloud by type of service currently used

(Bar chart showing % of users of 1/2/more cloud providers that also buy each service)

	SaaS	IaaS	PaaS
1 cloud provider only	68%	74%	56%
2 cloud providers	79%	74%	63%
3+ cloud providers	88%	91%	73%
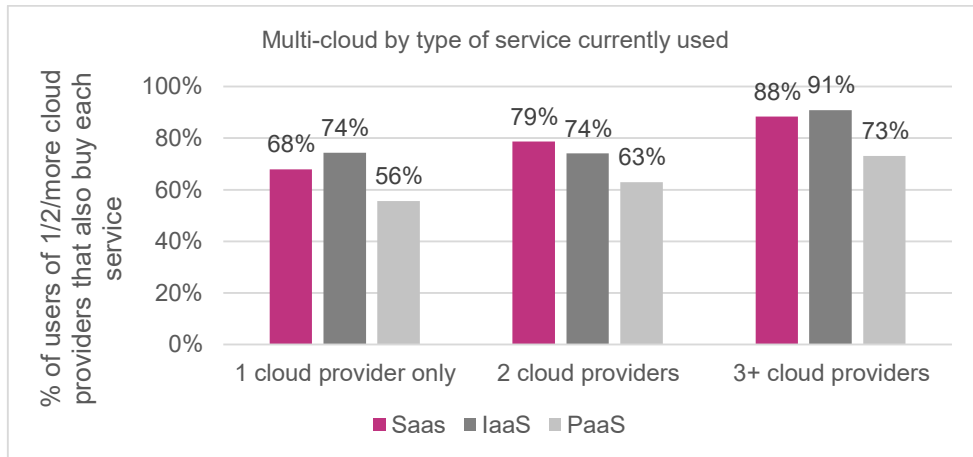
(12) Multi-cloud is also common across all suppliers.[83] While it is not possible to estimate what percentage of each supplier's customers in the sample multi-home, the below figures show that at least 30-45% of hyperscalers' customers multi-home, including with other cloud providers.[84]

Figure 7: Other cloud providers used by Azure users

Other suppliers of Azure users, n=419

Provider	%
AWS	29%
Google	27%
Oracle	12%
IBM	20%
VMWare/Pivotal	7%
Alibaba	2%
Digital Ocean	2%
BT	8%
Equinix	3%
SAP	8%
Lumen	2%
Rackspace	3%
OVHCloud	1%
Cleura	1%

---

[83] See responses to Q23 of the Ofcom data table.

[84] It offers IoT services based on fixed and mobile networks. It is not clear which BT services are included within the scope of Ofcom's survey. Sources: BT Cloud Voice | VoIP Phone System | BT Business; Insights | BT IoT.
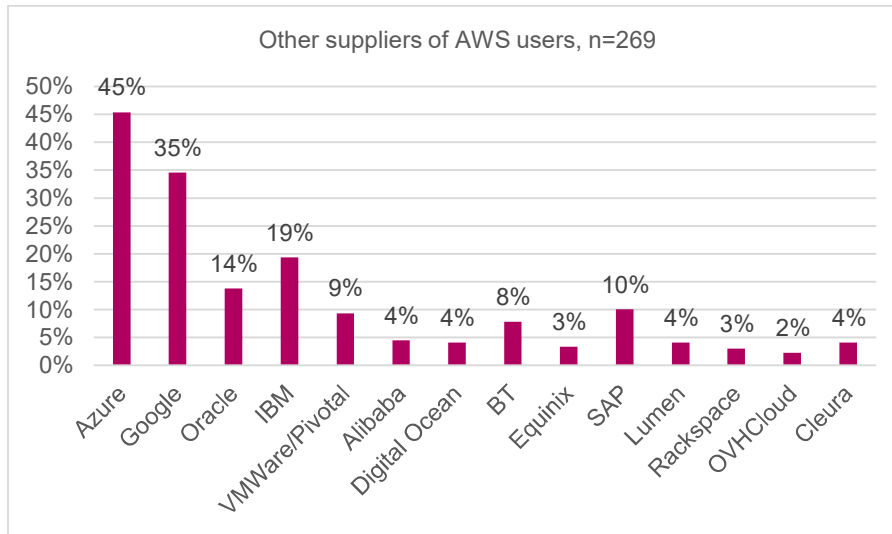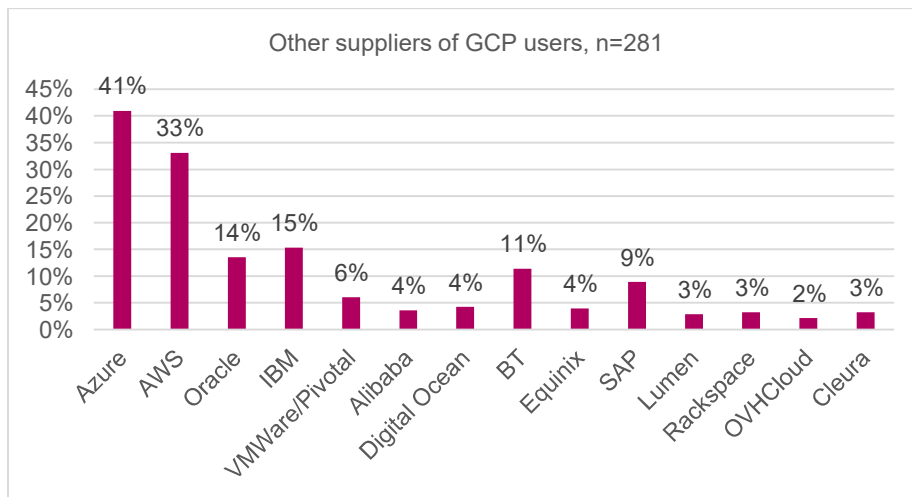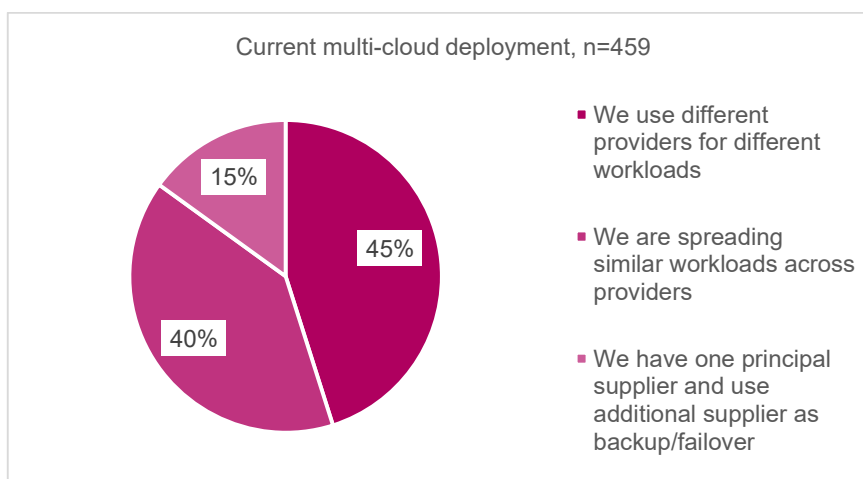
Figure 8: Other cloud providers used by AWS users

Other suppliers of AWS users, n=269

Azure 45%, Google 35%, Oracle 14%, IBM 19%, VMWare/Pivotal 9%, Alibaba 4%, Digital Ocean 4%, BT 8%, Equinix 3%, SAP 10%, Lumen 4%, Rackspace 3%, OVHCloud 2%, Cleura 4%

Figure 9: Other cloud providers used by GCP users

Other suppliers of GCP users, n=281

Azure 41%, AWS 33%, Oracle 14%, IBM 15%, VMWare/Pivotal 6%, Alibaba 4%, Digital Ocean 4%, BT 11%, Equinix 4%, SAP 9%, Lumen 3%, Rackspace 3%, OVHCloud 2%, Cleura 3%

(13)   Among multi-cloud user respondents, an equal proportion deploy similar workloads across cloud providers and use different providers for different workloads – see figure below. As Ofcom acknowledges, spreading similar workloads across providers may include both siloed and integrated workloads.[85] Integrated multi-cloud workloads may also be included within the category of respondents who use a principal and secondary supplier, whereby they may integrate some workloads. It is unclear how prevalent integrated multi-cloud is, from either the quantitative or the qualitative survey.

---

[85]   Interim Report Annex, para A7.12.

Figure 10: Current multi-cloud deployment

Current multi-cloud deployment, n=459

- 45% — We use different providers for different workloads
- 40% — We are spreading similar workloads across providers
- 15% — We have one principal supplier and use additional supplier as backup/failover

1.1.4 Price changes at renewal

(14) Ofcom found that 58% of user-provider relationships in their market research had renewed or renegotiated a contract at some point in the past.[86] Ofcom found that out of those who renewed/renegotiated, 56% experienced a price rise.[87] The average price rise was found to be 20%, the median price rise was 10%. At the same time, on the lower end of price increases, 10% of customers saw price rises of 4% or lower.[88]

(15) The way the survey results are reported is somewhat misleading for the following reasons:

(i) The relevant economic question is how the price changed for a *given* solution (consumption quantity, deployment model, etc.) upon renewal. The survey question does not clarify this for the respondents and, therefore, could potentially be picking up changes in spending rather than price changes. Neither has any timeframe been provided to survey respondents to assess the reported price changes. This also seems to be inconsistent with the generally held view in the industry that like-for-like prices always decrease over time.

(ii) The question was very open-ended while trying to collect very specific figures of price rise. The question asked was, "*You said you experienced a price rise. Please estimate by how much overall did the price rise? (% increase)*"[89]

(iii) The survey also did not provide an option for respondents to indicate whether they had experienced a price decrease.[90] It is possible that some customers may interpret the question as a price "change". Ofcom has only provided the average price rise for each supplier; no further distribution is provided to test whether any respondents indicated a price decline.[91]

---

[86] Interim Report, para 4.29.

[87] Interim Report, para 4.31. It has not been possible to replicate these figures from the Ofcom data table.

[88] Interim Report, para 4.32.

[89] See Q44 of the Ofcom data table.

[90] See Q43 of the Ofcom data table. The question was "*Please say which of the following apply to the prices and contracts you have from the providers you are using. Please select all that apply.*" Options were: (a) We found the renewal/renegotiation easy; (b) We experienced a price rise for some or all services; (c) We discovered that (some of) our preferred service(s) were not available anymore; (d) The new contract was more restrictive in terms of what we could do with the services purchased; (e) We were encouraged to buy more products than we needed.

[91] Interim Report, para 4.33.

(iv)   Ofcom highlights that the average price rise experienced for hyperscalers was 19% (Azure), 20% (AWS) and 24% (Google). The respective median price changes were 10%, 10% and 15%. However, Ofcom's study shows that the price rises were higher or similar for some non-hyperscale providers like Digital Ocean and VMWare (see figure below). This is not consistent with price rises being associated primarily with the hyperscalers' business practices.

Figure 11: Average price rise (%) experienced by respondents on renegotiation
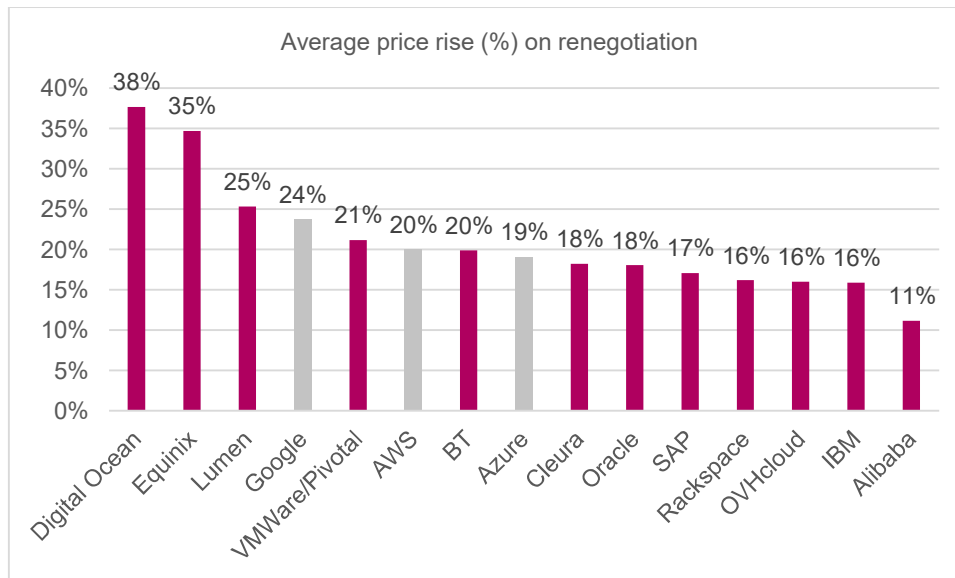
1.1.5   Existing relationships

(16)   Ofcom considers that Microsoft's position in traditional IT and SaaS makes its cloud services particularly attractive for enterprises that are already using Microsoft's products on premises. It notes that "*Some customers identify Azure as a natural choice when already using Microsoft for other services.*"[92] Ofcom's market research found that 'existing relationships for other services' was one of the top ten reasons for choosing a cloud provider.[93] Its market research also suggested that it is not considered one of the top six reasons for choosing AWS and Google, and concluded that Google and AWS benefit from relationships in other services to a lesser degree than Microsoft.[94]

(17)   The figure below shows the percentage of each supplier's users that chose a supplier due to already having a relationship with it.[95] It shows that the proportion of Microsoft's customers who chose it due to existing relationships for other services was similar to those of IBM, Digital Ocean, Equinix, Lumen and other smaller providers. It is therefore not a feature unique to Microsoft.
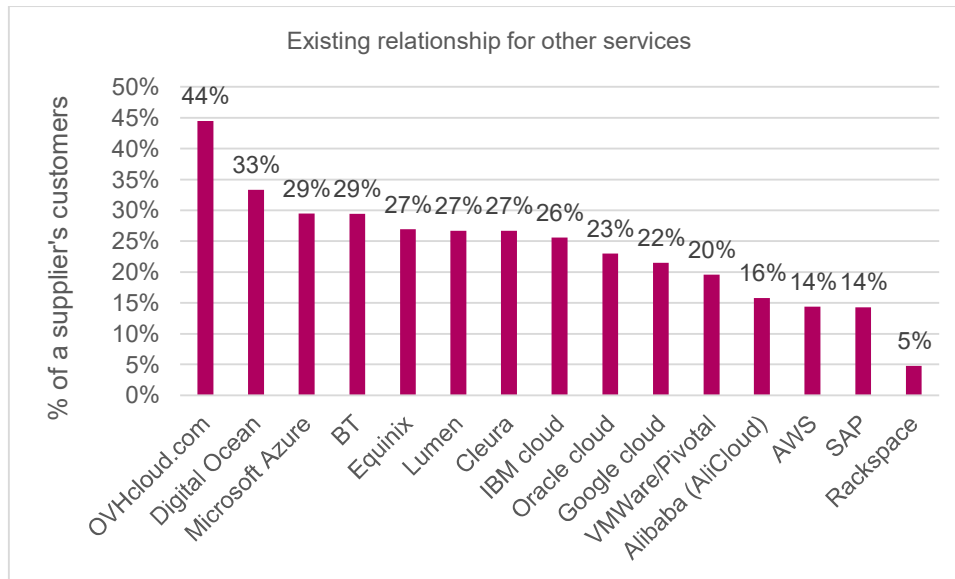
---

[92]   Interim Report, para 5.261.

[93]   Interim Report, para 5.259.

[94]   Interim Report, para 5.267.

[95]   See responses to Q25 of the Ofcom data table.

Figure 12: Share of customers who chose their cloud service provider owing to an existing relationship in other services

(18) Further, the survey suggests that cloud providers (hyperscalers and others) primarily compete on service quality, best value for money, supplier reputation, security and the number of features, rather than existing relationships.

Figure 13: Reasons for choosing cloud providers, most important and also important
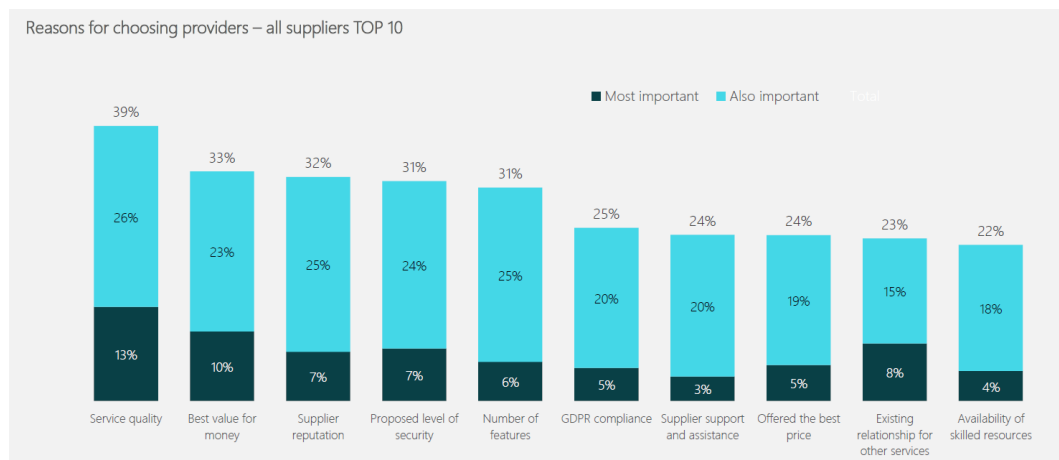

Source: Context Consulting report, page 68.

(19) For its assessment, Ofcom also asked survey respondents whether their contract or purchase of IaaS/PaaS cloud services was separate from other IT purchases, or whether these purchases were bundled with other (non-cloud) products or services. It found that, overall, about 4 in 10 "user-provider relationships" were reported to have other services bundled in, and in 56% of cases, it was a cloud-only purchase.[96] Ofcom found that purchases from AWS were more likely to be stand-alone (62%) than from Microsoft (53%) or Google (52%).[97]

---

[96] It has not been possible to replicate these figures from the Ofcom data table as it doesn't provide user-provider level data.

[97] Interim Report, para 4.20.

This is, however, a common practice across cloud suppliers who offer a multi-component array of technologies that function together as a system. Ofcom's survey confirms that customers purchase a range of non-cloud products along with IaaS/PaaS from hyperscalers as well as other cloud providers. In its survey, Ofcom asked respondents, "*Is your contract/purchase of IaaS/PaaS cloud services separate from other IT purchases, or were other products or services bundled in with them (for example, software licenses or hardware inc laptops, servers)?*".[98] As shown in the figure below, purchases from hyperscalers were just as likely to be bundled with other services compared to other cloud providers.

Figure 14: Share of customers who purchased a bundle of services, by cloud service provider
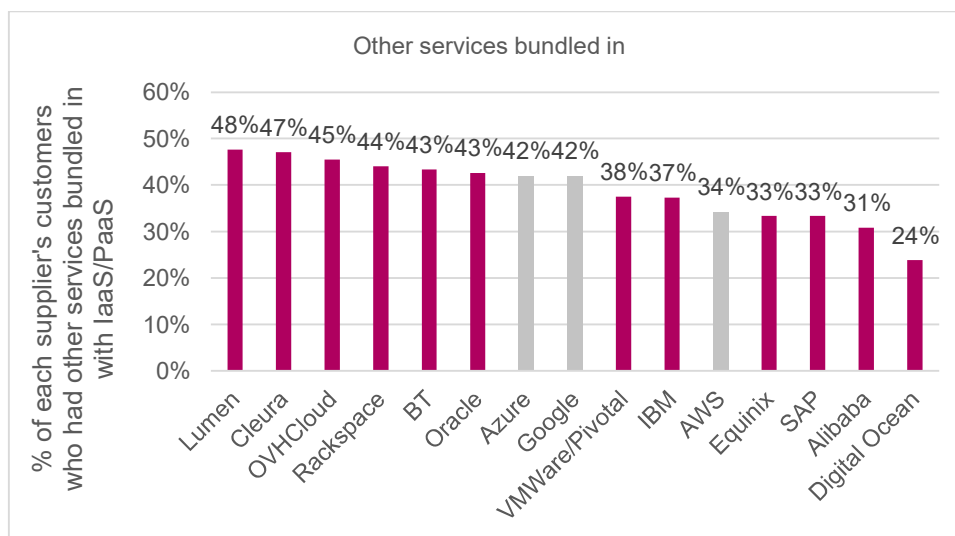
1.1.6 Prevalence of switching

(20) Ofcom acknowledges that the survey shows that about 18% of respondents said they switched IaaS/PaaS providers in the past.[99]

(21) Further, the survey shows that switching is more common among the respondents who are customers of hyperscalers than other providers and among large customers.[100] This is also at odds with claims that hyperscalers' business practices increase switching costs.
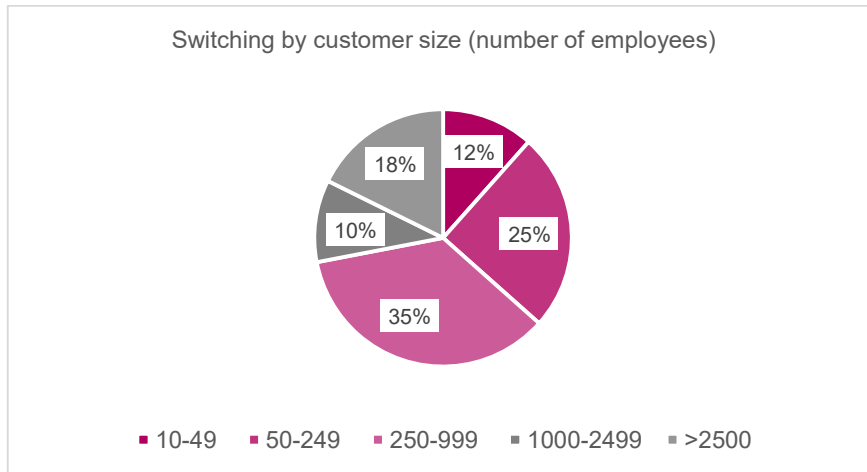
Figure 15: Switching by supplier

---
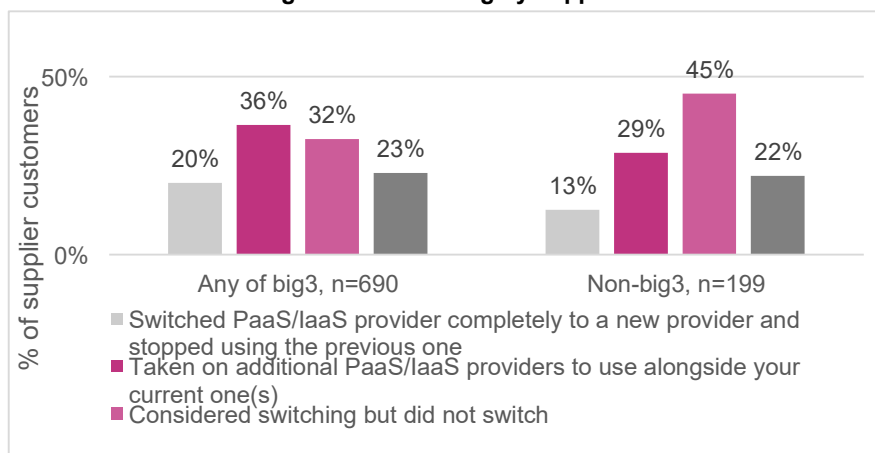
[98] See responses to Q39 of the Ofcom data table.

[99] Interim Report, para 5.41.

[100] See responses to Q47 of the Ofcom data table.

Figure 16: Switching by customer size

Switching by customer size (number of employees)

12% · 25% · 35% · 10% · 18%

- 10-49
- 50-249
- 250-999
- 1000-2499
- >2500

Annex 4
Facilitating multi-cloud solutions & interoperability

1   Tools that facilitate multi-cloud solutions

(1)   As businesses seek to avoid vendor lock-in and take advantage of 'best in class' cloud services,[101] there has been an emergence of tools that facilitate multi-cloud strategies. The following is an illustrative list of the range of tools, including open-source options.

(2)   **Cloud Management Platforms (CMPs)**: CMPs simplify the management of multi-cloud architectures.[102] They provide a unified interface for managing multiple cloud environments. Many vendors have their own CMPs, from hyperscalers such as AWS Systems Manager, Azure Arc, Google Anthos, to smaller providers and open-source options such as Scalr, Flexera Cloud Management Platform, IBM Multicloud Manager, Nutanix Beam and OpenStack.[103]

(3)   **Containers**: Containers essentially wrap the software process to make it executable in all computing environments.[104] They allow users to run applications across multiple clouds in a single virtual machine, making it easier to share resources.[105] Importantly, containerisation allows applications to be "*written once and run anywhere*." This portability speeds development and prevents cloud vendor lock-in.[106] Containers support many applications. Indeed, all Google applications, like Gmail and Google Calendar, are containerised.[107]

(4)   **Kubernetes**: Kubernetes is open-source orchestration software that provides an API to control how and where containers will run.[108] Originally designed by Google, this software specialises in grouping containers that make up an application into logical units for easy management and discovery.[109] With Kubernetes, users can manage their applications consistently across multiple cloud infrastructures. There are many providers that support Kubernetes services, such as Google Kubernetes Engine.[110]

(5)   **Cloud-to-cloud migration tools**: Cloud-to-cloud migration tools facilitate the migration of data in a multi-cloud environment.[111] They are available at all layers of the cloud architecture stack. Customers who wish to migrate between clouds can choose from tools offered by the cloud service provider they are switching to, or a range of third-party and open-source options. Examples include the Cisco Multicloud Portfolio, IBM Aspera service, NooBaa, and Rclone.

(6)   **Multi-cloud consultants**: IT teams can hire multi-cloud specialist consulting firms to help with managing multi-cloud solutions or migrating multi-cloud data. The biggest systems integrators have relationships with all the major cloud providers and can provide this kind of service. For example, BMC provides multi-cloud consulting services that can plan your

---

[101] What Is multicloud? Definition and benefits | Google Cloud.
[102] 8 Best Multi Cloud Management Platforms & Tools (bunnyshell.com).
[103] Best Cloud Management (CM) Platforms 2023 | SoftwareReviews.
[104] https://www.ridge.co/blog/what-are-containers/.
[105] What are Containers and How Do They Work? | Ridge Cloud.
[106] Containerization Explained | IBM.
[107] What are Containers and How Do They Work? | Ridge Cloud.
[108] Kubernetes vs Docker | Microsoft Azure.
[109] Containers in the Multi-Cloud – BMC Software | Blogs.
[110] Google Kubernetes Engine (GKE) | Google Cloud.
[111] Cloud-to-cloud migration tools offer comprehensive services | TechTarget.

migration, identify your existing assets, forecast cloud service costs and ensure that your migration remains secure and compliant.[112]

(7) **Multi-cloud ISVs**: ISVs offer solutions that are designed to work across multiple cloud providers. This makes it easier for users to adopt multi-cloud solutions while accessing the best services and features offered by each provider. For example, Snowflake.ai is a fully managed data warehouse that centralises data from various clouds and on premises sources while maintaining a consistent interface no matter where the data originated.[113]

2   Interoperability of Microsoft Azure

(8) Microsoft has developed a range of solutions to facilitate the interoperability of its cloud products with those of other cloud providers. Below is a (non-exhaustive) list of examples.

(9) **Open APIs**: Microsoft publishes volumes of information on its APIs and technical documentation on its website and GitHub to allow third-party services to communicate with its applications. This facilitates the ability of developers to build Azure-compatible systems, which in turn enhances the attractiveness of the platform; in other words, Microsoft's incentives are to be open.

(10) **Proprietary APIs can be called**: Even Microsoft's proprietary first-party APIs can be 'called' by another a service operating on another cloud. For example, a customer can store their data in AWS and send it to Azure for processing with relatively little effort (i.e. by changing one or two lines of code to connect with the endpoint of the Azure API.)

(11) **Azure PaaS compatibility**: Microsoft provides a range of first-party solutions that operate across clouds. For example, Defender for Cloud provides security assessments across Azure, AWS, and Google Cloud. SQL Server, Microsoft's leading proprietary database offering, is also available both on premises and on multiple clouds (Azure, AWS, GCP, OCI).

(12) **Azure Arc cloud-native services**: "Arc-enabled" Azure services are lightweight cloud-native services that can be run across Azure, other clouds, and on premises, without even 'calling into' Azure. Though Arc-enablement can facilitate a customer's multi-cloud strategy, it only works for services that do not require a complex architecture or substantial compute power to run, so customisation of the cloud environment is limited.

(13) **Linux support**: Azure allows Linux VMs to run on its cloud infrastructure, with only a small number of customers choosing to use Windows Server. This facilitates the deployment of open-source middleware and applications.

(14) **Partnership with Oracle**: Microsoft's partnership with Oracle Cloud Infrastructure improves their cross-cloud capability such that customers are able to run the entire solution stack in a single cloud, while leveraging the most advanced services from each provider.[114,115] For example, customers can link their Oracle and Microsoft clouds using the Azure's ExpressRoute and Oracle's FastConnect, which provides extremely low-latency, private connections between workloads distributed across OCI and Azure. They can also develop integrated multi-cloud architectures.[116]

---

[112] Cloud-to-cloud migration tools offer comprehensive services | TechTarget.
[113] Snowflake Architecture and It's Fundamental Concepts (projectpro.io).
[114] Connect Azure ExpressRoute with Oracle Cloud Infrastructure - Azure Virtual Machines | Microsoft Learn.
[115] Learn About Connecting Oracle Cloud with Microsoft Azure.
[116] Integrate Microsoft Azure with Oracle Cloud Infrastructure - Azure Virtual Machines | Microsoft Learn.

Annex 5
Multi-cloud customer examples

(1) The following table show examples of UK and European customers that were identifiable as using more than one cloud provider based on the customer webpages of cloud providers. These customer crossover lists were assembled based on the full customer story or case study lists from the web pages for AWS,[117] Azure,[118] GCP[119] and OCI.[120]

(2) The research shows that while multi-cloud is common among large enterprises, many small and medium-sized [✂]. [✂] uses Microsoft's Azure cloud for scalability, global deployment, compliance, fault tolerance and capacity. They also use Azure Kubernetes Services to deploy and manage containerised applications.[121] [✂] also uses GCP as a parallel provider to support a multi-cloud infrastructure for access to high availability products.[122]

---

[117] Customer Success Stories: Case Studies, Videos, Podcasts, Innovator stories (amazon.com).

[118] Microsoft Customer Stories Search (Microsoft customer stories were filtered for "Cloud Platform").

[119] Customers | Google Cloud.

[120] Vodafone (oracle.com).

[121] [✂]

[122] [✂]

Table *1*: European and UK companies with customer stories featured by more than one provider[123]

AWS & Microsoft	AWS & Google Cloud	AWS & Oracle	Microsoft & Oracle	Microsoft & Google Cloud	Oracle & Google Cloud
[✂]	AI21 Labs	Deutsche Bahn	[✂]	[✂]	Orange
[✂]	BBVA	Siemens	[✂]	[✂]	
[✂]	Brainly		[✂]	[✂]	
[✂]	Deliveroo*		[✂]	[✂]	
[✂]	Delivery Hero		[✂]	[✂]	
[✂]	Green Flag*		[✂]	[✂]	
[✂]	HSBC*			[✂]	
[✂]	ITV*			[✂]	
[✂]	Landbay*			[✂]	
[✂]	MAPFRE			[✂]	
[✂]	Monzo*			[✂]	
[✂]	News UK*			[✂]	
[✂]	Northumberland City Council*				
[✂]	Scout24				
[✂]	Sky Italia				
[✂]	Starling Bank*				
[✂]	Trax				
[✂]	TUI				
[✂]	Wix				
[✂]	Zalando				
[✂]					

---

[123] UK-based customer, where identified by the cloud provider or company headquarters.

Annex 6
Pricing evidence and lock-in

Annex 7
Market Shares

1   The economic assessment of competition cannot be wholly UK-centric

(3)   UK-specific metrics used by Ofcom (such as market shares, or survey responses) should be considered in the larger context of a fundamentally *global* market, a small proportion of which operates in the UK.

(4)   Ofcom estimates UK cloud spending for IaaS and PaaS services of up to £5.0 billion in 2021[124] which compares to a Gartner estimate of global cloud customer spending on IaaS and PaaS services of £229.32 billion in 2022.[125] Assuming an annual growth from 2021 to 2022 in UK IaaS and PaaS revenues of 30%,[126] this would mean that the UK accounted for around 2.18% of the global cloud market in IaaS and PaaS services for 2022.

(5)   Public cloud fundamentally involves cross-border networks and services; most R&D innovation is physically centred in the US, and competition between suppliers on innovation roll-outs, service quality, and other important non-price dimensions, is essentially global. Pricing is also generally set globally (e.g. Microsoft sets global prices in USD which is then converted to GBP and other currencies). Even if they pay in sterling, the vast majority of UK customers pay globally-set prices.

(6)   To be sure, there are in the UK, as is the case worldwide, a subset of Azure's largest customers, typically the largest multinational/FTSE100 enterprises, that benefit from discounts individually negotiated primarily in the UK. This is core to how Microsoft competes on price against AWS, GCP and others for its largest and most sophisticated cloud customers. However, there are no "price discrimination" or other idiosyncratic UK features that apply to this UK subset of customers that make the analysis of them discrete, for purposes of grouping individual customers into categories for a competition assessment, from large enterprise customers in North America, Europe and globally.

2   Global and UK IaaS market shares

(1)   Historically, Amazon Web Services (AWS) has held a strong position in the global IaaS market, with a share of revenue topping 50% as recently as 2017 (see below figure). While AWS has been challenged by competitors such as Microsoft, Alibaba, and Google, it was still almost double the size of Microsoft in 2021.

---

[124] Consultation: Proposal to make a market investigation reference into the supply of public cloud infrastructure services in the UK (ofcom.org.uk), para 3.13.

[125] Gartner Press Release, Gartner Forecasts Worldwide Public Cloud End-User Spending to Reach Nearly $600 Billion in 2023, 31 October 2021 accessed 3 May 2023. Gartner's 2022 figure for IaaS, PaaS and BPaaS has been converted into GBP using the GBP-USD average exchange rate for 2022: Average for the year to 31 December 2022 – GOV.UK (publishing.service.gov.uk).

[126] Consultation: Proposal to make a market investigation reference into the supply of public cloud infrastructure services in the UK (ofcom.org.uk), para 3.13.

Figure 1: Vendor revenue share of the IaaS market worldwide from 2015 to 2021
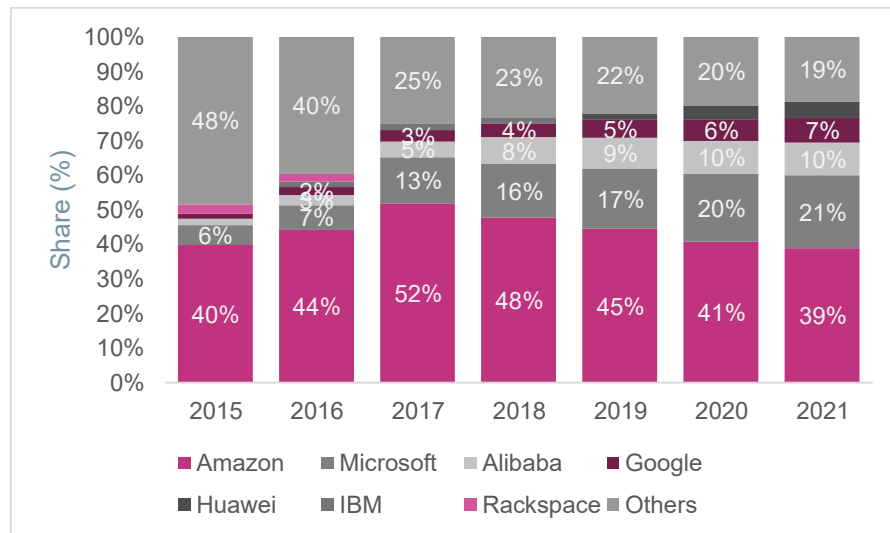
(2)   The UK IaaS segment is consistent with the global picture, with Amazon holding a strong market position. As of 2021, Amazon had a 44% share of UK IaaS revenue, with Microsoft and Google accounting for 25% and 8%, respectively.

Figure 2: Revenues from IaaS market UK from 2020 to 2021, by vendor
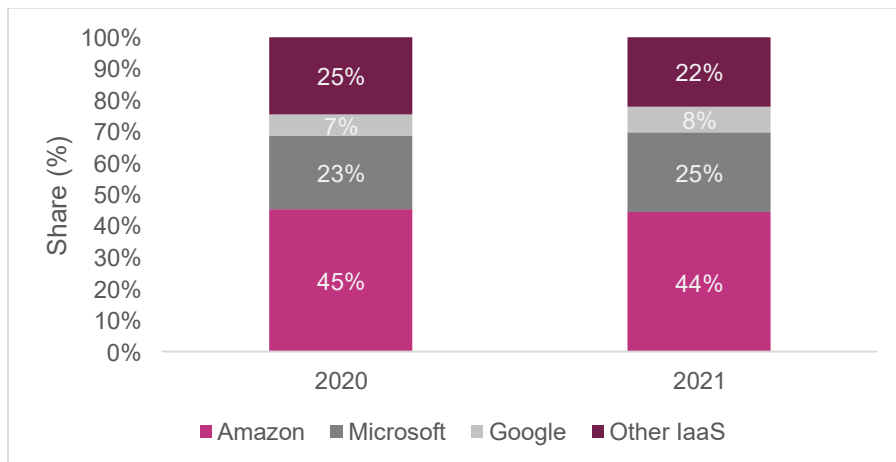

Source: Gartner: Market Share: IT Services.

(3)   While it is not possible to replicate Ofcom's share estimates due to use of information from other providers, the methodology described in the Interim Report may be overestimating Microsoft's IaaS shares.[127] The reason is that in response to Ofcom's RFIs, some providers attribute IaaS fees triggered by PaaS usage to the PaaS service (for example AWS and Google), while others (including MSFT) attribute such IaaS revenues to IaaS services. As a result, Ofcom's estimates of Microsoft's IaaS-only shares are likely to be over-estimated.

---

[127] Interim Report Annex, Table A5.2.

3 Global and UK PaaS shares

(4) There has been considerable growth in global PaaS spend in recent years – increasing from four billion dollars in 2015 to over one hundred billion dollars in 2022. UK PaaS revenue has also experienced considerable growth, with revenue projected to grow at a CAGR of 11% from 2022-2027.[128]

Figure 3: Public cloud PaaS end-user spending worldwide from 2015 to 2023
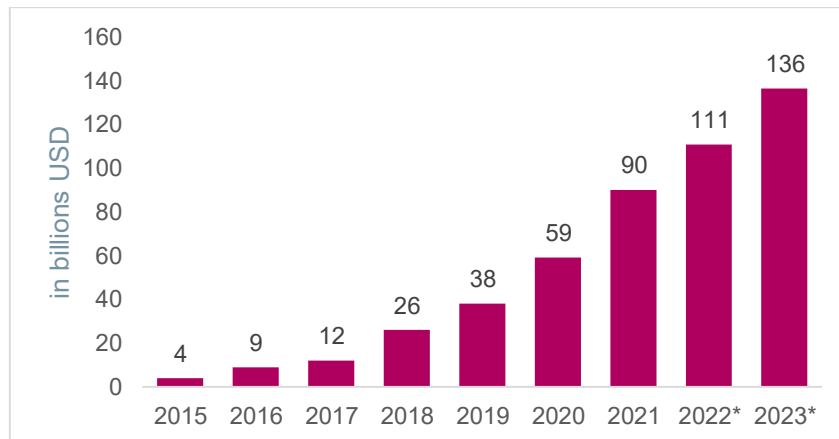
Source: Public cloud PaaS market size worldwide 2024 | Statista.

(5) Like IaaS, AWS also leads in terms of global revenue share for PaaS services, with a 25% share in 2020 and 2021. Microsoft, Salesforce, and Google follow as the next largest PaaS vendors, with market shares of 14%, 9% and 7%, respectively, as of 2021.

Figure 4: Vendor share of the PaaS market worldwide from 2020 to 2021
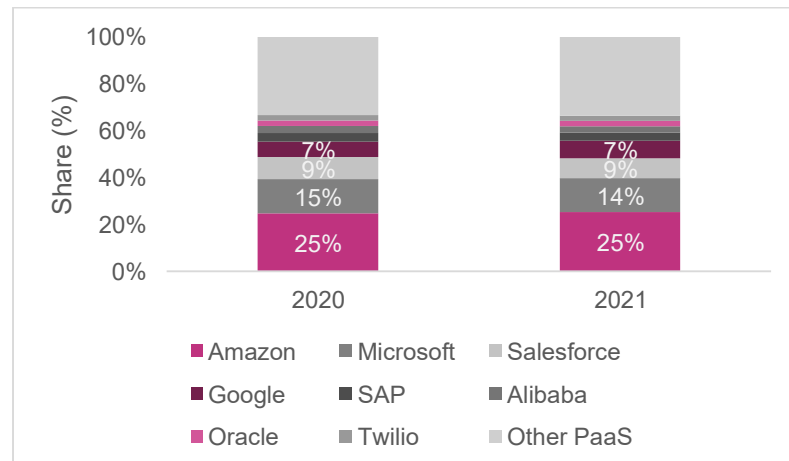
Source: Gartner: Market Share: Platform as a Service.