

BROADCOM'S ANTICIPATED ACQUISITION OF VMWARE

RESPONSE TO THE CMA'S ISSUES STATEMENT

**CONFIDENTIAL
CONTAINS BUSINESS SECRETS**

10 May 2023

CLEARY GOTTlieb STEEN & HAMILTON LLP

CONFIDENTIALITY

This Response (including annexes and any supporting documents) contains business secrets. Publication or disclosure of such information would significantly harm the legitimate business interests of Broadcom and/or VMware. Accordingly, no such information should be published or disclosed to any third party without the prior written consent of the relevant party. Should the CMA consider it necessary to disclose any part of this Response, the parties request that they first be given the opportunity to make further confidentiality representations to the CMA before any final decision is taken under Part 9 of the Enterprise Act 2002.

In addition, certain information provided in this response is confidential as between the parties and must not be disclosed to the other party.

Broadcom/VMware Merger Inquiry
Response to the CMA's Issues Statement

This Response is submitted to the CMA in response to its issues statement of 21 April 2023 (the **Issues Statement**) in relation to Broadcom's anticipated acquisition of VMware (the **Transaction**). It is submitted on behalf of Broadcom and VMware (together, the **Parties**).

1. EXECUTIVE SUMMARY

- 1.1 Most of VMware's and Broadcom's respective operations have no practical or competitive interaction. Of Broadcom's 16,000 products and services, only *three* hardware devices – Ethernet network interface cards (**NICs**), storage adapters, and Fibre Channel host bus adapters (**FC HBAs**) – have any interaction with VMware's server virtualisation software (**vSphere**).
- 1.2 A fourth type of hardware device – Fibre Channel (**FC**) switches – does not interact with vSphere. Instead, it is managed by software that can collect information from VMware using public APIs.
- 1.3 The Issues Statement suggests that the Transaction may substantially lessen competition in relation to each of these four hardware devices. For the reasons set out in this Response, this is a highly implausible outcome.
- 1.4 VMware sells vSphere to enterprise customers, which enables them to increase the efficiency and manageability of the computer servers installed in datacentres. Broadcom supplies three hardware components that can interoperate with vSphere: NICs, storage adapters, and FC HBAs. These components also interoperate with a wide range of other hardware and software used in a datacentre. Indeed, Broadcom supplies the components to server manufacturers (**OEMs**), which combine them with hundreds – if not thousands – of other components.
- 1.5 The three hardware components are input/output (**I/O**) devices that are physically attached to the motherboard of a server. I/O devices enable data transfer from the server to other datacentre infrastructure, namely to: (1) other servers (for NICs); (2) local storage devices (for storage adapters); and (3) storage devices on an FC network (for FC HBAs). Server virtualisation software, like vSphere, runs “on top of” a physical server. Since I/O devices allow the transfer of data to/from the server, they must be able to interact – or interoperate – with the virtualisation software. These devices interact using device drivers, which are straightforward pieces of software code that provide a mapping between the I/O device and the application programming interface (**API**) of the virtualisation software.
- 1.6 Broadcom also sells FC switches to storage OEMs. Unlike I/O devices that are installed in servers, FC switches are standalone hardware devices that are located elsewhere in the datacentre. FC switches direct data transfer to/from devices on an FC network, like a switchboard operator. FC switches communicate with servers using the industry standard FC protocol. In contrast to I/O devices, FC switches do not need to interoperate with server virtualisation software.
- 1.7 The Issues Statement advances two theories of harm:

- (a) The first (**ToH1**) has two parts. Part one (**ToH1(a)**) considers whether, after the Transaction, Broadcom (or the **Combined Entity**) may foreclose rival I/O device vendors by degrading interoperability between vSphere and rivals' device drivers. Part two (**ToH1(b)**) considers whether the Combined Entity may foreclose rival FC switch vendors by degrading interoperability between VMware APIs and the management software used by rivals' FC switches.
- (b) The second (**ToH2**) posits that competition in the supply of NICs, storage adapters, and FC HBAs may be harmed by the flow of commercially sensitive information from Broadcom's hardware rivals to the Combined Entity.
- 1.8 This Response explains why neither theory of harm can be sustained. In doing so, it summarises the explanations and evidence that the Parties have provided, and continue to provide to the CMA, in more detail. The evidence shows that the Transaction will not give rise to a substantial lessening of competition (SLC).
- 1.9 ToH1(a) cannot be substantiated in light of the evidence relating to the Parties' past conduct, their future strategies, their economic incentives, and other important market realities:
- **Interoperability is critical for all virtualisation software.** Hardware neutrality is essential for virtualisation software to provide its two core values: (1) making more efficient use of physical server resources, and (2) enabling enterprises uniformly to manage the diverse range of hardware and software components in datacentres, by creating an abstraction layer above the heterogeneous hardware. Neither can be achieved without hardware neutrality, as VMware, customers, rivals, partners, industry reports, and regulators have recognised.
 - **Degrading interoperability is contrary to Broadcom's strategy.** Seeking to advantage aspects of its hardware business by degrading interoperability of rivals' hardware has never been part of Broadcom's strategy. Broadcom is purchasing VMware for \$61 billion and will invest billions more per year in order to enhance and expand penetration of VMware's products. That objective would be fatally undermined if Broadcom were artificially to restrict the hardware on which vSphere could operate and thereby limit its pool of customers. Such a strategy would make no commercial sense and would jeopardise Broadcom's investment.
 - **OEMs would retaliate against any attempt to degrade interoperability.** Broad interoperability is a founding principle of the datacentre. No one in a datacentre environment has ever attempted to degrade interoperability. Even if the Combined Entity were to attempt to preference its own hardware products by degrading interoperability with rivals' products, OEMs would take measures to counter it. OEMs account for c. [X]% of Broadcom's I/O device revenues (c. \$[X] in 2021) and c. \$[X] of annual VMware bookings.¹ OEMs control relationships with enterprise customers and can influence hardware options that are made available to them. They demand interoperability so that they can

¹ Final Merger Notice dated 19 January 2023 (FMN), paragraphs 15.9 and 20.128.

multi-source hardware from their vendors and offer their customers choice. Reducing the number and types of servers on which vSphere can run also reduces the number and type of servers that OEMs can sell. Interoperability is therefore essential for OEMs' competitiveness, commercial success, and supply chain resilience. OEMs would punish the Combined Entity swiftly and harshly if it tried to degrade interoperability. So would enterprise customers and rivals.

- **Enterprises have many credible options for their workloads.** The premise of ToH1(a) is that vSphere customers have no real alternatives. This is not consistent with [X] and is not supported by the evidence available:
 - ***On-premises rivals are close competitors***, including Microsoft Hyper-V, IBM/Red Hat, and Nutanix, as well as the private cloud solutions introduced by cloud service providers (CSPs), such as AWS Outposts, Azure Stack, and Google Anthos. The significant constraint posed by these competitors is substantiated by VMware's internal documents, third-party evidence, and market shares.
 - ***CSPs are a strong and growing constraint***, in particular Amazon, Microsoft, and Google, but also Oracle, IBM and others. This is also supported by the Ofcom Cloud Market Study, the CMA's market investigation, industry reports, customer surveys, and CSPs' revenue growth, all of which show that the public cloud is a preferred choice for all types of enterprise workloads. VMware's internal documents indicate that public cloud represents [X] of VMware's churn.
 - ***Containerisation is increasingly used as an alternative*** to virtualisation software, as confirmed by the Ofcom Cloud Market Study, the CMA's market investigation, and industry reports of significant expected growth. Containerisation also facilitates switching among virtualisation providers.
- **The Combined Entity could not degrade interoperability for I/O devices in existing servers**, as the CMA's Phase 1 decision of 22 March 2023 (the **Phase 1 Decision**) confirms. To do so would disrupt the critical business applications on which enterprises depend. Nor could the Combined Entity gain sales from somehow *lessening* interoperability for existing servers. Although the Phase 1 Decision describes lessening interoperability as a theoretical concept, it does not explain what degradation is envisaged. I/O devices are extremely mature products; drivers for existing devices are already certified, such that degrading interoperability essentially means withholding incremental bug fixes and patches for security vulnerabilities. Doing so would not only damage VMware's reputation but also cause disruption and impose an enormous burden on VMware's customers. Enterprises would be faced with a choice of switching away from vSphere or shutting down their (potentially thousands of) servers and manually replacing their I/O devices that are already physically installed inside servers, which would take tens of thousands of hours, and risk damaging their servers. It would be wholly implausible.
- **The Combined Entity could not degrade interoperability for new sales of existing I/O devices.** VMware has no ability to distinguish between (1) existing

devices in existing servers and (2) new sales of existing devices (in new servers). Once VMware certifies a device driver, the **same** driver supports all instances of a device installed in existing servers over the past 10 years or more – covering multiple product generations – and new sales of the existing I/O device generation. Any updates to that driver similarly impact all these devices. The Combined Entity could not degrade the driver for new sales without also impacting the installed base of existing devices.

- **The Combined Entity could not degrade interoperability with new I/O devices that are certified by equivalency.** Almost all I/O devices are certified via a well-established, tick-box process by which VMware certifies a device as functionally equivalent to one that has already been certified. [90-100]% of Broadcom’s I/O devices are certified through equivalency. VMware is not alone in this – it is common industry practice, and a product of necessity. The Combined Entity could not refuse equivalency without also announcing its radical departure from hardware neutrality and thereby destroying VMware’s core value and reputation.
- **The Combined Entity could not in practice degrade interoperability even for new I/O devices that require full certification.** VMware has never refused certification for a driver that meets its published specifications. If the Combined Entity were to do so, it would reveal VMware’s deviation from hardware neutrality. The Combined Entity cannot lessen interoperability because there is minimal engagement with VMware given the automated nature of the certification process and published test kits. Hardware vendors can usually fix any issues that arise themselves, as they do for Linux drivers.
- **The Combined Entity has no reason to degrade interoperability for I/O devices,** which are many times less profitable than vSphere. Broadcom’s I/O devices average c. \$[X] incremental margin per server which compares to vSphere’s \$[X]. The potential gains to be realised from selling more I/O devices are therefore minimal relative to the potentially huge losses from losing vSphere licences. Indeed, the Combined Entity could, at most, gain c. \$[X] million in year one and c. \$[X] million per year in the longer term. The maximum gains from a partial foreclosure strategy are c. \$[X] per year. To put that in context, Broadcom wants to grow VMware revenues by \$[X] over three years through increased product deployment and utilisation. And yet, based on a survey of ~1,200 VMware customers, the Combined Entity would likely stand to lose at least c. \$[X] million in year one and at least c. \$[X] billion per year in the longer term from a foreclosure strategy. Broadcom cannot therefore profitably degrade interoperability.
- **Foreclosure attempts would have no anticompetitive effect on rival I/O device vendors.** Any putative foreclosure strategy (whether total or partial) would impact, at most, the ~[10-20]% of servers that run vSphere. The Combined Entity cannot foreclose non-VMware servers. Server OEMs would also ensure that foreclosure has no market-wide effects, given the importance of multi-sourcing to OEMs’ own competitiveness and commercial success. Because only [10-20]% of servers run vSphere, OEMs could neutralise any foreclosure attempt by increasing rivals’ NIC, storage adapter, and FC HBA

sales in the [90-100]% of servers that would be unaffected by any foreclosure. OEMs can do so at little to no cost.

1.10 As for ToH1(b), it misunderstands the (lack of any) relationship between vSphere and FC switches:

- **Neither FC switches nor their management software interoperate with vSphere.** There is no interaction at all between FC switches and vSphere. FC switch management software does offer FC switch customers the option to collect some basic inventory information from vSphere's management component, vCenter, using a public API. But that information is not used for the management of FC switches and there is no interoperation. Moreover, the information that can be collected from the API is available elsewhere, and most customers choose not to collect it in this way in practice. ToH1(b) is therefore implausible.
- **The Combined Entity has no reason to degrade interoperability with its public vCenter API.** Doing so would not impact rival FC switch vendors, but it would harm VMware's enterprise customers that use the same API to manage their IT environments. It would also harm OEMs that use the API for their management and monitoring software. It is a foreclosure strategy with substantial downside and no upside.

1.11 Finally, ToH2 significantly overstates the sensitivity of information shared with VMware for driver development and certification with respect to NICs, storage adapters, and FC HBAs, and wrongly dismisses industry-standard protections that apply across all vendors who routinely share information that is needed to ensure interoperability:

- **No competitively sensitive information is required for driver development and certification.** The information in question is only that which is needed to ensure that a driver can communicate with vSphere. In essence, all that is required is the driver source code, which provides the simple software mapping to vSphere's APIs. The driver source code does not impact the [§<] of the I/O device. That is determined by the device's [§<]. The driver source code reveals nothing about the chip or firmware, neither of which is ever shared with VMware.
- **Driver source code is mostly public before it is shared with VMware.** Hardware vendors develop drivers for open-source Linux before VMware. As a result, the source code is publicly available. 90% of Broadcom's VMware source code uses the Linux driver code. The remaining 10% reflects VMware's API specifications, which are also public.
- **NICs, storage adapters, and FC HBAs are mature products** that require limited to no detailed technical interaction with VMware. Indeed, over the past five years, vendors have made very few technical requests to VMware in relation to these products.
- **Product roadmaps are rarely shared with VMware and are not required for driver development.** In the past 10 years for NICs, storage adapters, and

FC HBAs, Broadcom has on one occasion provided a single document to VMware that could be described as a product roadmap. Over that period, Broadcom has certified thousands of devices with VMware. Product roadmaps are not shared because they are not necessary to enable the driver to communicate with vSphere. The only document of this type that Broadcom has shared with VMware gave only a high-level update on an industry-wide development that impacted Broadcom's FC HBAs, which was already widely known.

- **Broadcom and VMware must maintain their reputations as trusted partners.** For the Combined Entity to succeed, its products must continue to interoperate with a wide range of hardware and software in the datacentre. This requires other hardware vendors to trust the Combined Entity with any commercially sensitive information that may be shared to enable interoperability.
- **VMware assiduously protected all partner information when owned by hardware companies.** VMware was previously owned by two leading hardware companies for the vast majority of its corporate existence (EMC and Dell). Despite this, information on rival hardware vendors (e.g., Pure Storage, Hitachi, and IBM) was not shared with the parent hardware companies. The information was protected by NDAs.
- **Broadcom has [X] agreed [X].**

1.12 Far from giving rise to an SLC, the Transaction is in fact pro-competitive. Broadcom plans to invest \$[X] more per year to improve VMware's products, expand deployments and utilisation, and thereby increase sales by \$[X] over three years. To achieve this, Broadcom must enable VMware to compete more effectively with CSPs for enterprise workloads and stem the loss of workloads from VMware to the public cloud. Degrading the quality of VMware's offering would be antithetical to Broadcom's plans.

1.13 The remainder of this Response is structured as follows. **Section 2** describes Broadcom's rationale for the Transaction and explains that it will *increase* competition with the giant CSPs for enterprise workloads. **Section 3** explains that interoperability is a business imperative for *all* hardware and software vendors that operate in datacentres (including server OEMs), and why. **Section 4** shows that Broadcom could not and would not degrade interoperability between vSphere and rival hardware devices. **Section 5** explains that information shared by NIC, storage adapter, and FC HBA vendors with VMware does not and could not impact competition. **Section 6** concludes with summary remarks.

2. THE TRANSACTION WILL INCREASE COMPETITION WITH CSPs FOR ENTERPRISE WORKLOADS

2.1 Broadcom's strategy is to grow through acquisitions and organic innovation. It invests in businesses that have strong technology and that are active in established and sustainable markets. VMware fits this model.

2.2 Broadcom has a strong track record of acquiring businesses and delivering growth through product innovation. It has adopted a business model of operational excellence, efficient centralised services, focus on core competencies, and a strong commitment to R&D. It applies this model to all of its business units. Broadcom’s internal documents show that Broadcom’s rationale for the Transaction is to [redacted] and to fulfil customer desire for [redacted].² To achieve this objective, Broadcom must ensure that it invests in VMware and develops its products in a way that allows it to (1) win more enterprise workloads in a large and growing cloud market, (2) compete more effectively against the CSPs, and (3) compete to retain workloads that would otherwise move to CSPs.

VMware has the technology to compete with the CSPs but is unable to execute its growth strategy

2.3 VMware has a strong core technology and operates in a proven and growing market for enterprise workloads. But VMware’s own internal documents describe its market position as a “[redacted]” [redacted].³ Broadcom believes it can provide VMware with the scale and capabilities to reverse this trend. Broadcom plans to win more enterprise workloads by [redacted]. With time and investment, Broadcom’s objective is to enable enterprises more easily to deploy workloads across cloud environments (both private and public) and to move workloads among those environments.⁴ In short, Broadcom intends to execute Phases 2 and 3 of VMware’s growth plan (as set out in **Figure 1** below) that VMware has failed to do and is unable to achieve alone.

Figure 1: Broadcom plans to execute VMware’s growth strategy

[redacted]

2.4 **VMware is losing enterprise workloads to the public cloud.** The public cloud is a preferred IT environment for many enterprises to deploy their business applications, run workloads, and store data. As explained in Section 4 below, enterprises are increasingly moving workloads to the public cloud. They are doing so for many reasons, including for ease of use, to increase flexibility and agility, to improve security, and to strengthen resilience.⁵ The CSPs – in particular Amazon, Microsoft, and Google – also enable enterprises to outsource their computing and thereby reduce or eliminate the need to own or lease physical hardware (e.g., servers, switches, and storage devices). CSPs have their own large datacentres, housing hundreds of thousands of servers, and provide enterprises with full compute, storage, and network virtualization offerings. CSPs use proprietary software to virtualize the entirety of their datacentres. And these “software-defined” datacentres are more efficient, more flexible, easier to manage, and require less hardware than traditional on-premises datacentres.

2.5 **VMware’s software can enable enterprises to build private clouds.** Like CSPs, VMware has the full suite of software necessary to virtualise the entire datacentre (its

² See e.g., Annex Q9(BM)-001 and Annex Q9(BM)-004 submitted with the FMN.

³ See VMware internal document, [redacted]. See **Annex ISR-001**.

⁴ See Ofcom, Cloud Services Market Study, Interim report, 5 April 2023 (the **Ofcom Cloud Market Study**), Section 5, finding that “[a] lack of interoperability and portability can hinder customers’ ability to switch and multi-cloud”.

⁵ See Ofcom Cloud Market Study, Figure 3.3.

Private Cloud software).⁶ VMware enables enterprises to build and operate their own software-defined datacentres on-premises or hosted by a third party off-premises (each of which constitutes a private cloud).⁷ VMware licenses software to enterprises rather than adopting the CSPs' outsourcing model. VMware's Private Cloud software is intended to provide enterprises with the equivalent flexibility and ease of use as CSPs' offerings.

- 2.6 **VMware has been unable to persuade its customers to license and deploy its Private Cloud software.** Of VMware's top [X] customers, only [20-30]% (~[X] customers) have purchased VMware's Cloud Foundation (VCF) software. This VCF Private Cloud software seeks to enable VMware customers to virtualise their entire datacentres (using vSphere to virtualise servers, vSAN to virtualise storage, and NSX to virtualise the network). Today, only [5-10]% of VMware's top [X] customers (~[X] customers) [X] VCF.⁸ This low uptake is despite the demand generally among enterprises to build private clouds as an alternative to public cloud. Indeed, as described below, the CSPs are moving into private cloud, which is a segment that is projected to grow substantially as part of the broader cloud market.⁹ Broadcom has learned from its due diligence that this is because VMware's Private Cloud software is [X]. Fewer than [20-30]% of enterprises that have [X].
- 2.7 **VMware lacks the scale needed to develop its Private Cloud software and increase consumption.** The growth of CSPs, described in the Ofcom Cloud Market Study,¹⁰ demonstrates the significant and increasing demand from enterprises to deploy workloads in a cloud environment. Broadcom believes that VMware is unable to compete with CSPs in part because its VCF Private Cloud software is [X]. While VMware's VCF is [X]. [X].
- 2.8 To address product deficiencies and increase utilisation, VMware needs to invest more in R&D and provide more deployment support to its enterprise customers. As a standalone public company, however, VMware lacks the scale to make the necessary investments. VMware's FY2023 cloud revenues [X] Amazon's c. \$80 billion, Microsoft's c. \$75 billion, and Google's \$26 billion FY2022 cloud revenues.¹¹ At a company-wide level, VMware's c. \$12 billion revenues in FY2021 were tiny in

⁶ VMware's Private Cloud software includes vSphere, vRealize, vSAN, NSX, vCenter, and SDDC Manager.

⁷ Off-premises private clouds are gated from other tenants using firewalls and other isolation methods. Enterprises often manage off-premises private clouds themselves.

⁸ Based on VMware company materials provided to Broadcom during due diligence.

⁹ IDC expects private cloud "to grow at a five-year CAGR of 151.8% and cross \$7.6 billion in revenue in 2025 (see [Dedicated Cloud Infrastructure as a Service, 2019–2025: Market Trends and Outlook](#)).

¹⁰ In the Ofcom Cloud Market Study, Ofcom found, e.g., that: 82% of survey respondents have increased their cloud spend in recent years and 79% expect to spend more on cloud in the next 18 months (paragraphs 3.12 and 3.13).

¹¹ Amazon, [Amazon.com Announces Fourth Quarter Results](#), p.1, Microsoft, [Annual Report 2022](#), and Alphabet, [2022 Annual Report](#), p.28.

comparison to Amazon's \$470 billion, Microsoft's \$168 billion, and Google's \$258 billion.¹²

Broadcom brings scale and will invest \$[X] more per year to grow VMware's position in private cloud

- 2.9 **Broadcom will increase R&D investment by \$1 billion per year.** While Broadcom is still small in comparison with the CSPs, it has sufficient scale to increase VMware's R&D investment by [X]% per year. This investment will be used to improve VMware's [X], by increasing the interoperability of the [X],¹³ and make it easier to deploy by enterprise professionals that are currently unfamiliar with the software. Broadcom plans to [X]. This will in turn facilitate the movement of workloads among different (private and public) clouds, thereby providing enterprises with greater control of their workloads.
- 2.10 **Broadcom will invest an additional \$1 billion per year to build deployment support capabilities.** Broadcom plans to double VMware's professional services capabilities from \$1 billion to \$2 billion per year and to offer professional services to VMware's customers [X]. Broadcom plans to train and certify up to [X] external professionals over three years at Global Systems Integrators (GSIs), such as [X], to enable them to provide enterprises with the necessary expertise to [X]. GSIs would also highlight and promote better ways for enterprises to deploy and embed workloads in a private cloud based at their on-premises datacentres (or hosted off premises) instead of moving workloads to the public cloud.
- 2.11 **Broadcom will increase the efficiency of VMware's go-to-market and administrative structure.** Broadcom has a strong track record of fiscal discipline, and plans to achieve c. \$[X] cost synergies by leveraging Broadcom's [X].¹⁴ Broadcom also plans to increase [X] over the course of its acquisitions of CA Technologies and Symantec. These cost savings will free up more resources to invest in R&D and the GSI partner network.

Broadcom will enable VMware to increase competition with the CSPs

- 2.12 **Enterprises want a private cloud solution that offers the ease of use, flexibility, and resilience of public cloud.** Despite having the technology and incentive to do so, VMware has been unable to satisfy enterprise demand, after many years of effort. Broadcom's investment will give VMware the scale, investment, and support it needs to deliver on its potential. It will create a credible private cloud alternative to the CSPs [X], which will require the Combined Entity to support the broadest hardware compatibility to ensure customers can choose the appropriate hardware for their needs (including lower-priced alternatives). Broadcom will make VMware's Private Cloud software more compelling and easier to deploy and use, which will enable it to compete

¹² Amazon, [Amazon.com Announces Fourth Quarter Results](#), Microsoft, [Annual Report 2021](#), Alphabet, [2021 Annual Report](#), p.45.

¹³ For the avoidance of doubt, the Combined Entity would continue to offer the individual components of the [X] on a standalone basis.

¹⁴ Broadcom projects these cost savings through synergies in [X] (see BCOM-CMA-00000047, slide 4).

more effectively with the CSPs for enterprise workloads. Broadcom also has plans to enable enterprises to use VMware's [X].¹⁵

Broadcom is focused on growing deployment and consumption of VMware's Private Cloud software

- 2.13 Through the strategies outlined above, Broadcom plans to grow VMware product revenues by \$[X] over three years through increased product deployment and utilisation.¹⁶ More generally, Broadcom is focused on winning more enterprise workloads and growing VMware's share of the \$400 billion plus cloud market.¹⁷
- 2.14 If Broadcom is to achieve these ambitious goals, it cannot increase prices or degrade interoperability, as some third parties have speculated (without basis). These putative strategies are not only unsupported by evidence, they would also make no commercial sense because they would *decrease* VMware sales and *accelerate* migration of enterprise workloads away from VMware. Put simply, they would achieve the opposite of Broadcom's actual and documented strategy. Indeed, increasing prices or degrading interoperability would damage Broadcom's and VMware's reputations and risk undermining Broadcom's \$61 billion investment in VMware.

3. INTEROPERABILITY IS THE FOUNDATION OF THE DATACENTRE

- 3.1 Datacentres house the equipment necessary to enable enterprises to deploy their business applications (e.g., email, booking systems, CRM, accounting software, etc.).¹⁸ Business applications run on operating systems (OSs) on servers and create data that must be stored either locally in the server, on a SAN (Ethernet or FC SAN), or in the cloud. Each server, and the datacentre more generally, comprises hundreds, if not thousands, of components from different vendors. All of them must interoperate.
- 3.2 Enterprises choose from a wide range of hardware and software options when configuring their servers and building their datacentres. Customers do not need to consider interoperability between different hardware components because they are designed to work together, based on open, industry standards. Similarly, when choosing OSs and, if enterprises so choose, virtualisation software to run on their servers, enterprises do not need to consider hardware/software interoperability. It is taken as a given fact.
- 3.3 Undermining interoperability would introduce the complexity and cost that industry participants have worked for years to mitigate. As the evidence below shows, interoperability is a business imperative for both Broadcom and VMware. Interoperability is, moreover, essential for server OEMs, whose business model

¹⁵ More specifically, Broadcom plans to develop a network of [X] that would use VMware's [X].

¹⁶ See BCOM-CMA-00000047, slide 4: Broadcom's operating model for VMware plans growth in software revenue from \$[X] to \$[X].

¹⁷ See Gartner, Worldwide Public Cloud Services End-User Spending Forecast.

¹⁸ A large datacentre may contain tens or hundreds of thousands of servers. A medium datacentre may have hundreds or thousands of servers. A small firm may have a single rack or row of servers. For context, VMware's top [X] customers have [X] servers on average and its top [X] customer have [X] on average.

depends on combining hardware and software from different vendors and offering enterprises the wide variety of choice they demand.

- 3.4 It is for these reasons that no hardware or software vendor has sought to degrade interoperability in the datacentre.

Experience shows that hardware rivals must ensure interoperability

- 3.5 There are numerous examples of hardware vendors ensuring interoperability with rival products even when, technically, they could have attempted not to. Four illustrative examples are provided below.

- 3.6 **VMware’s past owners ensured that vSphere remained hardware neutral.** VMware was owned by two of the largest hardware companies in the world until it was spun off in November 2021. Neither EMC nor Dell deviated from ensuring the broadest possible interoperability between storage arrays and servers, respectively, and vSphere.¹⁹ Dell and EMC had more to gain from foreclosing rivals; they supplied server and storage products that cost as much as a million dollars more than the few hundred dollars that Broadcom’s I/O devices sell for. But both kept VMware neutral because interoperability is paramount. The speculation in the Phase 1 Decision to the contrary is entirely unsubstantiated.²⁰ In fact, VMware worked with rivals of EMC such as NetApp and Pure Storage that promoted disruptive technology that threatened EMC’s market position.

- 3.7 **Intel and AMD ensure interoperability between their CPUs and rivals’ NICs.** Server CPUs are the “*brains of computer systems.*”²¹ They “*perform all types of operations [including] running software, analysing data, managing network traffic, and fetching data from memory.*”²² Enterprises typically build their servers around their chosen CPU and they are generally the most expensive components of a server.²³ Despite Intel and AMD accounting for approximately 95% of all server CPUs²⁴ and offering their own Ethernet NICs, neither has attempted to degrade interoperability between their CPUs and Ethernet NICs supplied by hardware rivals (*e.g.*, Broadcom, NVIDIA, Cisco, and Marvell) in a putative attempt to boost their NIC sales.

- 3.8 **NVIDIA ensures interoperability between its GPUs and rivals’ NICs.** GPUs process graphic images or computations that require massive parallel execution of computational tasks, and are typically used in large scale datacentres and key artificial intelligence applications.²⁵ Despite NVIDIA accounting for 90-100% of datacentre

¹⁹ For instance, VMware worked with EMC rivals Pure Storage and NetApp, both of whom had disruptive technology that fundamentally threatened EMC’s market position (*see* paragraph 2.11 of the Parties’ response to the Issues Letter dated 3 March 2023).

²⁰ Phase 1 Decision, paragraph 221.

²¹ *AMD/Xilinx*, Case COMP/M.10097, Commission decision of 30 June 2021, paragraph 9.

²² *Ibid.*

²³ The price of a CPU ranges from c. \$75 to over \$3,000 (*see, e.g.*, the processor options available for the PowerEdge XR11 Rack Server on Dell’s [server configurator](#)).

²⁴ FMN, paragraph 15.450.

²⁵ *NVIDIA/Mellanox*, Case Comp COMP/M.9424, Commission decision of 19 December 2019, paragraph 19.

GPUs²⁶ and offering its own Ethernet NICs, it has not attempted to degrade interoperability between its GPUs and Ethernet NICs supplied by hardware rivals.

- 3.9 **Broadcom’s business model ensures interoperability and neutrality.** Broadcom’s business model is based on independent, sustainable business franchises. Broadcom ensures that its products and components work with the wide range of hardware and software found in a customer’s datacentre. Broadcom has not degraded interoperability on any occasion in relation to any products, even in circumstances which are similar to those being considered by the CMA today (e.g., [§]). It has never even been contemplated. Doing so would depart from the golden rule of interoperability in the datacentre, and be ruinous to Broadcom’s reputation and sales.
- 3.10 In short, software and hardware vendors in the datacentre environment do not and would not degrade interoperability. It would cross a well-established red line and would result in fierce backlash from rivals, OEMs, and end customers.

Broad hardware neutrality is essential for virtualisation software

- 3.11 The two core values of virtualisation software are: (1) making more efficient use of physical server resources, and (2) enabling customers to manage uniformly the diverse range of hardware and software components in datacentres by creating a software abstraction layer. Virtualisation software cannot perform its core functions without interoperating with the various hardware components that enterprises have chosen over time to use in their servers. It is unsurprising, therefore, that VMware, its customers, its rivals, its partners, industry reporters, and regulators have recognised the critical importance of interoperability to VMware.
- 3.12 **VMware recognises that interoperability is crucial for its success.** VMware’s financial statements over many years explain that “[t]he success of [its] products depends upon the cooperation of hardware and software vendors to ensure interoperability with [its] products and offer compatible products and services to end users.”²⁷ VMware’s internal documents also explain that having the [§]²⁸ and that VMware’ market position is [§].²⁹ These statements show that VMware considers interoperability to be a business imperative and offering the broadest interoperability to be a parameter of competition.
- 3.13 **VMware customers have stated that interoperability is a critical part of the value of VMware for them.** As discussed further below in Section 4, Management Insight Technologies (MIT) recently carried out a survey of approximately 1,200 VMware customers (the **2023 MIT survey**) exploring how VMware customers would respond to any hypothetical plan to degrade the interoperability of VMware with non-Broadcom NICs, storage adapters, and FC HBAs. Open-ended responses to the survey, by themselves, show clearly that VMware would lose customers if it deviated from hardware neutrality: “*limiting VMware to Broadcom hardware is an absolute deal breaker for us. VMware would be gone from our environment never to return*”;

²⁶ *Ibid*, paragraph 264.

²⁷ VMware Form 10-K, FY2020-2022.

²⁸ VMware Annex Q15-011 to the FMN.

²⁹ CMA Phase 1 Decision, paragraph 122.

“[h]ardware vendor lock-in would result in **a complete exit** of VMware to alternative providers”; and VMware’s dropping “support for some of the common vendors of critical hardware” would result in a “**massive customer exodus**” (emphasis added).³⁰

- 3.14 **Software rivals and hardware partners identify interoperability as essential for virtualisation software.** [X], a virtualisation rival, has explained that “[X]”.³¹ [X] I/O device vendor, has also indicated that “[X].”³² [X].
- 3.15 **Industry press reports highlight the importance of interoperability.** Technology-news service, The Register, has commented that “[r]estricting user choice for hardware would [represent] an act of self-harm” for VMware.³³ MarketResearch likewise explains that VMware “competes mainly on the basis of reliability, interoperability [...] and the ability to offer products that support multiple hardware platforms”.³⁴
- 3.16 **Both the Commission and the CMA have recognised the importance of interoperability.** In its *Dell/EMC* merger investigation, the Commission found that VMware has always adopted a “hardware/software-neutral approach” and that “in order to promote a large adoption of its product, VMware **had to endorse an open and non-discriminatory architecture policy**” (emphasis added).³⁵ [X].³⁶ In its Phase 1 Decision, the CMA “agrees that a complete breakdown in interoperability [with products using a driver that has previously been certified] would cause undesirable disruption to customers.”³⁷

OEMs multi-source and demand interoperability to offer choice to enterprises

- 3.17 OEMs play an important role in the datacentre supply chain. They have the primary relationship with enterprise customers and can influence the hardware options that are made available to enterprises. OEMs compete intensely with each other, *inter alia*, on price, range of choice, and reliability. To do so, OEMs must multi-source hardware components. OEMs also multi-source to reduce costs and increase profit margins. Put simply, multi-sourcing is critical for OEMs’ commercial success.
- 3.18 **Enterprise customers want choices when configuring servers.** It is clear from OEMs’ online server configurators – web pages that allow customers to build servers

³⁰ RBB, Implications of the MIT Switching Survey (RBB paper on 2023 MIT survey). See [Annex ISR-002](#).

³¹ The Parties’ response to the Commission’s Article 6(1)(c) Decision (the **6(1)(c) Response**), paragraph 200.

³² See 6(1)(c) Response, paragraph 37. See too the Parties response to the Commission’s Statement of Objections (the **SO Response**), paragraph 384.

³³ The Register, [EU probes Broadcom/VMware deal over impact on hardware, not price hikes or cloudy concerns](#).

³⁴ MarketResearch, Global Software Defined Storage Market Research Report, April, 2023. See [Annex ISR-003](#).

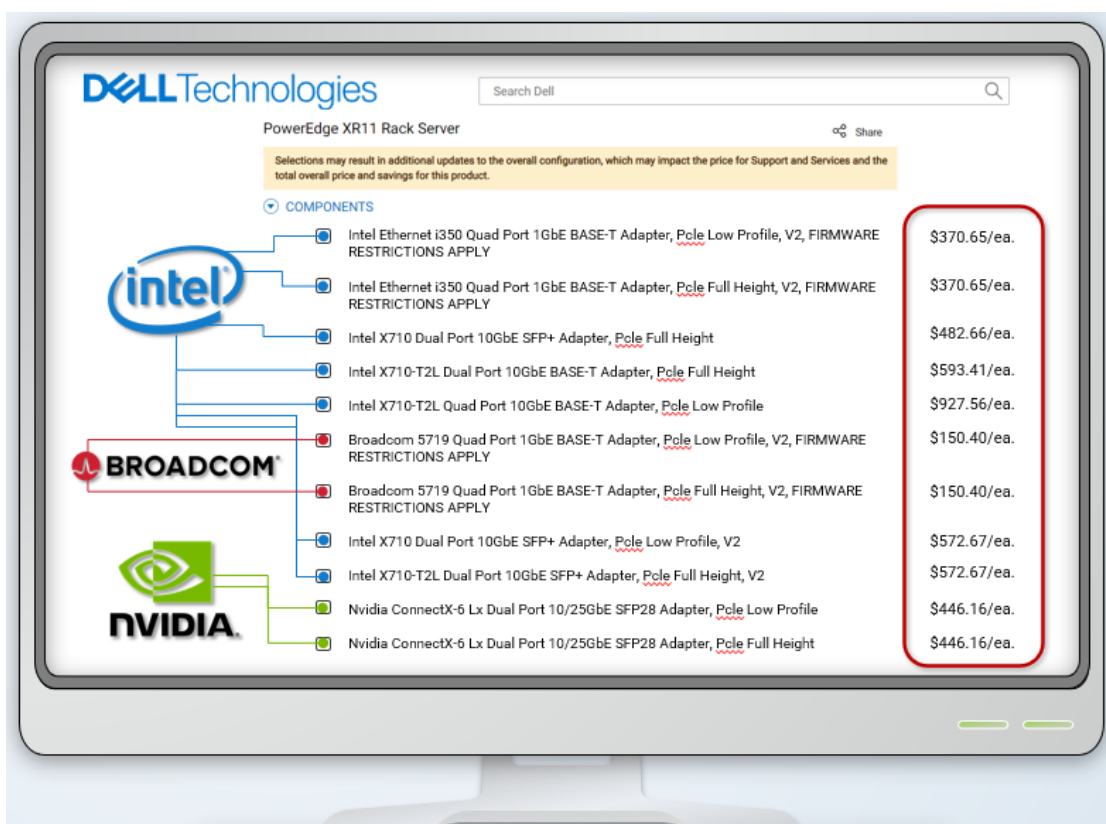
³⁵ *Dell/EMC*, Case COMP/M.7861, Commission decision of 29 February 2016, paragraphs 177-178.

³⁶ SO, paragraph 422.

³⁷ CMA, Phase 1 decision, paragraph 188.

by selecting their chosen components – that offering a wide variety of choices is important to them and their customers (see **Figure 2** below).

Figure 2: It is important for OEMs to provide customers with choices



3.19 For instance, Dell’s PowerEdge XR11 Rack Server allows customers to choose from **25** different NICs with different architectural features and price points.³⁸ The same is true for FC HBAs and storage adapters. For example, Lenovo’s server configurator offers an array of options including FC HBA SKUs released more than 5 years ago.³⁹ OEMs offer choices to cater for the multitude of different enterprise preferences, including in relation to price, form factor, vendor, and product generation.⁴⁰ [38].

3.20 **Maintaining price competition between two or more suppliers is important for OEMs to drive down costs.** OEMs are sophisticated customers that would not leave themselves at the mercy of a monopolist hardware component supplier. This is confirmed by numerous sources of evidence:

- (a) *The CMA’s Phase 1 market investigation confirms the importance of multi-sourcing for price competition.* The Phase 1 Decision explains that “[t]hird parties [...] emphasised the importance of multiple supplier options, with one

³⁸ See [Dell server configurator](#).

³⁹ See SO Response, paragraph 520.

⁴⁰ See [Annex ISR-004](#), which explains in further detail end user and OEM preferences in FC HBAs, NICs and storage adapters, supported by Broadcom’s internal documents.

server OEM stating that it is important to have multiple suppliers of the components in order to obtain the best pricing.”⁴¹

(b) [REDACTED].

(c) ***Broadcom’s internal documents show that OEMs regularly pit Broadcom against rivals to obtain lower prices.*** For example, Broadcom [REDACTED]⁴² Similarly, Broadcom [REDACTED].⁴³

(d) [REDACTED].⁴⁴ [REDACTED].⁴⁵ [REDACTED].⁴⁶

3.21 **Maintaining price competition on I/O devices specifically is important for OEMs to maintain the profitability of their server businesses.** Server OEMs operate on thin margins for standalone servers and rely on “options” like I/O devices to increase – or sometimes make – profits. This is evident from an email from [REDACTED] to [REDACTED]: “[REDACTED]” (emphasis added).⁴⁷ Degrading interoperability would therefore impact OEMs’ profitability directly.

3.22 **Multi-sourcing is important for OEMs to ensure they have an alternative if a supplier has supply difficulties.** [REDACTED]. [REDACTED].⁴⁸

3.23 Degrading interoperability would increase OEMs’ costs, reduce their margins, weaken their competitiveness, and expose them to supply chain risks. If the Combined Entity were to attempt such a strategy, OEMs would retaliate harshly and swiftly in a variety of ways, including by removing products from configurators (or demoting them), switching purchases away from Broadcom, and stopping promotion of VMware (as explained in Section 4 below).⁴⁹

4. THEORY OF HARM 1: FORECLOSURE OF HARDWARE COMPETITORS

4.1 The Issues Statement’s ToH1 considers whether the Combined Entity could “leverage VMware’s market power in server virtualisation software to reduce the competitiveness of Broadcom’s hardware rivals.”⁵⁰ More specifically, ToH1 concerns (a) whether the Combined Entity could impair “the certification of competitors’ drivers” for I/O devices and (b) whether it could impair “access to VMware’s API for competitors’ FC switches.”⁵¹ This Response addresses the CMA’s theory in relation to I/O devices and

⁴¹ CMA, Phase 1 decision, paragraph 218(a).

⁴² See **Annex ISR-005**.

⁴³ See **Annex ISR-006**.

⁴⁴ See [REDACTED] (**Annex ISR-014**), paragraph 8.

⁴⁵ *Ibid.*, paragraph 9.

⁴⁶ *Ibid.*

⁴⁷ See **Annex ISR-007**.

⁴⁸ [REDACTED], paragraph 9.

⁴⁹ See paragraph 4.26 below.

⁵⁰ Issues Statement, paragraph 22.

⁵¹ *Ibid.*

FC switches separately because their respective communications with vSphere are fundamentally different.

(A) No SLC in I/O devices: storage adapters, FC HBAs and NICs

4.2 I/O devices communicate with vSphere through device drivers, which provide a software mapping between the I/O device and the vSphere APIs. As the CMA accepts, the only way the Combined Entity could, in theory, degrade interoperability with I/O devices is via the driver. The Issues Statement contends that the Combined Entity may foreclose rival vendors of NICs, FC HBAs, and storage adapters in the following ways:

- (a) **Total foreclosure of new devices.** The Issues Statement only questions whether the Combined entity could **prevent** “*interoperability between rivals’ new hardware products that have not yet had drivers developed and certified.*”⁵² It accepts, rightly, the CMA’s Phase 1 finding that the Combined Entity cannot prevent interoperability for products that are “*currently installed in servers*” which would cause “*disruption to customers’ existing servers*”;⁵³ and
- (b) **Partial foreclosure of new and existing devices.** The Issues Statement contends that the Combined Entity could refuse, delay, or hamper “*information exchange regarding driver updates or technical support for rivals’ new or existing hardware.*”⁵⁴

4.3 The evidence and explanations below demonstrate, however, that the Combined Entity would have no ability or incentive to engage in these putative foreclosure strategies and that any supposed “*harm to competitors*” would in any event not “*result in a substantial harm to overall competition*” in the supply of storage adapters, NICs, or FC HBAs.⁵⁵

The Combined Entity would have no ability to foreclose hardware rivals

4.4 The Combined Entity could not foreclose rival I/O device vendors. Not only would it break the golden rule of the datacentre, cause retaliation from OEMs, and destroy Broadcom’s and VMware’s reputations, it would accelerate switching of enterprise workloads away from VMware to on-premises and cloud rivals. The premise of ToH1(a) is that enterprises have no real alternatives for their workloads and, therefore, would switch to Broadcom hardware despite it having engaged in a foreclosure strategy of degraded interoperability. In reality, however, enterprises have numerous credible alternatives for their workloads.

4.5 **VMware is constrained by rival environments for deploying workloads.** Enterprises could “*easily switch away from [vSphere] to a range of effective alternative suppliers*” if the Combined Entity were to attempt to degrade interoperability:⁵⁶

⁵² Issues Statement, paragraph 22(a).

⁵³ CMA, Phase 1 Decision, paragraph 189.

⁵⁴ Issues Statement, paragraph 22(a).

⁵⁵ CMA, Merger Assessment Guidelines (MAGs), paragraph 7.35.

⁵⁶ CMA, MAGs, paragraph 7.14(a). This paragraph relates to the CMA’s input foreclosure analysis, but the MAGs confirm that the “*CMA may use the same framework in similar situations where the merged*

- (a) ***VMware’s share is declining because enterprises are choosing alternatives for their workloads.*** The most relevant measure of concentration is based on the number of new virtualised licences shipped globally, which reflects the current market conditions and competitive constraints.⁵⁷ The Phase 1 Decision contends that “*shares based on the number of CPUs installed are likely to be more appropriate*” because foreclosure “*could affect both existing and/or new servers*”, but this is mistaken for the reasons explained below.⁵⁸ Even when looking only at the datacentre, VMware’s share has declined from [20-30]% to [20-30]% between 2019 and 2021.⁵⁹ For all new deployments (including public cloud), VMware’s share declined from [10-20]% to only **[10-20]%**.⁶⁰ These IDC shares are broadly consistent with information VMware relies on in the ordinary course, including market shares that VMware presented to its Board of Directors more than a year before the Transaction was announced, to inform VMware’s strategy for competing with the CSPs. Indeed, the market shares presented to the Board of Directors show that, even on the (less appropriate) installed-base basis, VMware had a share of only [30-40]% of workloads on-premises (including the datacentre and private cloud) and only [10-20]% of workloads across all deployments (including public cloud).⁶¹ These modest and declining shares indicate that enterprises have credible alternatives for their workloads. They are far from a level that could be associated with “*market power*”.⁶²

entity could use its presence in one market to directly harm the competitiveness in another [e.g.,] by using control of a complementary product to deteriorate its interoperability with competitors” (paragraph 7.11).

⁵⁷ The CMA has consistently placed limited reliance on market shares that do not reflect the current state of current or future competition: in *Viasat/Inmarsat*, the CMA found historic shares were of “*limited evidentiary value in assessing suppliers’ current competitive strength*” because they resulted from competition “*many years ago*” (see Anticipated Acquisition by Viasat, Inc. Of Connect Topco Limited, Final Report, paragraph 8.116); and in *Sabre/Farelogix*, the CMA found market shares were “*less indicative of future competitive constraint*” (see CMA, Anticipated acquisition by Sabre Corporation of Farelogix Inc., Final Report, paragraph 11.23). The Phase 1 Decision relies on data from the Parties’ internal documents that are less reliable. For instance, the CMA highlights an Annex ([§]) in which VMware estimates that it has a share of [§]% by number of virtualised instances. [§]. The CMA also highlights an Annex ([§]) in which VMware estimates that it has a [§]% share in the overall server virtualisation market, which rises to an estimate of [§]% in the paid-for market. However, the version of the document that was ultimately *presented* to VMware’s Board of Directors (see the document submitted previously to the CMA as [§] *did not* include this market share reference – which proves that any relevant reference was not deemed sufficiently reliable. The [§]% market share reflects [§] *revenue-based shares*, which the Phase 1 Decision accepts are less reliable. [§].

⁵⁸ Phase 1 Decision, paragraph 110(a). See paragraph 4.7.

⁵⁹ See ‘**Annex RFI 13 Q1**’, page 3 provided at **Annex RFI4Q5-003** in response to the CMA’s request for information dated 25 November 2023.

⁶⁰ See annexes **RFI4Q5-001 - RFI4Q5-004** to the FMN. If “*VMware’s actual internal data*” is combined “*with IDC*” data, as the Phase 1 Decision suggests would be “*more reliable*” (see paragraph 110(c)), VMware’s market share on premises declined from [30-40]% to [20-30]% from 2019 to 2021 and its share of all deployments declined from [10-20]% to [10-20]% over the same period (see SO Response, Table 1).

⁶¹ See **Annex Q10(VM) – 011**, pp. 114-15.

⁶² Issues Statement, paragraph 22.

- (b) ***VMware’s declining share of enterprise workloads is confirmed by its many lost opportunities each year.*** vSphere lost more than \$[redacted] of opportunities for new vSphere licences annually during FY2020-2022 (notably, total losses were substantially larger due to lost associated annual revenues, including for services and support).⁶³ The Phase 1 Decision suggests that VMware’s opportunities data is consistent with its having a high share of supply.⁶⁴ But this is both incorrect – opportunities data is not a measure of share of supply – and beside the point. What matters for ToH1 is whether enterprises have credible alternatives for a sufficient number of their workloads at the point where VMware is bidding for their workloads to render foreclosure unprofitable. The opportunities data show clearly that they do.⁶⁵
- (c) ***vSphere competes intensely with strong on-premises rivals,*** which is confirmed by a variety of sources.
- (i) VMware’s internal documents confirm that Hyper-V and IBM/Red Hat are close alternatives to vSphere. For instance, internal documents describe Hyper-V as a source of [redacted],⁶⁶ and state that Hyper-V is [redacted]⁶⁷ and less [redacted] – making VMware products [redacted].⁶⁸ VMware is also constrained by [redacted]⁶⁹ [redacted], as evidenced by the need to give heavy (c. [redacted]%) discounts to [redacted],⁷⁰ including to [redacted].⁷¹ Evidence gathered by the Commission also confirms the strong constraint. The parties are seeking the Commission’s consent to share that evidence with the CMA.
- (ii) Internal VMware documents also identify Nutanix as a [redacted] to [redacted].⁷² VMware considers Nutanix to be the [redacted] virtual hypervisor.⁷³ Nutanix is seen as [redacted].⁷⁴ A customer described Nutanix as [redacted] on

⁶³ See [redacted].

⁶⁴ Phase 1 Decision, paragraph 126.

⁶⁵ Examples of lost opportunities include: [redacted].

⁶⁶ VMware internal document, email exchange with subject [redacted]. See [Annex ISR-008](#).

⁶⁷ VMware internal document, email exchange with subject: [redacted]. See [Annex ISR-009](#).

⁶⁸ VMware internal document, [redacted]. See [Annex ISR-010](#).

⁶⁹ VMware internal document, email exchange with subject [redacted]. See [Annex ISR-011](#).

⁷⁰ VMware internal document, email exchange with subject [redacted]. See [Annex ISR-012](#). The email notes that [redacted].

⁷¹ *Ibid.*

⁷² VMware internal document, [redacted]. See [Annex ISR-013](#).

⁷³ VMware internal document, email exchange with subject [redacted]. See [Annex ISR-009](#).

⁷⁴ VMware internal document, email exchange with subject [redacted]. See RSLV_00028772.

VMware.⁷⁵ Workloads are also moving to KVM. For example, [REDACTED].⁷⁶ The constraint from Nutanix and KVM is also evident from [REDACTED].⁷⁷

- (iii) **Customers routinely multi-source from two or more on-premises solutions, increasing competitive pressure.** This is evident from VMware’s internal documents. For instance, [REDACTED].⁷⁸ [REDACTED].⁷⁹ It is also evident from Broadcom’s experience switching workloads away from VMware.⁸⁰ Customers that already use a second virtualisation provider can more readily switch workloads way from vSphere.
- (d) **CSPs represent a strong and growing competitive constraint.** The Phase 1 Decision acknowledges that public cloud has been “*rapidly growing in recent years.*”⁸¹ Despite this, it suggests that the constraint from public cloud is “*likely to be limited.*”⁸² This is not supported by the evidence available:
 - (i) **The Ofcom Cloud Market Study confirms increasing cloud migration.** Having carried out “*50 one-hour discussions and over 1000 survey interviews with UK decision-makers in UK businesses,*” Ofcom calculated that, of these UK respondents: >40% are migrating more workloads to the cloud; >80% expect to increase their cloud spend; >70% use cloud for storage databases, and back up; >50% use cloud for hosting websites, social media, business applications, software development testing, and business intelligence; and >50% move to public cloud for flexibility, agility, and security.⁸³
 - (ii) **VMware’s internal documents confirm that the public cloud is the strongest competitive threat** and show that VMware customers are “*shifting workloads to the public cloud.*”⁸⁴ For example, a document presented to VMware’s Board of Directors shows that, of ~[REDACTED] million enterprise workloads running across all deployments in 2021, around [REDACTED] million ([REDACTED]%) are on public cloud, which is projected to grow to

⁷⁵ [REDACTED].

⁷⁶ VMware internal document, [REDACTED]. See [Annex ISR-015](#).

⁷⁷ See slide 30 of Broadcom’s Site Visit Presentation.

⁷⁸ VMware internal document, email exchange with subject [REDACTED]. See [Annex ISR-016](#).

⁷⁹ VMware Internal Document, email exchange with subject line [REDACTED]. See [Annex ISR-017](#).

⁸⁰ As explained in the response to question 3 of Phase 2 RFI 1.

⁸¹ Phase 1 Decision, paragraph 151.

⁸² Phase 1 Decision, paragraph 61. The Phase 1 Decision relies on third-party evidence that “*most customer [...] either do not consider public cloud to be an alternative [...] in the short run, or only for some of their workloads, or not at all.*” However, this formulation obfuscates what proportion of third parties indicated that public cloud is not an alternative at all, as opposed to not being an alternative in the short-run or for all of their workloads. Indeed, if the majority of third parties identified CSPs as an alternative for some of the workloads in the long-run, it would show that they represent a strong and growing constraint, not a limited one.

⁸³ Ofcom Cloud Market Study, Figures 3.3, 3.4, 3.5, and 3.6.

⁸⁴ Phase 1 Decision, paragraph 155.

[§<] million workloads by 2024.⁸⁵ As noted above, of those ~[§<] million workloads, VMware has a share of only ~[§<]% (which is taking account of its installed base).⁸⁶ VMware estimated in 2022 that [§<] *i.e.*, that [§<]% of workloads switching away from VMware go to the public cloud.⁸⁷ Moreover, VMware expects [§<].⁸⁸ This document also [§<].⁸⁹ Other customers choosing the public cloud are [§<],⁹⁰ and [§<], which is [§<].⁹¹

- (iii) ***VMware documents and industry reports show that CSPs are also growing as a private cloud and multi-cloud provider.*** VMware describes AWS Outposts as [§<],⁹² [§<],⁹³ [§<].⁹⁴ The 2023 Flexera State of the Cloud report concludes that “[m]ost organizations are taking a multi-cloud, hybrid approach in which private cloud plays an essential role. Microsoft Azure Stack ranked first, with 41% currently running workloads, up from 37% last year. AWS Outposts switched places with VMware vSphere/vCenter for second place.”⁹⁵ Azure Stack is also identified as a strong competitor in VMware documents. For example, Van Havermaet, a Belgium-based accounting firm, “had historically run its workloads and applications on-premises, and for the past decade on VMware ESX,” but “turned to Azure Stack” to build an “on-premises environment as close to Azure as possible.”⁹⁶ Similarly, Picanol Group “wanted to manage [their] whole environment with a single solution. That’s what [they] got with Azure Stack HCI—the flexibility to run workloads in the cloud or on-premises and decide at each moment where [they] want [their] resources.”⁹⁷ VMware’s CEO described Google Anthos in 2022 as [§<]⁹⁸ In general, a study by a consulting firm states that: [§<].⁹⁹

85 VMware Site Visit Presentation, page 13. See [§<].

86 *Ibid.*

87 [§<]. See **Annex ISR-018**.

88 *Ibid.*, at page 52.

89 *Ibid.*

90 VMware internal document, [§<]. See **Annex ISR-019**.

91 VMware internal document, [§<]. See **Annex ISR-020**.

92 VMware internal document, email exchange with subject [§<]. See **Annex ISR-021**.

93 *Ibid.*

94 VMware internal document, [§<]. See **Annex ISR-022**.

95 Flexera, 2023 State of the Cloud Report, p. 66. See **Annex ISR-023**.

96 Microsoft, Customer Stories: Van Havermaet scales up innovation to deliver superior accounting advice with Azure Stack HCI (7 April 2023).

97 Microsoft, Customer Stories: Picanol Group finds flexibility and faster time to market with Azure Stack (28 February 2023).

98 VMware internal document, email exchange with subject [§<]. See **Annex ISR-024**.

99 [§<], provided to the CMA on 16 January 2023 as Annex – Response to 6(1)(c) Decision - 6.

- (iv) ***CSPs and other “key competitors” are increasing competitive pressure by offering full stack solutions (including hardware).*** A VMware internal document states that: [§].¹⁰⁰
- (v) ***Industry surveys show cloud migration across industries.*** VMware’s 2021 Workloads Study of 1,700 respondents shows that nearly [§]% of *net new workloads* deployed in the previous two years by VMware customers were deployed on virtual machines in the public cloud.¹⁰¹ It also shows that [§]% of respondents, across industry categories, moved workloads to the public cloud in the two previous years and [§]% migrated all evaluated workloads.¹⁰² Gartner’s 2022 “Cloud Shift” research forecasts that enterprise IT spending on public cloud computing will overtake spending on traditional datacentre IT in 2025.¹⁰³ Contrary to the suggestion in the Phase 1 Decision that IT systems in regulated sectors “*may always remain on-premise*”,¹⁰⁴ a 2021 consulting firm report estimates that [§]% of workloads even in highly-regulated industries are expected to move away from vSphere, many of them to public cloud.¹⁰⁵ The same report includes customer statements that VMware (1) lacks the [§], and (2) needs to [§].¹⁰⁶
- (vi) ***Broadcom itself moved [§] from vSphere to the public cloud,*** in particular to Google Cloud Platform (GCP).¹⁰⁷ Broadcom moved [§] virtual machines (VMs) to GCP between [§].¹⁰⁸ The move was facilitated by [§].¹⁰⁹ As of January 2023, over [§]% of Broadcom’s VMs were deployed in the public cloud.¹¹⁰ Broadcom estimates that, in the past two years, approximately [§]% of its net new workloads were deployed in the public cloud.
- (e) **vSphere competes directly with containers for workloads.** This is confirmed by Ofcom which recently found that “[*v*]irtual machines and containers offer similar functionalities.”¹¹¹ The Phase 1 Decision also confirms that some customers “*consider containerization to be an alternative for all their*

¹⁰⁰ VMware internal document, [§]. See [Annex ISR-022](#).

¹⁰¹ [Annex Q15-011](#) to the FMN.

¹⁰² *Ibid.* See also RBB paper on 2023 MIT survey, paragraph 3.4 and Figure 9, which, consistent with the Workloads Study, shows that irrespective of industry and company size, c. [§]% of respondents indicate that they are planning to migrate workloads away from vSphere in the ordinary course.

¹⁰³ Gartner, [Gartner Says More Than Half of Enterprise IT Spending in Key Market Segments Will Shift to the Cloud by 2025](#) (9 February 2022).

¹⁰⁴ Phase 1 Decision, paragraph 154.

¹⁰⁵ [§].

¹⁰⁶ *Ibid.*

¹⁰⁷ GCP was Broadcom’s preferred public cloud solution due to [§].

¹⁰⁸ Broadcom internal document, [§]. See [Annex ISR-026](#).

¹⁰⁹ Broadcom internal document, [§]. See [Annex ISR-027](#).

¹¹⁰ As explained in the response to question 3 of CMA Phase 2 RFI 1.

¹¹¹ Ofcom Cloud Market Study, footnote 44.

workloads in the short run” (emphasis added).¹¹² It is reasonable to infer that a greater proportion of customers consider containers to be an alternative for *some* of their workloads in the *longer term*, not least because the containerisation market is growing rapidly.¹¹³ Evidence gathered by the Commission in its investigation also shows that [×]. When containers are used in conjunction with virtualisation software, they “*reduce technical barriers to switching*”¹¹⁴ and thereby strengthen the constraint on VMware from rival hypervisors and CSPs.

(f) ***2023 MIT survey demonstrates that most customers would switch workloads away from VMware in response to a degradation of interoperability.*** The survey covered 1,204 VMware customers and examined how they would react to a foreclosure strategy. Responses show:

(i) ***Majority of VMware customers would switch further existing workloads away from VMware.*** [50-60]% of respondents for NICs, storage adapters, and FC HBAs indicated they would switch more workloads away from vSphere or plan to start switching workloads away in response to the next version of vSphere being incompatible with future generations of rival hardware components. They indicated that they would move or plan to move around [50-60]% of their vSphere workloads in response to incompatibility, on top of their current migration plans.¹¹⁵

(ii) ***Majority of customers would deploy fewer net new workloads on VMware in the following two years.*** [50-60]% of respondents for NICs, storage adapters, and FC HBAs indicated they would switch all or some of their planned net new workloads away from vSphere should vSphere become incompatible with rival hardware components. They indicated that this would involve [50-60]% of net new workloads currently planned for deployment on vSphere.¹¹⁶

(iii) ***VMware customers would switch workloads to both on-premises and cloud rivals,*** as shown in **Figure 3** below. On this basis, it is not sustainable to suggest, as the Phase 1 Decision did, that “*evidence on migration away from VMware on-premise deployments does not necessarily imply that customers would switch to public cloud in response to Broadcom foreclosing hardware competitors*”.¹¹⁷

¹¹² Phase 1 Decision, paragraph 64.

¹¹³ The application containerisation market is estimated to increase from \$1.5 billion in 2020 to \$9.7 billion by the end of 2027 (see Enterprise Storage Forum, [The Containerization Market in 2022](#) (6 July 2022)).

¹¹⁴ Ofcom, Cloud Market Study, paragraph 6.22.

¹¹⁵ RBB paper on 2023 MIT survey, p. 6.

¹¹⁶ RBB paper on 2023 MIT survey, p. 9.

¹¹⁷ Phase 1 Decision, paragraph 192.

Figure 3 Where workloads would be deployed if the Combined Entity were to degrade interoperability

[REDACTED]

- (iv) ***Only [10-30]% of VMware servers using non-Broadcom hardware would switch to Broadcom hardware in response to foreclosure.***¹¹⁸
These results imply that the actual switching rate would be far below the critical switching rate (which is 97% as explained further below).¹¹⁹

4.6 **The Combined Entity could not degrade interoperability with existing I/O devices in VMware servers.** The Issues Statement agrees with the Parties' submission that existing hardware devices could not be totally foreclosed. It contends, however, that existing devices could be partially foreclosed. The suggestion appears to be that *lessening* interoperability through hampering information exchange would not cause disruption to workloads yet would lead to switching to Broadcom hardware. This is incorrect and the reasoning is internally inconsistent.

- (a) ***Partial foreclosure of existing devices is implausible.*** Once hardware devices are installed in servers, there is minimal engagement between VMware and hardware vendors. VMware updates or support may be required when it is necessary to fix bugs or address issues that raise security or data corruption issues. Such instances are rare, however, and no other engagement is required beyond this. The Issues Statement speculates that the Combined Entity may try to leverage these opportunities to suggest to enterprises that the issues could be avoided by purchasing Broadcom hardware. But it is implausible that enterprises would respond to partial foreclosure by switching to Broadcom. If VMware failed to address bugs or fix security holes, it would be seen as unreliable and Broadcom would be seen as a bad actor that does not behave in customers' best interests. It would be readily apparent to OEMs and enterprises that the Combined Entity was entirely at fault for the disruption because rival products would visibly continue working without any issues in non-VMware environments (*i.e.*, in [90-100]% of servers). In these circumstances, customers that are concerned about reliability and trust, which would likely be a significant proportion, would migrate workloads away from VMware.¹²⁰ Partial foreclosure of existing devices would therefore harm VMware, not Broadcom's rivals.
- (b) ***Degrading interoperability with installed devices cannot result in any new Broadcom hardware sales.*** This foreclosure strategy is also implausible because it would impose an enormous burden on enterprises to buy new servers (which are several orders of magnitude more expensive than the individual I/O devices) or to change the hardware installed in their existing servers. The time and cost of replacing hardware components in servers is enormous. The manual

¹¹⁸ [10-30]% reflects the switching rate in the longer term, which compares to [REDACTED]% in year one (*see* RBB Switching rate paper, Tables 3 and 5).

¹¹⁹ *See* paragraph 4.19.

¹²⁰ *See* paragraph 4.5(f).

replacement alone would take an estimated 3-4 hours per server per I/O device. It would therefore take 3 to 5 years for a person to replace the NICs alone in 10,000 servers.¹²¹ In addition to the cost of buying and replacing the hardware, there would also be (1) the cost and disruption of server downtime to change and test equipment, and (2) the likelihood that OEMs' server warranties would be voided. For VMware to force such a burden on to customers would be nonsensical and commercially ruinous.¹²² Customers would simply not entertain replacing existing I/O devices, whether in response to partial or total foreclosure. Indeed, if a customer would not incur the burden of replacing existing hardware in response to a "complete breakdown in interoperability",¹²³ it is self-evident that it would not incur the same burden if interoperability were only reduced.

- (c) ***Total or partial foreclosure of existing devices would be disastrous for VMware's reputation.*** The Phase 1 Decision correctly acknowledges that "a complete breakdown in interoperability would cause undesirable disruption to customers".¹²⁴ It would be disastrous for VMware to cause disruption to enterprise workloads (e.g., preventing banks, governments, and hospitals from running software applications). This is particularly the case if the disruption was *intended* by the Combined Entity as part of a strategy to coerce customers to buy more hardware products from Broadcom. Whether disruption is caused by a total or partial foreclosure strategy, the result would be the same. This reality was accepted by the Commission's SO.¹²⁵
- (d) ***VMware has contracted to provide services and support to existing customers*** and therefore delaying, hampering, or refusing support for existing devices may breach its agreements with customers and expose VMware to liability. Moreover, partial foreclosure strategies would increase VMware's costs because it would receive more support calls from enterprises seeking troubleshooting. Each support call costs VMware c. \$[X]. The Phase 1 Decision mistakenly assumes that a customer experiencing unreliable I/O device operation or that has a security intrusion would make only a single support call.¹²⁶ But if issues persist – which they would need to for the foreclosure strategy to have an effect – a single customer (and different users from the same customer) may call numerous times. Across the installed base, this cost increase would be substantial. If 2,000 customers were to call VMware an additional five times per year, it would increase VMware's costs by \$[X] per year, which is far more than the potential gains from partial foreclosure.¹²⁷
- (e) ***OEMs would retaliate against Broadcom.*** Customers would alternatively ask server OEMs to resolve the issues. It would be clear to the server OEMs that

¹²¹ VMware's top [X] customers have c. [X] servers each on average.

¹²² See, for example, the SO, paragraph 560 (the Commission [X]).

¹²³ Phase 1 Decision, paragraph 188.

¹²⁴ Phase 1 Decision, paragraph 188.

¹²⁵ SO, paragraph 425.

¹²⁶ Phase 1 Decision, paragraph 223.

¹²⁷ See paragraph 4.24 below.

the cause of the issue was VMware if the component in question was working well with non-VMware servers. OEMs would swiftly retaliate unless VMware fixed the issue (in the ways discussed at paragraph 4.26 below).

- 4.7 **The Combined Entity could not degrade interoperability with new sales of existing I/O devices.** To minimise change and the risk of disruption, device drivers support multiple hardware generations. Customers’ adoption of new generations of devices is slow, meaning that a single driver supports hardware devices with market lifetimes that could span 10 years or more. Hypervisors also ensure that their new software versions and updates work with devices that were certified for prior software versions. Breaking “backwards compatibility” is not possible because it would disrupt the installed base.
- 4.8 Most new NIC, storage adapter, and FC HBA sales each year are of devices that have already been sold – and therefore certified – in previous years.¹²⁸ In other words, they are sales of devices that are already installed in existing servers. Since device drivers are not customer- or workload-specific, as the Phase 1 Decision accepts,¹²⁹ and the *same driver* covers devices in the installed base and new sales, the Combined Entity could not degrade interoperability with the new sales of existing products without directly impacting the installed base. The new sales of existing devices are therefore also protected from foreclosure.
- 4.9 **The Combined Entity could not degrade interoperability with I/O devices that are certified by equivalency.** Hardware vendors use already-certified device drivers in new devices. VMware certifies these new devices via a process of “*equivalency*”, which is a tick-box exercise with no product testing or troubleshooting with VMware, [X].¹³⁰ VMware established this process to reduce the administrative burden of certification, given the number of new devices with existing drivers. All VMware does is certify that the new device is functionally equivalent to a pre-existing certified device. Broadcom’s data show that **[90-100]%** of devices are certified in this way. The Combined Entity could not refuse equivalency requests outright because doing so would publicise VMware’s deviation from hardware neutrality and harm VMware’s reputation and value to enterprises. Nor could the Combined Entity refuse, delay, or hamper “*information exchange*” or “*technical support*”¹³¹ because no information or support is needed to be provided to vendors for equivalency certifications.
- 4.10 **There is no rational basis to degrade interoperability for new I/O devices requiring full certification,** which account for a negligible share of devices. Broadcom’s share of revenues from NICs, FC HBAs, and storage adapters requiring a full certification via

¹²⁸ Around [5-10]%, [0-5]%, and [5-10]% of NIC, storage adapter, and FC HBA sales each year, respectively, are of devices that were not already certified; see Compass Lexecon, No Incentive for the Merged Entity to Reduce VMware’s Interoperability with Rival Hardware (**Compass Lexecon Incentives Paper**), dated 10 May 2023, paragraphs 3.3 to 3.5, provided as **Annex ISR-028**.

¹²⁹ Phase 1 Decision, paragraph 250 and footnote 372.

¹³⁰ After initial certification, I/O devices do not generally need to seek certification with new vSphere releases because their certification is automatically carried forward. Any failure by the Combined Entity to carry forward certification would represent a radical deviation from past practice and endanger the interoperability of existing servers.

¹³¹ Issues Statement, paragraph 22(a).

VMware's VIVA process in 2021 were [0-5]%, [0-5]%, and [0-5]%, respectively.¹³² In theory, the Combined Entity could refuse or delay certification of rivals' devices, but it is apparent that it could not in practice.

- (a) ***VMware has never refused a request for certification through VIVA*** for a hardware device that meets the relevant product specifications. If the Combined Entity were to do so for the first time, it would immediately reveal its radical deviation from hardware neutrality and harm VMware's reputation and value to enterprises.
- (b) ***The VIVA process is largely automated***, which means there is minimal scope for partial foreclosure strategies. Broadcom estimates that [90-100]% of its VIVA certifications are completed without any support from VMware. The process is designed to minimise friction given the number of devices being certified.¹³³ [X]. This is what happens in the Linux environment, where there is no support provided from Linux for driver development.
- (c) ***Even significant delays in full certification will not affect rivals' sales***. The Commission has [X]. Such a foreclosure strategy is unrealistic and would be ineffectual.¹³⁴
 - (i) ***Hardware vendors certify new generation devices with VMware long before they are released***. Vendors have ample time to certify new device generations. For instance, Broadcom's latest generation FC HBA was first certified in [X] was ready for shipment to server OEMs in [X]. Moreover, if Broadcom were concerned about delays in obtaining certification from VMware before a server release date, it could begin the certification process earlier. To illustrate the point, for the latest generation of NICs and storage adapters, if Broadcom had chosen to begin certification with VMware at the same time as it did with Linux, it would have started engagement with VMware 11 months and 14 months earlier than it did, respectively. Accordingly, even delays of full certification for more than a year would have no effect on rivals' new generation sales.¹³⁵
 - (ii) ***New I/O devices are adopted slowly***. In the first year after launch, sales of new device generations represent fewer than [X]% of overall sales,

¹³² Compass Lexecon Incentives Paper, paragraph 4.7.

¹³³ As explained on slide 23 of VMware's Presentation to the Commission on Interoperability & Certification (provided to the CMA on 24 February 2023), each month, VIVA has [X] users, over [X] certifications, and over [X] page views.

¹³⁴ As with other putative foreclosure scenarios, there is no way for the Combined Entity to delay full certification without publicly deviating from the hardware neutrality (and perception thereof) that is so critical for VMware. It will be readily apparent to rivals that certification is taking longer than it has in the past, particularly if it means that a vendor misses a targeted server release window.

¹³⁵ Notably, the Combined Entity could not delay certification by more than a year without also publicly deviating from hardware neutrality. *See also* Compass Lexecon Incentives Paper, paragraph 4.5.

due to slow customer adoption.¹³⁶ Indeed, new device generations represent the minority of sales for several years after launch (around [X] years for NICs and more than [X] for storage adapters and FC HBAs).¹³⁷ Accordingly, already certified devices remain the predominant choice for customers in the short run.

- (iii) ***I/O devices are mature products with long lifecycles.*** Once certified, NICs, storage adapters, and FC HBAs will compete effectively in the market for up to a decade.¹³⁸ For example, based on when they were first certified, Microchip’s latest generation storage adapters will compete for at least the next [X] years, Marvell’s latest generation FC HBAs will compete for at least the next [X] years, and NVIDIA and Intel’s latest generation NICs will compete for at least the next [X] years.¹³⁹ So even if a rival’s new generation device were to miss a new server release window, it may have zero effect on rivals’ overall sales.
 - (iv) ***New I/O device generations do not typically feature on new server releases.*** OEMs adopt a so-called “[X]” rule, whereby they [X] with new generation servers. This is intended to minimise interoperability errors with other hardware components within a server. And even when an OEM [X] in its configurator for a new server, it would typically also offer numerous existing and established I/O devices on its configurator.¹⁴⁰
 - (v) ***No first-mover advantage.*** Sales of a new device generation depend on industry adoption of the generation, not the release date. This is evident from Broadcom’s release of its 16 Gbps FC HBA in 2011, which [X].¹⁴¹
 - (vi) ***OEMs can and do add I/O devices to their configurators when they are certified.*** Even if a rival hardware device were to miss a new server release, there would be minimal impact on new device sales because only [X] of new server sales coincide with release and OEMs could and would include I/O devices as an option when they are certified.¹⁴² [X].
- (d) ***Devices in VMware’s Compatibility Guide are fully interoperable.*** When VMware certifies a device as compatible, this means it interoperates with

¹³⁶ Only [X]% of Broadcom’s sales of NICs between 2019 and the first half of 2021 related to new device generations that were made in the first year after launch; [X]% for FC HBAs and [X]% for storage adapters (see Compass Lexecon Incentives Paper, paragraph 3.6).

¹³⁷ See Compass Lexecon Incentives Paper, Figure 1.

¹³⁸ As explained in the response to question 4 of Phase 2, RFI 1, Broadcom and Marvell have FC HBAs in the market that will compete at least until [X]; Microchip and Marvell have storage adapters that will compete at least until [X], and Intel and NVIDIA have NICs in the market that will compete at least until [X].

¹³⁹ As explained in the response to question 4 of Phase 2, RFI 1.

¹⁴⁰ See, e.g., Dell server configurator at paragraph 3.19.

¹⁴¹ See SO Response, Figure 43.

¹⁴² [X]. See [Annex ISR-014](#).

VMware just as well as any other device does. VMware certification status is binary; there is no way for VMware to certify certain drivers to a “lesser” or “greater” extent.

- 4.11 **The Combined Entity could not degrade interoperability for new I/O devices in non-VMware environments.** As shown in **Figure 4** below, the vast majority of I/O devices are used in non-virtualised servers or with virtualisation software other than vSphere.¹⁴³

Figure 4: Vast majority of NICs, storage adapters, and FC HBAs are used in non-VMware servers

[X]

- 4.12 Since vSphere is not present on c. [90-100]% of servers,¹⁴⁴ the Combined Entity could not leverage “*VMware’s position in server virtualisation software*” to foreclose hardware rivals’ sales into those servers. The Phase 1 Decision speculates, however, that “*at least some*” server OEMs may choose to extend the putative foreclosure strategy beyond VMware servers by selecting only Broadcom hardware “*even in cases where the server is not ultimately used with*” vSphere.¹⁴⁵ There are four main flaws in this theory:

- (a) ***OEMs would resist a Broadcom monopoly.*** As explained in Section 3 above, OEMs multi-source hardware components to provide choices to enterprises, to foster competition between suppliers, and to ensure security of supply. Even if having “*different server configurations for different OSs would incur additional costs for OEMs*”, which Broadcom does not accept,¹⁴⁶ these costs would be negligible in comparison to the costs of sourcing exclusively from one I/O device component supplier.
- (b) ***It is inconsistent with the evidence in the Phase 1 Decision.*** As noted above, “[t]hird parties [...] emphasised the importance of multiple supplier options, with one server OEM stating that it is important to have multiple suppliers of the components in order to obtain the best pricing.”¹⁴⁷ Another server OEM indicated that it “*might continue to use rival hardware [if] a sufficient number of customers request such servers*”, which the Parties consider to be likely given the small share of VMware servers.¹⁴⁸ Only a single server OEM said that it “*would stop or significantly reduce purchases from Broadcom rivals [to]*

¹⁴³ See Compass Lexecon Incentives Paper, Figures 8 to 10; Broadcom’s Site Visit Presentation, slide 88.

¹⁴⁴ **Annexes RFI4Q5-001 - RFI4Q5-004** to the FMN.

¹⁴⁵ Phase 1 Decision, paragraph 216.

¹⁴⁶ As discussed in Section 3 above, OEMs must offer choice to enterprises to compete effectively. As such, “*testing and qualifying different products*” (see Phase 1 Decision, paragraph 21) is a routine and streamlined process for them. Indeed, Dell and Lenovo offer more than 10 options of FC HBAs and storage adapters in their server configuration, spanning from models released in 2013 to models released in 2022.

¹⁴⁷ Phase 1 Decision, paragraph 218(a).

¹⁴⁸ Phase 1 Decision, paragraph 217.

ensure all customers have the option to use VMware” (emphasis added).¹⁴⁹ But there is no explanation of the server OEM’s decision nor any evidence on the proportion of rivals’ sales to non-VMware environments that would be impacted by this OEM’s decision.

(c) *It is inconsistent with* [redacted].

(d) *It is inconsistent with* [redacted].¹⁵⁰ [redacted].¹⁵¹ [redacted].¹⁵²

4.13 **Foreclosure would accordingly not harm Broadcom’s hardware rivals.** Given the small segment of I/O devices that could *potentially* be foreclosed (as shown in **Figure 4** above), almost all demand for NICs, storage adapters, and FC HBAs would be unaffected by any foreclosure strategy. As a result, the Combined Entity could not “deprive them of a substantial volume of sales.”¹⁵³ Given the size of the rivals concerned, it is not plausible to contend that degrading interoperability with vSphere would “harm rivals’ competitiveness.”¹⁵⁴ Indeed, for NICs, Intel and NVIDIA have FY2021 market shares of [redacted]% and [redacted]%, and revenues of US\$[redacted] million and US\$[redacted] million, respectively. For storage adapters, Dell and Microchip have FY2021 market shares of [redacted]% and [redacted]%,¹⁵⁵ and revenues of US\$[redacted] million and US\$[redacted] million, respectively. For FC HBAs, Marvell has a FY2021 market share of [redacted]% and revenues of US\$[redacted] million.

4.14 **In any event the Combined Entity could not degrade interoperability with Microchip and Marvell.** [redacted]¹⁵⁶. [redacted]¹⁵⁷:

(a) [redacted],¹⁵⁸ and

(b) [redacted].¹⁵⁹

4.15 [redacted]:

(a) [redacted],¹⁶⁰ and

¹⁴⁹ *Ibid.*

¹⁵⁰ [redacted], paragraph 8.

¹⁵¹ [redacted], paragraph 26.

¹⁵² *Ibid.*

¹⁵³ MAGs, paragraph 7.33(c).

¹⁵⁴ MAGs, paragraph 7.33.

¹⁵⁵ The market share for Microchip includes the share of competitor boards using Microchip controllers, which the Phase 1 Decision considers to be the appropriate metric (*see* paragraph 282). The Parties disagree for the reasons explained elsewhere (*see* slide 37 of Broadcom’s Site Visit Presentation).

¹⁵⁶ [redacted].

¹⁵⁷ The Commission dropped concerns at Phase 2 in relation to NICs.

¹⁵⁸ [redacted].

¹⁵⁹ [redacted].

¹⁶⁰ [redacted].

(b) [REDACTED].

4.16 [REDACTED].¹⁶¹

The Merged Entity would have no incentive to foreclose hardware rivals

4.17 The CMA’s Guidelines state that a “merged entity may be more likely to pursue a [particular foreclosure strategy] if its business strategy involves this approach, it has a history of doing this with similar products or its deal rationale involves plans to do so post-merger.”¹⁶² None of these considerations applies in this case. On the contrary, Broadcom’s business model is based on independent, sustainable franchises.¹⁶³ VMware’s business model is also based on ensuring hardware neutrality.¹⁶⁴ Like other hardware vendors in the datacentre, Broadcom has ensured and must ensure interoperability with rival hardware products.¹⁶⁵ And in contrast to a strategy of degrading interoperability, Broadcom’s documented strategy and investment depends on increasing deployment and utilisation of VMware products which would be fatally undermined if it were to restrict enterprise customers’ hardware choices.¹⁶⁶ It is therefore unrealistic to suggest that the Combined Entity might attempt to degrade interoperability with its hardware rivals.

4.18 Aside from this crucial evidence in relation to Broadcom’s actual behaviour, the exercise of inferring the Combined Entity’s “behaviour from their financial incentives”¹⁶⁷ also shows clearly that degrading interoperability would make no commercial sense.

4.19 **vSphere’s margins are “much higher” than Broadcom’s margins on I/O devices,** as the Phase 1 Decision accepts.¹⁶⁸ Broadcom’s average incremental margin per server for NICs, storage adapters, and FC HBAs that would be gained from foreclosure amounts to just \$[REDACTED], which compares to VMware’s average vSphere margin of \$[REDACTED] per server.¹⁶⁹ This disparity means that the critical switching rate for the foreclosure strategy to be profitable is “high” as the Phase 1 Decision acknowledges.¹⁷⁰ More specifically, the Combined Entity would need at least 97% of affected customers to

¹⁶¹ [REDACTED].

¹⁶² MAGs, paragraph 7.34(a).

¹⁶³ As explained in paragraph 2.1 above.

¹⁶⁴ As explained in paragraph 3.12 above.

¹⁶⁵ As explained in paragraph 3.9 above.

¹⁶⁶ As explained in paragraph 2.1 above.

¹⁶⁷ MAGs, paragraph 7.34(a).

¹⁶⁸ Phase 1 Decision, paragraph 206.

¹⁶⁹ Compass Lexecon Incentives Paper, Table 1 and paragraph 4.7.

¹⁷⁰ Phase 1 Decision, paragraph 206.

choose Broadcom hardware for their new devices, rather than switching away from VMware to an alternative on-premises or public cloud solution.¹⁷¹

- 4.20 **Potential hardware gains from total foreclosure are minimal whereas potential vSphere losses are huge.** If, in response to a foreclosure strategy targeting all new devices,¹⁷² 100% of VMware customers that would otherwise have purchased non-Broadcom hardware were to switch to Broadcom for their new I/O devices in year one, the Combined Entity would gain \$[X] million (as shown in **Figure 5** below). The maximum gain in year one is so low because of (1) the small share of devices that could be foreclosed (which excludes new sales of already-certified devices)¹⁷³ and (2) the slow roll out of new I/O devices.¹⁷⁴ Even if, as the Phase 1 Decision suggests, the “majority of VMware’s customers [using rival hardware] would prefer to switch to Broadcom’s hardware” when “faced with the loss of interoperability”,¹⁷⁵ the Combined Entity could still lose more than \$[X] million in year one.¹⁷⁶ If only between [10-20]% and [20-30]% of affected customers were to switch to Broadcom hardware in year one, as VMware customers responding to the 2023 MIT survey suggest,¹⁷⁷ the Combined Entity could lose **between \$[X] million and \$[X] million** in year one.¹⁷⁸ Confronted with these potential gains and losses, no company would consider the foreclosure strategy.

Figure 5: Profit/loss from degrading interoperability between vSphere and rival I/O devices¹⁷⁹

[X]

- 4.21 **The asymmetry of gains and losses would persist over the longer term.** The maximum gain from foreclosure over the long term is c. \$[X] million per year, which

¹⁷¹ The critical switching rate for year one (97%) differs from the rate in the longer term (96%) due to variation in the relative adoption rates of NICs, FC HBAs, and storage adapters the short run and the long run (see Compass Lexecon Incentives Paper, paragraphs 2.7 and 2.8).

¹⁷² This includes, conservatively, devices certified by equivalency, which, as noted elsewhere, cannot be foreclosed (see paragraph 4.9).

¹⁷³ As explained above, the Combined Entity could not foreclose devices using already-certified drivers without impacting the installed base of existing servers. The Parties consider that the maximum gain figures provided overstate how much the Combined Entity might gain in practice, because they include, conservatively, devices certified by equivalency as potentially being foreclosed. As explained at paragraph 4.9 above, the Combined Entity could not degrade interoperability for these devices in practice.

¹⁷⁴ Around [X] of NICs, FC HBAs, and storage adapters shipped by Broadcom are devices which have been launched within the past [X] years, with the remaining [X] of shipped devices launched more than [X] years before. See Compass Lexecon Incentives Paper, paragraph 3.6.

¹⁷⁵ Phase 1 Decision, paragraph 208.

¹⁷⁶ If [X]% of affected servers were to switch to Broadcom hardware, the Combined Entity would lose \$[X] million in the first year (see Compass Lexecon Incentives Paper, paragraph 3.12).

¹⁷⁷ RBB, Estimation of actual switching rates based on the MIT Switching survey (RBB Switching rate paper on 2023 MIT survey) provided as **Annex ISR-002**.

¹⁷⁸ See Compass Lexecon Incentives Paper, Figure 5.

¹⁷⁹ The estimated actual switching rate of [X]% presented on slide 86 of Broadcom’s CMA Site Visit Presentation has been updated in Figure 5 following refinements to the incidence of new device sales in the RBB Switching rate paper on the 2023 MIT survey.

would take more than seven years to materialise. If the majority of customers were to switch away from VMware, the Combined Entity could lose at least \$[X] billion per year.¹⁸⁰ If only [10-30]% of customers were to switch to Broadcom hardware, the Combined Entity would lose between \$[X] billion and \$[X] billion per year.¹⁸¹ Moreover, a prolonged foreclosure strategy would increase the likelihood that customers not directly affected by foreclosure would switch away from VMware due to its deviation from neutrality.¹⁸² It would also afford enterprises more time to switch to the public cloud, which the Phase 1 Decision suggested the majority of customers might consider only for some of their workloads in the “*short run*”.¹⁸³

4.22 **Minimal switching away from VMware would render total foreclosure unprofitable.** Given the small gain that might be realised by selling more hardware relative to the much-higher value vSphere revenue, a foreclosure strategy would be unprofitable if VMware loses only a small share of workloads to rivals. To illustrate how fragile a foreclosure strategy would be, any of the following scenarios would render the strategy unprofitable:

- (a) Foreclosure would be unprofitable in the first year if *any* of the following occurs:
 - (i) Hyper-V’s annual growth rate increases by only [X]% at VMware’s expense;
 - (ii) Hyper-V captures just [X]% market share from VMware, increasing its market share from [X]% to [X]%;
 - (iii) IBM/Red Hat captures just [X]% market share from VMware, increasing its market share from [X]% to [X]%;
 - (iv) IBM/Red Hat’s annual growth rate increases by only [X]% at VMware’s expense;
 - (v) The public cloud annual sale growth rate increases by just [X]%, at VMware’s expense; or
 - (vi) The containerization annual sales growth rate increases by only [X]%, at VMware’s expense.
- (b) Foreclosure would be unprofitable even once new potentially-foreclosable devices finally replace nearly all existing devices, if *any* of the following occurs:

¹⁸⁰ The loss of \$[X] billion assumes [X]% of affected servers switch away from VMware virtualisation software. This loss would be even greater if additional servers moved away from VMware.

¹⁸¹ Compass Lexecon Incentives Paper, Figure 7.

¹⁸² Open-ended responses in the 2023 MIT survey show that respondents have a strong preference for hardware neutrality and caution against a foreclosure strategy. For example: “*The compatibility problem is definitely critical for non-Broadcom users*”; “*If VMware forces HW only from Broadcom it could be a [sic] error*”; “*Hardware vendor lock-in would result in a complete exit of VMware to alternative providers*”.

¹⁸³ Phase 1 Decision, paragraph 61(a).

- (i) VMware retains less than [X]% of affected servers;
- (ii) Just [X]% of existing workloads leave VMware;
- (iii) Hyper-V's annual sales growth rate increases by only [X]%, at VMware's expense;
- (iv) IBM/Red Hat captures just [X]% market share from VMware, increasing its market share from [X]% to [X]%;
- (v) Hyper-V captures only [X]% from VMware, increasing their market share from [X]% to [X]%;
- (vi) IBM/Red Hat's annual sales growth rate increases by only [X]%, at VMware's expense;
- (vii) The public cloud annual sales growth rate increases by just [X]%, at VMware's expense; or
- (viii) The containerization annual sales growth rate increases by only [X]%, at VMware's expense.

4.23 Given the numerous credible alternatives available to enterprises on-premises and in the public cloud, it is inevitable that the foreclosure strategy would be unprofitable.

4.24 **Partial foreclosure is even less likely to be profitable, because maximum gains from partial foreclosure are negligible and do not increase over time.** The only realistic method of partial foreclosure is the delay of full certification for new devices with new drivers. As explained above, existing devices and new sales of existing devices cannot be partially foreclosed without causing disruption to the installed base, and equivalency certifications cannot be partially foreclosed because there is no information or support to hamper or delay.¹⁸⁴ Since partial foreclosure cannot be achieved secretly,¹⁸⁵ it raises the same reputational risks for the Combined Entity as total foreclosure. But it has an even smaller potential upside. Since VMware is not involved in [X]% of new device certifications,¹⁸⁶ fewer than 500 devices – representing revenue of **c. \$[X] per year** – are potentially subject to partial foreclosure.¹⁸⁷ Unlike in a total foreclosure theory, the number of new devices potentially subject to partial foreclosure would not increase over time.¹⁸⁸

4.25 **Gains and losses from foreclosure must be assessed over the same time horizon.** The Phase 1 Decision criticised the Parties' incentives analysis for failing to take

¹⁸⁴ See paragraphs 4.8 and 4.9 above.

¹⁸⁵ See paragraphs 4.6(a) and 4.10(c) above.

¹⁸⁶ There is no engagement with VMware for certifications by equivalency or for the vast majority of certifications through the VIVa process (see paragraphs 4.9 and 4.10).

¹⁸⁷ The 500 devices potentially subject to partial foreclosure represent the portion of new devices per year that may require full certification and engagement with VMware (see Compass Lexecon Incentives Paper, paragraph 4.6).

¹⁸⁸ Compass Lexecon Incentives Paper, paragraph 4.14 *seq.*

account of longer-term dynamic considerations.¹⁸⁹ However, it dismisses the constraint from containers and the public cloud on the basis of third-party evidence indicating that they are not alternatives for *all* their workloads in the *short run*.¹⁹⁰ And it contends that enterprises face significant switching costs when moving away from VMware “*within a short time-horizon*”.¹⁹¹ But the relevant period for assessing the gains to the merging parties against the losses to rivals must be the same.¹⁹²

4.26 **OEMs and rivals would neutralise and retaliate against any efforts to reduce interoperability.** OEMs would retaliate harshly and swiftly against the Combined Entity if it were to degrade interoperability and thereby undermine OEMs’ multi-sourcing business models, their competitiveness, and profitability. With respect to rivals, the Phase 1 Decision was incorrect to contend that rivals would not retaliate,¹⁹³ which is evident from third-party evidence on the CMA file which cites “*retaliation as a possible strategy*”.¹⁹⁴ In more detail:

- (a) ***OEMs could neutralise any attempt to reduce interoperability at little to no cost.*** The Phase 1 Decision makes the unsupported assertion that an OEM’s “*retaliation strategy would incur additional costs to server OEMs.*”¹⁹⁵ But this fails to appreciate that OEMs could remove Broadcom hardware from their server configurators (or threaten to do so) – either generally or for non-VMware environments specifically – at near-zero cost. OEMs offer a large number and wide variety of hardware combinations in their server configurators. The hardware choices that are made available to customer can and do vary depending on a customer’s selections.¹⁹⁶ It would therefore be trivial for an OEM to remove, for example, Broadcom I/O devices as an option for customers that select to preload a rival hypervisor. Doing so for a small fraction of servers would neutralise the foreclosure strategy.¹⁹⁷ Server OEMs could also change the pre-loaded hypervisors in their configurators to remove VMware or select alternatives as defaults.
- (b) ***Any of Broadcom’s main OEM customers could easily erase the maximum gains from foreclosure.*** Broadcom generates substantial annual profits from selling hardware to Cisco (c. \$[<]), Dell (\$[<]), HPE (\$[<]), Lenovo (\$[<]).¹⁹⁸ They would only have to switch a small amount of purchases away from Broadcom to rivals (<[0-5]% for Cisco, Dell, and HPE, and <[5-10]% for

¹⁸⁹ Phase 1 Decision, paragraph 220.

¹⁹⁰ Phase 1 Decision, paragraphs 62 and 152.

¹⁹¹ Phase 1 Decision, paragraph 169.

¹⁹² See, e.g., *Microsoft/Activision Blizzard*, Addendum to the Provisional Findings, paragraph 4.9.

¹⁹³ Phase 1 Decision, paragraph 224.

¹⁹⁴ Phase 1 Decision, paragraph 224.

¹⁹⁵ Phase 1 Decision, paragraph 218(a).

¹⁹⁶ For example, some Dell servers may be configured for NVMe drives if a certain type of storage controller is added: See, e.g., Dell, [PowerEdge R6625 Rack Server Configurator](#).

¹⁹⁷ See Compass Lexecon Incentives Paper, Table 3.

¹⁹⁸ These figures reflect Broadcom’s total profits from these OEMs, not for I/O devices specifically.

Lenovo), which their multi-sourcing strategy enables them to do, to make foreclosure unprofitable.¹⁹⁹

- (c) ***Foreclosing rival storage adapter vendors would require the Combined Entity to foreclose server OEMs that manufacture storage adapters.*** OEMs compete with Broadcom in storage adapters by building their own boards with controllers from Broadcom, Microchip and/or Marvell. Broadcom sells [X] storage controllers as storage adapters. OEMs can and do handle the certification of their own storage adapters. While OEMs may rely on hardware supplier's drivers and assistance, they also have their own engineers and resources to build their own drivers, which they routinely do when certifying their servers. It is inconceivable that the Combined Entity would attempt to reduce interoperability vis-a-vis OEMs – their main customers and counterparts – thereby causing not only a loss of reputation, but also retaliation by server OEMs.
- (d) ***Broadcom's rivals in Ethernet NICs are the dominant suppliers of CPUs and GPUs, on whose support VMware relies.*** Any attempt by the Combined Entity to foreclose rival NIC vendors would be particularly futile given the identity and strength of those rivals. Broadcom has only [10-20]% share of NICs sales. The market is dominated by NVIDIA with a [40-50]% share and Intel with [30-40]%. AMD also competes with a [0-5]% share. Those same companies are also the dominant suppliers of CPUs and GPUs, the brains of the server and the most important component. Approximately 95% of x86 servers are sold with Intel and AMD CPUs, and NVIDIA is the dominant GPU supplier with nearly a 100% share. Since the vast majority of vSphere deployments run on Intel's and AMD's x86 CPUs and NVIDIA's GPUs, VMware is today – and Broadcom will be post-Transaction – entirely dependent on the support of these CPU and GPU suppliers to develop its virtualization products. A lack of interoperability with these products would be fatal to VMware's business as it would effectively be shut out of the server market and unable to sell vSphere. Indeed, for these reasons, the Commission has previously recognized in a similar context that Intel's and AMD's market positions in CPUs would “*most likely eliminate all incentives by [their rivals] to degrade interoperability.*”²⁰⁰
- (e) ***Hardware rivals could also retaliate in other ways, including:***
- (i) ***By publicising the Combined Entity's foreclosure attempts.*** Given their extensive experience developing and certifying drivers, rivals would quickly detect any attempt to deny or degrade interoperability.²⁰¹ At almost no cost, they could publicise the foreclosure attempt and hurt Broadcom's and VMware's reputations.

¹⁹⁹ See Broadcom's Site Visit Presentation, slide 92 and Compass Lexecon Incentives Paper, paragraph 5.24.

²⁰⁰ See also *NVIDIA/Mellanox*, Case COMP/M.9424, Commission decision of 19 December 2019, paragraph 242: The Commission considered that the risk of Intel and AMD degrading the interoperability of their CPUs with the NVIDIA/Mellanox entity's NICs eliminated the entity's incentive to degrade interoperability, particularly given that “*the vast majority of NVIDIA's GPUs and Mellanox's network interconnect products sales depend on being interoperable with Intel's and AMD's CPUs*”.

²⁰¹ [X] would also be able to detect any attempt to degrade interoperability [X].

- (ii) ***By leveraging their supply relationships with Broadcom.*** In addition to competing with Broadcom in storage adapters, Microchip supplies Broadcom with [§]. Broadcom sources products [§].²⁰² Similarly, Marvell, Broadcom’s main rival in FC HBAs, [§].²⁰³

4.27 **Other Phase 1 criticisms of the Parties’ incentives analysis do not alter the conclusion that foreclosure would be unprofitable.** The criticisms can be addressed as follows:

- (a) ***VMware would not recoup losses if workloads migrate to the cloud.***²⁰⁴ It would be illogical for VMware’s customers to respond to foreclosure by moving from one VMware product to another. In any event, VMware Cloud (VMC) captures [§].²⁰⁵ Customers perceive VMC as [§]²⁰⁶ and a temporary solution while transitioning to public cloud. One customer noted that: [§].²⁰⁷ Another noted, [§].²⁰⁸ Moreover VMC has consistently underperformed, failing to meet projections even after the company adjusted its forecasts downward. Strategy documents from 2019 forecasted that VMC would achieve \$[§] in annual bookings in FY2022. This forecast was adjusted downwards in VMware’s FY2020 corporate long-range plan, which predicted VMC would book a more conservative \$[§] revenue in FY2022.²⁰⁹ In fact, VMC achieved only \$[§] in bookings in FY2022, substantially missing both projections. Recognizing this underperformance, [§].²¹⁰ Even then, VMC continued to miss its revised quarterly and annual bookings projections, leading VMware to lower even further VMC’s projections in the FY2023 AOP by [§]%. Separately, the Phase 1 Decision also fails to account for the fact that Broadcom customers use less Broadcom hardware when they switch to the public cloud, such that the Combined Entity would also not be able to recoup losses through additional hardware sales.²¹¹
- (b) ***Customers could face non-trivial barriers to switch to Broadcom hardware.*** The Phase 1 Decision indicates that “switching costs for hardware are ‘limited’, with switching ‘made easier by standards and protocols’”.²¹² This is based on a misreading of the Parties’ submissions, which relate to OEMs’ ability to switch between alternative suppliers given that they “already multi-source from

²⁰² A Broadcom global enterprise risk management deck notes that [§]. See [Annex ISR-029](#), slide 11.

²⁰³ See Broadcom internal document, email exchange with subject [§]. See [Annex ISR-030](#).

²⁰⁴ Phase 1 Decision, paragraph 213.

²⁰⁵ SO response, paragraph 473.

²⁰⁶ VMware internal document, email exchange with subject [§]. See [Annex ISR-031](#).

²⁰⁷ VMware internal document, [§]. See [Annex ISR-025](#).

²⁰⁸ *Ibid.*

²⁰⁹ VMware internal document, [§]. See [Annex ISR-032](#).

²¹⁰ In addition to the long-range plans (LRPs) that project annual bookings going forward for the next three fiscal years, in advance of each fiscal year VMware prepares an Annual Operating Plan (“AOP”) that makes an updated projection of bookings for the coming fiscal year.

²¹¹ Less than 10% of CSPs’ servers use storage adapters and less than 1% use FC HBAs.

²¹² Phase 1 Decision, paragraph 209.

different suppliers.”²¹³ The same is not true for all enterprise customers. Indeed, switching hardware may be challenging and expensive. For instance, certain enterprise users of FC HBAs write custom scripts to simplify and automate their server builds. A custom script written for a Marvell FC HBA installation would not work with a Broadcom FC HBA. There are examples of Broadcom failing to win sales as a result of this barrier to switching (e.g., [§<]).²¹⁴ Some hardware customers also enter into framework agreements which may tie them exclusively to a particular vendor for several years.²¹⁵

- (c) ***The use of FC HBAs and FC switches together does not reinforce the Combined Entity’s incentives to pursue a foreclosure strategy.*** Broadcom does not bundle FC HBAs and FC switches together. End customers have different purchasing cycles for FC switches and FC HBAs. They are not typically purchased together.²¹⁶ The FC protocol also ensures that FC HBAs and FC switches from different vendors interoperate fully together. Moreover, customers build FC SAN networks with one vendor’s FC switches due to the ability to manage them with common software. Thus enabling foreclosure in this manner would require a customer replacing all of its existing Cisco FC switches, which would often be a [§<] exercise with no benefit for the customer.
- (d) ***Rivals would maintain sufficient scale to compete.*** Even if it were plausible, a foreclosure strategy would have minimal impact on rivals’ sales. Overall, the markets for NICs, FC HBAs, and storage adapters are mature, with long product lifecycles and slow adoption rates for new generation products. Virtually no technical engagement with VMware is required. Accordingly, existing players have VMware-certified products in the market that will compete effectively for the next decade.²¹⁷ Furthermore, the investment required for the next generation of devices is [§<] relative to total revenues over the lifecycle of the previous generation. Since fixed costs represent [§<], rivals would remain substantially above any realistic estimate of minimum efficient scale.²¹⁸ Even under total foreclosure, and assuming no countervailing reactions from customers, OEMs and rivals, [90-100]% of demand for NICs, [90-100]% for

²¹³ FMN, paragraph 20,131.

²¹⁴ [§<] buys around [§<] FC HBAs a year, [§<]% of which are Marvell FC HBAs ([§<]) has expressed to Broadcom that [§<] (see [Annex ISR-004 Appendix 2](#)).

²¹⁵ For example, [§<] purchases about [§<] FC HBAs a year from Marvell. Broadcom’s internal documents show that [§<] (see [Annex ISR-004 Appendix 6](#), p. 2) and that [§<] agreed to a framework agreement to purchase [§<] (see [Annex ISR-004 Appendix 7](#), p. 2).

²¹⁶ See also the Commission decision in *Broadcom/Brocade*, Case COMP/M.8314, Commission decision of 12 May 2017, paragraph 223. The replacement cycle of FC HBAs is driven by the server replacement cycle. But procurement of FC SAN switches is largely driven by end customers' storage arrays requirements and replacements, which typically follows a longer life cycle than the servers.

²¹⁷ See the Parties’ response to RFI of 25 April 2023, question 4(b).

²¹⁸ As explained in the response to question 4 of Phase 2, RFI 1, In 2021, Broadcom’s fixed costs for FC HBAs were around [§<]% of its revenue, for storage adapters and NICs around [§<]% of revenue, and for FC switches, around [§<]% of revenue.

storage adapters, and [90-100]% for FC HBAs, would be *unaffected* by any foreclosure strategy in the first year.²¹⁹

- (e) ***The Combined Entity would be unable to raise hardware margins.*** The Phase 1 Decision’s speculation that the Combined Entity may “*be able to raise hardware margins*”²²⁰ is unsubstantiated and inconsistent with market reality. Since the Combined Entity cannot degrade interoperability for hardware devices in non-VMware environments, or for existing devices, or devices certified by equivalency in VMware environments (as discussed above),²²¹ these devices would continue to compete with the tiny proportion of (supposedly) foreclosable new devices with new drivers. This competition would prevent Broadcom from increasing margins, particularly given the slow adoption rate of new devices.

4.28 In summary, the Transaction will not result in an SLC under ToH1 for storage adapters, FC HBAs, or NICs. The evidence clearly shows that the Combined Entity would not have the ability or incentive to degrade interoperability. Even if it did attempt to foreclose rivals, such a strategy would be bound to fail: it would have no effect on rivals or competition, and would be ruinous to Broadcom’s business and reputation.

(B) No SLC in FC switches

- 4.29 The Issues Statement posits that the Combined Entity could leverage vSphere’s “*position in server virtualisation software*” to impair “*access to VMware’s API for competitors’ FC switches*”.²²² It contends that the Combined Entity could foreclose rivals by reducing “*VMware’s engagement in bilateral discussions which facilitate the implementation of VMware’s APIs, avoid[ing] or delay[ing] making any new APIs available to rivals [and/or] choos[ing] not to make APIs public in the future.*”²²³
- 4.30 ToH1 in relation to FC switches fundamentally misunderstands the significance of VMware APIs in the operation of FC switches. Contrary to the evidence of the single third party on which the Phase 1 Decision relies,²²⁴ FC switches do not interact at all with vSphere and do not need access to VMware APIs to operate. Moreover, the vCenter APIs that FC switch management software calls on are not specific to FC switches, but rather are generic APIs called by hundreds of other types of software.

The Combined Entity has no ability to foreclose FC switch vendors

- 4.31 **FC switches do not interact with vSphere.** FC switches communicate with servers and storage devices via FC HBAs only, using the industry-standard FC protocol. It is

²¹⁹ Compass Lexecon Incentives Paper, paragraph 5.14. These figures assume that devices certified by equivalency cannot be foreclosed for the reasons set out at paragraph 4.9.

²²⁰ Phase 1 Decision, paragraph 220.

²²¹ See paragraph 4.9.

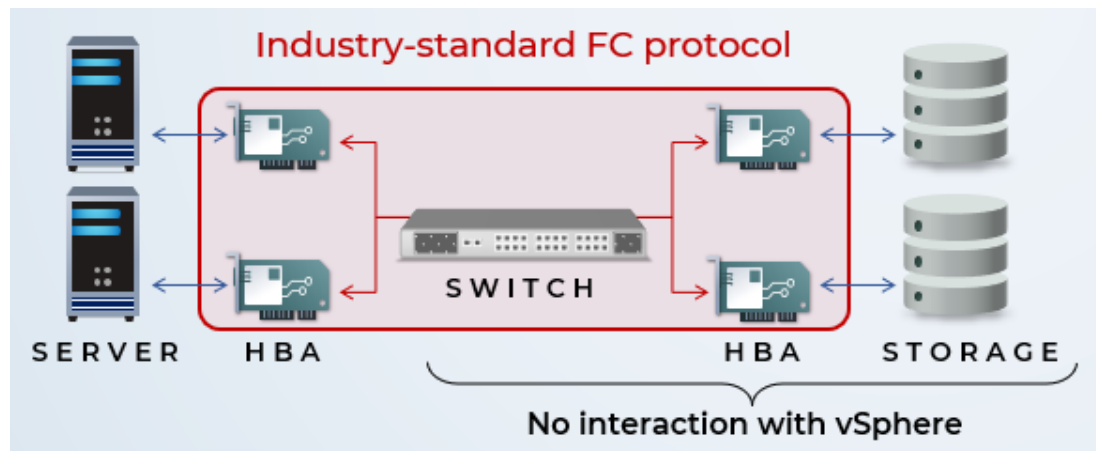
²²² Issues Statement, paragraphs 22-23.

²²³ Issues Statement, paragraph 22(b).

²²⁴ Phase 1 Decision, footnotes 223-224.

the FC HBA that interoperates with vSphere in VMware-virtualised servers (as shown in **Figure 6** below) via a device driver, not the FC switch.

Figure 6: FC switches do not interact with vSphere



- 4.32 **VMware’s vCenter provides public, standardised APIs that can be used by datacentre hardware management software.** VMware offers public APIs in vCenter – the vSphere management component – to provide inventory information on VMs that are configured using vCenter.²²⁵ These APIs can be used by any customer or partner, and by software that wants access to the information, including monitoring software (e.g., Splunk, Zabbix, Ansible),²²⁶ enterprise management software (e.g., [X]),²²⁷ storage and server OEM management software (e.g., Dell’s Cloud IQ, IBM Storage Insights, HPE Store Virtual Storage),²²⁸ and FC switch management software (e.g., Broadcom’s SANnav, Cisco Fabric Manager).
- 4.33 **VMware has no ability to identify the function or purpose of the software using its APIs.** They are standardised and are not specific to any type or category of hardware. The APIs do not actively collect information on the third-party software calling on it (e.g., whether it is Dell, Cisco, or Broadcom software), and VMware does not monitor the creation or use of software interfacing with these APIs (nor does it seek to do so). Plugins can also be used to extend the vCenter interface with additional features, such that the APIs could potentially be used for a wide range of purposes by a wide range of software (which obviously VMware cannot detail because it does not have visibility).
- 4.34 **Neither FC switches nor FC switch management software requires information from VMware to operate.** All FC switches operate without any direct or indirect access to information from VMware. While some FC switches are managed by software that collects information from VMware (e.g., on ESXi hosts, CPU, memory, network and disk utilization), that information is neither needed nor used for FC switch configuration or management (i.e., allowing and controlling communications and

²²⁵ 95% of VMware customers configure virtual machines used vCenter (see VMware Site Visit Presentation, slide 81).

²²⁶ Splunk, [Splunk Infrastructure Monitoring Features](#), Zabbix, [Zabbix 6.4](#), Red Hat Ansible, [Integration: Ansible and LogicMonitor](#).

²²⁷ Broadcom understands that [X] built and uses its own datacentre management and monitoring software.

²²⁸ See the Parties’ response to the RFI of 25 April 2023, Question 8.4-8.5.

managing and monitoring the fabrics, switches, switch ports, and server and storage ports on the FC SAN).

- 4.35 **Any VM-specific information that can be collected by FC switch management software is already separately available to customers through vCenter.** The VM-related information that is available via public APIs is also available directly through the vCenter client (as shown in **Figure 7** below), which is vCenter’s own management and monitoring portal. Access to the public API requires only log-in credentials for a vCenter account (a username, password and server ID). The information may also be available from FC HBA or storage vendors, or third-party monitoring software. Broadcom provides the information in SANnav as a convenience, [§].

Figure 7: VMware API provides data that is already available to customers

[§]

- 4.36 **Very few FC switch customers collect information about VMs.** Broadcom estimates that [§]% of FC switch customers do not collect information at the virtual level.²²⁹ This is because the information collected via the VMware public APIs are not required for the management of the FC SAN). The proportion of support calls that Broadcom receives relating to VM-related information illustrates the limited use of this type of information to FC switch customers. Brocade Storage Networking (**BSN**), the Broadcom division that supplies FC switches, received over [§] support calls from 2022 to date.²³⁰ Of these, only [§] concerned SANnav, and only [§] ([§]%) mentioned a detail related to VMware products or any VM-related information.
- 4.37 **No opportunity for partial foreclosure strategies.** Because FC switches can operate without any access to VMware’s APIs, any foreclosure strategy (whether total or partial) would have no effect. In any event, contrary to the suggestion in the Issues Statement, there is no engagement between FC switch manufacturers or software developers and VMware that could be hampered or delayed. Broadcom does not discuss implementation of the public APIs with VMware, which is evident from the fact that no VMware personnel were even aware that Broadcom’s SANnav was using VMware APIs. Broadcom developed its “*management software products*”²³¹ without any contact with VMware and believes this to be [§]. Insofar as VMware is aware, it does not discuss implementation of the APIs with other vendors.

The Combined Entity has no incentive to degrade APIs for third-party FC switch management software

- 4.38 **VMware cannot target new FC switches.** Customers build FC SAN fabrics with only one vendor’s FC switches, due to the ability to manage them with common software. Even if it were possible to interfere only with rivals’ API access, the same management software is used to control all of a customer’s FC SAN switches. Interfering with its API use would therefore impact both a customer’s existing and new FC SAN switches.

²²⁹ Broadcom [§].

²³⁰ As explained in the response to question 6 of Phase 2 RFI 1.

²³¹ Phase 1 Decision, paragraph 192.

Thus, adopting such a strategy would force a customer to replace all of its existing Cisco FC switches, which would in many cases cost [X], with no benefit for the customer.

- 4.39 **VMware cannot selectively interfere with FC switches.** As noted above, VMware’s public APIs are standardised: they are not specific to FC switches or even FC networks. The same APIs are used by VMware customers using enterprise management software, storage and server OEMs, as well as other third-party software. While the information collected via the APIs is not required or important for FC switches, it is for many users of other hardware. For instance, VMware customers use the vCenter APIs to automate their IT environment, including to migrate VMs from one server to another or to decommission VMs. Indeed, VMware estimates that around 95% of VMware’s vSphere customer use vCenter. And server and storage OEMs, which are important customers for Broadcom and VMware, use the APIs to provide monitoring software to their enterprise customers. Degrading these APIs would directly affect all of these customers and would therefore be disastrous for the Combined Entity’s reputation and customer relationships.
- 4.40 **Any harm to Cisco would also harm Broadcom.** As noted, the Combined Entity has no ability to foreclose Cisco’s FC switches using VMware’s API. Nevertheless, hypothetically, if degrading the API did impact Cisco FC switches, it would also impact Broadcom’s own FC switches, including those that are installed in existing FC SANs.²³² This would be just as “*undesirable*” as a “*complete breakdown in interoperability*” for I/O devices in existing servers.²³³
- 4.41 In summary, the Transaction will not give rise to an SLC under ToH1 relating to FC switches. There is simply no way to impact the functioning or competitiveness of FC switches through degrading VMware’s APIs.

5. THEORY OF HARM 2: EXCHANGE OF COMMERCIALY SENSITIVE INFORMATION

- 5.1 The Issues Statement’s ToH2 concerns “*whether competition could be harmed by the flow of commercially sensitive information [...] from Broadcom’s hardware competitors to VMware [...] in the supply of Ethernet NICs, FC HBAs, and storage adapters*” as part of the I/O device certification process.²³⁴ Specifically, the Issues Statement posits that the Combined Entity may “*have the incentive to compete less aggressively*” and/or may “*otherwise put its hardware rivals at a competitive disadvantage, particularly in terms of product development/innovation.*”²³⁵
- 5.2 The evidence below shows that the information exchanged during driver development and certification is not *competitively* sensitive and, accordingly, that Broadcom’s access

²³² Developing specific APIs to foreclose Cisco is unrealistic because VMware has no visibility on the software calling on its public API. Indeed, the reason that VMware uses a public API is to avoid the substantial cost and inconvenience of having to develop distinct APIs for the wide variety of software that may want to collect information from vCenter.

²³³ Phase 1 Decision, paragraph 188.

²³⁴ Issues Statement, paragraph 25: (“*the information passed to VMware includes product samples, product roadmaps, driver source code, and other technical information.*”).

²³⁵ Issues Statement, paragraph 29(d).

to that information cannot harm competition or innovation. Moreover, it would be detrimental for the Combined Entity if it were to fail to protect the confidentiality of hardware rivals' *commercially* sensitive information (if it were shared).

No competitively sensitive information is shared during driver development and certification

- 5.3 **I/O device vendors do not compete in relation to device drivers.** The main parameters of competition for NICs, storage adapters, and FC HBAs are [X]. I/O devices simply transfer data and convert it to the appropriate protocol to be understood by the host devices. [X] are determined by the [X] that configures and manages the device. Drivers on the other hand provide a software mapping between the I/O device and the hypervisor APIs. The driver software (*i.e.*, the source code) does not reveal anything about the [X] of an I/O device; it is not competitively sensitive. Hardware vendors have no reason to and do not share their [X] designs with OSs and hypervisors.
- 5.4 **Innovation and investment reside in the [X].** In order to increase [X], I/O device vendors must invest in re-designing and re-engineering their [X]. This is done using software code:
- (a) ***The controller chip [X].*** Hardware vendors like Broadcom use RTL software to design their physical chips, which are too complex to design by hand. For example, Broadcom's latest generation FC HBA contains [X] physical transistors (or silicon gates), which required Broadcom to write around [X] lines of RTL code. Given the complexity of the task, developing a chip typically takes around [X] years.
 - (b) ***The firmware manages the device components and determines how data moves through the I/O device.*** Like RTL, the firmware comprises a large and complex code base. Broadcom's latest generation FC HBA contained around [X] lines of code. Developing firmware typically takes around [X] years.
 - (c) ***The board incorporates the chip and firmware and other components.*** The board, which is physically inserted into servers, takes approximately [X] months to develop and test. The form factor of the board is standardised to enable it to fit in and work with all OEMs' servers.
 - (d) ***The driver sits outside the board and interfaces with the relevant OS or hypervisor.*** Drivers are developed using relatively few and [X] lines of code. For the latest generation FC HBA, Broadcom wrote around [X] lines of driver code. Broadcom estimates that the firmware and RTL code were [X] times and [X] times, respectively, more complicated than the driver code. In other words, writing [X] lines of driver code is the equivalent to writing [X] lines of firmware code and [X] lines of RTL.
- 5.5 **Broadcom's resourcing confirms that driver development is simple.** The way Broadcom allocates its engineers demonstrates that driver development is a straightforward part of I/O device development. For storage adapters, out of [X] engineers, [X] engineers work on driver development, and only [X] work on VMware drivers. For FC HBAs, [X] of Broadcom's [X] engineers work on driver development ([X] for VMware). For NICs, [X] engineers out of [X] work on driver development

([redacted] to [redacted] for VMware). The fact that Broadcom allocates so few engineers to driver development demonstrates the simplicity of the task.

- 5.6 **Driver source code is mostly public information.** Hardware vendors develop drivers to enable their devices to interface with operating systems and hypervisors that are used by enterprises. Since ~95% of servers run Microsoft and Linux operating systems,²³⁶ Broadcom and other vendors prioritise these environments for driver development (see **Figure 8** below).²³⁷ Linux, which is the most widely used operating system, typically comes first, which enables vendors to develop their drivers and fix any issues with help from the open source community. The driver source code for new devices is made public on kernel.org. Hardware vendors then typically use this driver source code to develop their drivers for Microsoft and then other OSs and hypervisors (including VMware). Broadcom has [redacted] agreed [redacted].

Figure 8: Certification of latest generation NICs and storage adapters demonstrates prioritization of Linux

[redacted]

- 5.7 Broadcom estimates that ~90% of its VMware driver source code uses the Linux OS driver source code. The ~10% difference between Broadcom’s Linux driver code and VMware driver code results from differences in VMware’s APIs, which are made available to all members of the relevant certification programme for NICs, storage adapters, and FC HBAs (which is called IOVP). The ~10% difference is not competitively significant. And even if it were, hardware vendors would want their devices to be optimized for the largest proportion of servers, which are those running Linux and Windows OS, not VMware.
- 5.8 **Product roadmaps are rarely shared, are not required for driver development, and are not competitively sensitive.** Broadcom generally provides no product roadmaps to VMware because they are not required for driver development and certification.²³⁸ Over the past 10 years, Broadcom has provided only one FC HBA document to VMware that could be described as a product road map. It has not provided any roadmaps for NICs or storage adapters. The single “roadmap” provided does not contain any competitively sensitive information, as explained in **Annex ISR-025**. Broadcom does [redacted]. [redacted]. Hardware vendors tend to work with other software providers, such as OS providers Linux and Microsoft, with respect to new devices before they work with VMware, such that VMware’s discussions of [redacted] with those vendors lag behind.
- 5.9 **Product samples are not required for driver development and certification and are not competitively sensitive.** The information required by VMware for the purposes of driver development and certification is essentially limited to the driver source code. Broadcom does not provide VMware with product samples. In any event, product samples are not competitively sensitive because (1) samples are widely

²³⁶ As explained in the response to question 7(b) of RFI 3.

²³⁷ Linux-based hypervisors and Hyper-V benefit from this prioritisation because they use the *same driver code* as Linux and Windows OS, respectively.

²³⁸ See the Parties’ response to the RFI of 25 April 2023, Question 16.

distributed to OEMs and others for demonstrations (*see* **Annex ISR-025**); (2) product samples are provided shortly before general release (*e.g.*, [§<]),²³⁹ and (3) product samples are often released with Linux and Windows before they are certified with VMware.²⁴⁰

- 5.10 **Technical information shared for driver development is not competitively sensitive.** As explained in **Annex ISR-025**, the technical exchanges with VMware for the purposes of driver development and certification only require information that is needed to test vendors' compliance with VMware's API specifications and to troubleshoot issues mapping drivers to APIs. Such engagement with VMware is required in less than 1% of cases.²⁴¹

Broadcom's access to rivals' commercially sensitive information would not impact competition

- 5.11 **Interoperability necessitates at least some sharing of commercially sensitive information among hardware and software vendors.** Hardware and software vendors in the datacentre environment routinely share information to ensure their products work together. As noted above, Broadcom exchanges information with its hardware rivals, [§<].²⁴² While it is not the "same" commercially sensitive information that rivals provide to VMware for driver certification – it is more sensitive than driver-related information – [§<].²⁴³ The datacentre environment depends on the sharing of commercially sensitive information to ensure interoperability.²⁴⁴

The Combined Entity would have no incentive to misuse rivals' commercially sensitive information

- 5.12 **Broadcom and VMware depend on their reputations as trusted partners.** Broadcom's and VMware's business models depend on working with a range of hardware and software partners, many of whom share commercially sensitive information with them. Although information shared by rivals for the purposes of driver development and certification is not *competitively* sensitive, it would be disastrous for the Combined Entity if they were to gain a reputation for failing to

²³⁹ As explained in the response to question 34 of RFI 4.

²⁴⁰ For Broadcom's latest generation of FC HBAs, Broadcom obtained verified Linux support for NVMe over FC with Linux on 11 May 2018, Windows Server 2012 and 2016 were updated to include NVMe over FC support for Broadcom's FC HBA on [§<], and then on [§<], vSphere posted support for NVMe over FC for Broadcom's FC HBA.

²⁴¹ See paragraphs 4.9 and 4.10 above.

²⁴² See paragraph 4.26(e).

²⁴³ Issues Statement, paragraph 29(b).

²⁴⁴ See, *e.g.*, *NVIDIA/Mellanox*, Case COMP/M.9424, Commission decision of 19 December 2019, paragraphs 282-310: where the Commission found that the NVIDIA/Mellanox merged entity would not have the ability nor the incentive to misuse GPU suppliers' potentially commercially sensitive information received by Mellanox to favour its own position on the GPU market.

protect, or misusing, *commercially* sensitive information.²⁴⁵ VMware receives information from vendors under NDAs and must ensure that it is protected.

5.13 **The importance of maintaining trust is illustrated by VMware’s neutral stance under Dell and EMC ownership.** While VMware was owned by Dell and EMC, as far as VMware is aware, Dell’s rivals (HP and Lenovo) and EMC’s rivals (Pure Storage and Net App) never complained or raised concerns to VMware about providing information for driver certification or working with VMware on new product developments. These rivals continued to innovate throughout the period. For example, NetApp introduced its flagship storage arrays.²⁴⁶ There is no evidence to suggest that Broadcom’s ownership of VMware would lead to different dynamics.

5.14 [REDACTED]

(a) [REDACTED].²⁴⁷

(b) [REDACTED].²⁴⁸ [REDACTED].²⁴⁹ [REDACTED].²⁵⁰

[REDACTED].

The Merger would not reduce rival hardware vendors’ incentives to innovate

5.15 **Rivals would continue to innovate post-Transaction** even if they did not want to share information with the Combined Entity. Since only c.[10-20]% of servers run on VMware, hardware rivals would continue to develop their products for the [90-100]% of servers running on bare metal or alternative virtualisation software. Rivals would not need to share information with the Combined Entity about these innovations.

5.16 In summary, the Transaction will not give rise to an SLC under ToH2 relating to the sharing of information for the purposes of certification of rival NICs, storage adapters, and FC HBAs.

6. CONCLUSION

6.1 For the reasons set out above, the Transaction will not result in an SLC within any markets in the United Kingdom. In fact, the Transaction is pro-competitive and will enable VMware to compete more effectively with the giant CSPs and to increase choice and flexibility for enterprise workloads.

²⁴⁵ Information sharing is common in the datacentre environment and any misuse of commercially sensitive information would severely damage an entity’s reputation in the market. *See, e.g., NVIDIA/Mellanox*, Case COMP/M.9424, Commission decision of 19 December 2019, paragraphs 307 and 309.

²⁴⁶ Issues Letter Response, paragraph 2.11.

²⁴⁷ [REDACTED].

²⁴⁸ [REDACTED].

²⁴⁹ [REDACTED].

²⁵⁰ [REDACTED].