EUROPEAN COMMISSION

Brussels, 19.12.2019
C(2019) 9467 final

PUBLIC VERSION

In the published version of this decision, some information has been omitted pursuant to Article 17(2) of Council Regulation (EC) No 139/2004 concerning non-disclosure of business secrets and other confidential information. The omissions are shown thus […]. Where possible the information omitted has been replaced by ranges of figures or a general description.

To the notifying party

Subject: **Case M.9424 – NVIDIA / MELLANOX**
Commission decision pursuant to Article 6(1)(b) of Council Regulation No 139/2004[1] and Article 57 of the Agreement on the European Economic Area[2]

Dear Sir or Madam,

(1) On 14 November 2019**,** following a referral pursuant to Article 4(5) of the Merger Regulation, the European Commission received notification of a proposed concentration pursuant to Article 4 of the Merger Regulation by which NVIDIA Corporation ("**NVIDIA**", USA) intends to acquire within the meaning of Article 3(1)(b) of the Merger Regulation control of Mellanox Technologies, Ltd. ("**Mellanox**", Israel) by way of purchase of shares (the "Transaction")[3]. NVIDIA is designated hereinafter as the "Notifying Party", and NVIDIA and Mellanox are together referred to as the "Parties", while the undertaking resulting from the Transaction is referred to as the "Merged Entity".

---

1 OJ L 24, 29.1.2004, p. 1 (the "Merger Regulation"). With effect from 1 December 2009, the Treaty on the Functioning of the European Union (the "TFEU") has introduced certain changes into Union law, such as the replacement of "Community" by "Union" and "common market" by "internal market". The terminology of the TFEU will be used throughout this decision.

2 OJ L 1, 3.1.1994, p. 3 (the "EEA Agreement").

3 Publication in the Official Journal of the European Union No C 398, 25.11.2019, p. 6.

Commission européenne, DG COMP MERGER REGISTRY, 1049 Bruxelles, BELGIQUE
Europese Commissie, DG COMP MERGER REGISTRY, 1049 Brussel, BELGIË

Tel: +32 229-91111. Fax: +32 229-64301. E-mail: COMP-MERGER-REGISTRY@ec.europa.eu.

1. THE PARTIES

(2) NVIDIA is a publicly traded Delaware Corporation, founded in 1993 and headquartered in Santa Clara, California. NVIDIA invented the graphics processing unit ("**GPU**") in 1999. NVIDIA specializes in markets in which GPU-based visual computing and accelerated computing platforms can provide enhanced throughput for applications. NVIDIA's products address four distinct areas: gaming, professional visualization, datacentre, and automotive. Only the datacentre area is relevant to the Transaction with Mellanox, as datacentre customers are the only ones also procuring components from Mellanox. In addition to GPU cards, NVIDIA produces software called "NVIDIA GRID" that allows computers that do not have their own GPU to use "virtual GPUs". Moreover, it offers one family of server systems (DGX-1, DGX-2, and DGX-Station) that perform GPU-accelerated Artificial Intelligence ("**AI**") and deep learning training and inference applications. Finally, NVIDIA also offers the NGC Software Hub, a cloud-based software repository for AI developers that provides deep learning software stacks.

(3) Mellanox is a publicly held corporation, founded in 1999 and headquartered in Sunnyvale, California and Yokneam (Israel). Mellanox offers network interconnect products and solutions that facilitate efficient data transmission between servers, storage systems and communications infrastructure equipment within datacentres, based on two network interconnect protocols: Ethernet and InfiniBand. Mellanox's network interconnect components include the following: network interface cards ("**NICs**" or "**network adapters**"), switches and routers, cables, and related software.

2. THE CONCENTRATION

(4) On 10 March 2019, NVIDIA, NVIDIA International Holdings Inc., ("**NVIDIA Holdings**", USA), Teal Barvaz Ltd. ("**Teal Barvaz**", Israel) and Mellanox entered into an Agreement and Plan of Merger ("**Merger Agreement**"). NVIDIA Holdings is a wholly-owned subsidiary of NVIDA. Teal Barvaz is a wholly-owned subsidiary of NVIDIA Holdings.

(5) Pursuant to the Merger Agreement, the Transaction will be implemented as follows: Teal Barvaz (Merger Sub) merges with and into Mellanox with Mellanox being the surviving entity, following which Teal Barvaz will cease to exist, and Mellanox will become a wholly owned subsidiary of NVIDIA Holdings. Each of Mellanox's shares will be transferred to NVIDIA Holdings in exchange for the right to receive an amount in cash equal to USD 125 (approximately EUR 106), representing a total acquisition value of approximately USD 6 900 million (approximately EUR 5 800 million) to be paid by NVIDIA.

(6) Therefore, NVIDIA (via NVIDIA Holdings) will acquire sole control over Mellanox and the Transaction constitutes a concentration within the meaning of Article 3(1)(b) of the Merger Regulation.

3. UNION DIMENSION

(7) The Transaction does not have a Union dimension within the meaning of Article 1(2) or Article 1(3) of the Merger Regulation as the EU turnover of one of the

Parties (Mellanox) in the last financial year for which data is available at the date of notification amounted to EUR […].

(8) On 14 June 2019, the Notifying Party informed the Commission by means of a reasoned submission that the Commission should examine the Transaction pursuant to Article 4(5) of the Merger Regulation. The Commission transmitted a copy of that reasoned submission to the Member States on 14 June 2019.

(9) In fact, the Transaction fulfils the two conditions set out in Article 4(5) of the Merger Regulation since it is a concentration within the meaning of Article 3 of the Merger Regulation and it is capable of being reviewed under the national competition laws of three Member States, namely Czechia, Germany and Hungary.

(10) As none of the Member States competent to review the Transaction expressed its disagreement as regards the request to refer the case, the Transaction is deemed to have a Union dimension pursuant to Article 4(5) of the Merger Regulation.

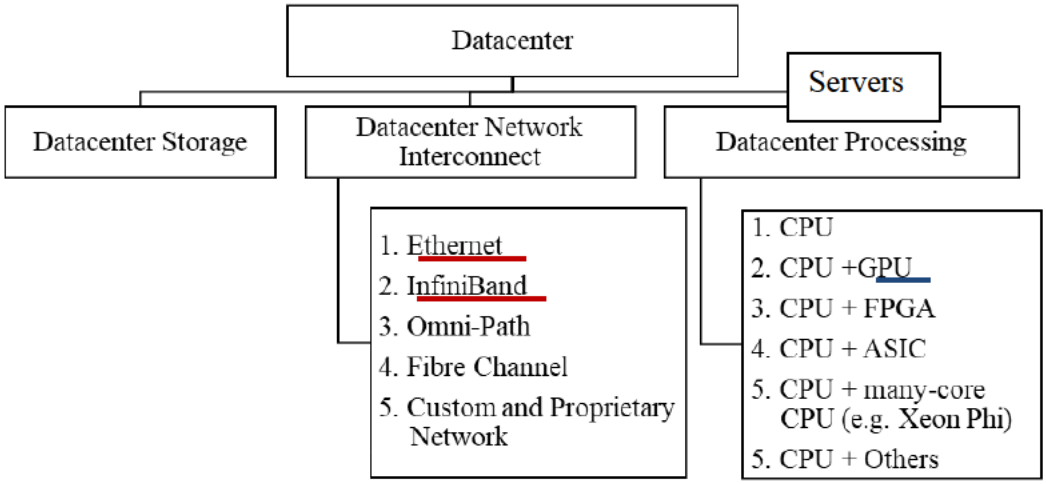
4. RELEVANT MARKETS

4.1. Introduction

(11) The Transaction concerns key components used in datacentre servers, in particular those used for high performance computing ("**HPC**"; also sometimes referred to as "**supercomputing**"). HPC datacentres deliver the computation power required for research and innovations in a number of key developing areas such as autonomous driving, weather forecast, oil exploration and astrophysics. In particular, HPC datacentres are key enablers for many AI applications. Both Parties supply different components that can be used in datacentre servers.

(12) Datacentres are a collection of servers that are connected by a network and that work together to process/compute workloads. Datacentres in general have three fundamental elements: (1) storage/memory, (2) network interconnect, and (3) processing/computing included in the servers, as Figure 1 illustrates.

Figure 1: Datacentre Overview



Source: based on Figure 1 in the Form CO, page 28.

(13) The components produced by the Parties are rarely bought separately by end customers. They are mainly sold through original equipment manufacturers ("**OEMs**") and original design manufacturers ("**ODMs**"), which build datacentre servers based on specific requests of end customers.[4] However, the Parties sometimes also sell their products directly to large end customers, such as cloud service providers ("**CSPs**").[5]

(14) The Transaction does not give rise to any horizontal relationships, as NVIDIA's GPUs and servers are not substitutable with Mellanox's network interconnects.

(15) However, the Transaction gives rise to a vertical relationship between the supply of datacentre servers (in which NVIDIA is active) and the supply of network interconnect cards (in which Mellanox is active).

(16) Finally, the Transaction results in a conglomerate relationship between NVIDIA's GPUs and Mellanox's network interconnects. These products perform complementary functions in datacentre servers, in particular for HPC applications, with GPUs performing so-called "workload acceleration" and network interconnects connecting accelerated servers.

(17) In light of the above, the Commission will examine the relevant market definition in relation to the following product areas: (i) supply of discrete GPUs for datacentres, (ii) supply of datacentre network interconnects, and (iii) supply of datacentre servers.

---

4 Datacenter end customers unilaterally determine the nature and scope of their datacenters working directly with OEM/ODM suppliers. They award contracts pursuant to individual tenders, whose scope is controlled and determined entirely by customers and OEMs/ODMs. See Form CO, paragraph 433.

5 In 2018, NVIDIA sold […]% of its GPUs through OEMs/ODMs and […]% directly to datacentre customers. Mellanox sold […]% of its network interconnects to OEMs/ODMs and […]% directly to datacentre customers. See Form CO, paragraph 582, figures 7 and 8.

4.2. Discrete GPUs for datacentres

4.2.1. Introduction

(18) GPUs are specialised semiconductor devices that are optimized for processing graphic images. They are made available to customers either as standalone, or "discrete", semiconductor devices or as integrated components of chips that contain other components, including a central processing unit ("**CPU**"). In datacentres, GPUs are used to accelerate the datacentre workload computing/processing. They are only used in datacentres that require acceleration. They are necessarily used in addition to CPUs, which are always present in datacentre servers.

(19) When CPUs and GPUs are combined in a datacentre server, they carry out complementary computational tasks. CPUs operate as the general purpose centralised "brains" of computer systems. They are able to perform all types of operations. In contrast, GPUs – although they have more limited computational capabilities – are much better suited to process graphic images or computations that require massive parallel execution of relatively simple computational tasks. This is why GPUs are increasingly used in HPC and key AI applications, which both require massive parallel execution of rudimentary arithmetic operations.

(20) Other types of accelerators are sometimes used in datacentre servers: (1) field programmable gate arrays ("**FPGAs**"); (2) application specific integrated circuits ("**ASICs**"); (3) CPUs with integrated GPUs or **many-core CPUs** (as opposed to the standalone/discrete GPUs offered by NVIDIA).

4.2.2. Product market definition

4.2.2.1. Commission precedents

(21) The Commission has not assessed the boundaries of the market for discrete GPUs for datacentres in past decisions. However, in its 2011 *Intel/McAfee* decision,[6] the Commission has defined a separate market for CPUs based on the x86 architecture ("**x86 CPUs**"), in which GPUs and other accelerators were not included. In its 2015 *Intel/Altera* decision, the Commission also defined a separate market for FPGAs, distinct from other complex programmable logic devices ("**CPLDs**") and from GPUs.[7]

4.2.2.2. Notifying Party's views

(22) The Notifying Party submits that the relevant product market is the market for datacentre processing, which includes all types of processing, i.e., GPUs, as well as CPUs, FPGAs, ASICs (including ASICs developed in-house by companies, *e.g.*, Google's Tensor Processing Unit ("**TPU**")) and any other processors/accelerators for datacentres.[8]

(23) The Notifying Party submits that from a demand-side perspective, for datacentre customers, GPUs are one acceleration choice amongst many and that the competition

---

6 Commission decision of 26 January 2011 in case M.5984 – *Intel/McAfee*, paragraph 30.

7 Commission decision of 14 October 2015 in case M.7688 – *Intel/Altera*, paragraph 41.

8 Form CO, paragraph 183.

between accelerators is fierce. The Parties submit that GPUs cannot perform any processing that cannot also be performed by CPUs, FPGAs, ASICs, and other accelerators. Customers would consider all datacentre processing options, not only GPUs, which competitively constrains GPUs. No accelerator, or even any category of accelerators, is essential and indispensable to any datacentre. When designing and constructing their datacentres, customers will consider the processing needs of the datacentre as a whole. Datacentre customers compare all processing options in accordance with multiple variables, most notably price, performance, efficiency, and scalability.[9] Additionally, the Notifying Party submits that no accelerator option is always best suited for all applications, or even for a given type of application.[10]

(24) Moreover, the Notifying Party submits that the different types of datacentre accelerators are also constrained by cloud-based computing, for the following reasons. On the one hand, when customers compare the overall price and performance of datacentres, they take into account both building a system on their own premises and buying cloud-based datacentre computing. This would in turn constrain the GPU's pricing and commitment to innovation. On the other hand, GPU sales to cloud service providers are constrained by the in-house solutions these providers develop.[11]

(25) The Notifying Party submits that there is also supply-side substitution. Several suppliers of datacentre processing have developed products that perform parallel processing, and act as accelerators, even if the suppliers do not name the products GPUs.[12] Finally, the Notifying Party claims that suppliers of GPUs for use beyond datacentres can also easily supply datacentre GPUs. This is because both types of GPUs are based on the same fundamental architecture.[13]

4.2.2.3. Results of the market investigation and Commission's assessment

(26) As a preliminary remark, the market investigation confirmed that discrete GPUs for datacentres and discrete GPUs for gaming are part of different markets.[14] While the two types of GPUs are based on the same architecture,[15] they have different levels of performance, due to technical limitations of the GPUs for gaming.[16] Moreover, contractual restrictions and driver support prevent purchasers of NIVIDIA's discrete GPUs for gaming from deploying these in datacentres.[17] As a result, the majority of

9 Form CO, paragraphs 213-215.

10 Form CO, paragraph 216.

11 Form CO, paragraph 212.

12 Form CO, paragraph 294.

13 Form CO, paragraph 295.

14 NVIDIA's GPUs are also used to create enhanced computer graphics for video games, professional visualisation and automotive applications. However, demand substitution seems to be limited because the GPU product lines used for such applications are not the same as the ones used for datacentre workloads. Moreover, NVIDIA has banned the use of its consumer-grade GPUs in datacentres (via licensing restrictions). From a supply side perspective, it is difficult to switch from supplying GPUs for graphics to supplying GPUs for datacentre due to the software barrier to entry created by the entrenchment of NVIDIA's CUDA software as the dominant programming GPU interface.

15 Form CO, paragraph 295.

16 See Replies to Questionnaire Q2 to OEMs, questions 8 and 8.1.; Questionnaire Q3 to End Customers, questions 9 and 9.1; agreed minutes of the conference call of 25 October 2019 with a CSP, paragraph 10.

17 See Replies to Questionnaire Q2 to OEMs, questions 8 and 8.1.; Questionnaire Q3 to End Customers, questions 9 and 9.1; .agreed minutes of the conference call of 25 October 2019 with a CSP, paragraph 10.

datacentre customers does not consider discrete GPUs for datacentres and discrete GPUs for gaming as substitutes.[18]

(27) As for the different types of datacentre processors, such as integrated GPUs, CPUs, FPGAs or ASICs, the responses to the market investigation showed that competitors, OEMs and most end customers consider different types of accelerators to be suitable for different kinds of HPC and deep leaning applications. [19] They are therefore likely not part of the same market as discrete GPUs for datacentres.[20]

(28) Competitors and OEMs indicated that there are specific datacentre parallel workloads for which GPUs have become the standard acceleration solution. They submitted that for the most powerful high performance computers as well as for AI applications, a combination of CPUs and GPUs achieves the maximum efficiency. End customers would therefore prefer it to CPU-only architectures.[21]

(29) Several CSPs submitted that some high performance cloud computing customers specifically ask for GPU accelerated servers for certain types of workloads, such as deep learning, physic simulation or molecular modelling. These customers would not be willing to perform these compute-intensive workloads on servers with other types of acceleration.[22] This could be because the workload of these customers is already optimized for NVIDIA's GPUs (*e.g.*, using NVIDIA's Compute Unified Device Architecture ("**CUDA**")).[23] Moreover, for some applications, the total cost of ownership is lower when using GPUs, while for other applications, other types of accelerators (TPUs, FPGAs) have a total cost of ownership advantage.[24] Therefore, the CSPs consider that in most cases, GPUs and other types of accelerators are not substitutable when considering the specific applications they are meant to serve.

(30) The above was confirmed by the replies of end customers. For most HPC applications, the majority of end-customers considered only discrete GPUs, integrated GPUs and adding more CPUs as suitable alternatives.[25] However, they considered both integrated GPUs and adding more CPUs to have certain limitations in comparison with discrete GPUs. End customers pointed out that integrated GPUs

---

18 See Replies to Questionnaire Q3 to End Customers, questions 9 and 9.1.

19 Replies to Questionnaire Q1 to Competitors, questions 6 to 7.1; Questionnaire Q2 to OEMs, questions 6 to 7.1; Questionnaire Q3 to End Customers, questions 6 to 7.1.

20 The Commission also assessed whether discrete GPUs for datacentres are constrained by cloud-based solutions. However, the large majority of end customers stated that they would not consider renting computing power "as-a-service" from cloud-based solutions using the cloud service supplier's own in-house accelerator, such as Google's TPU, even if the prices for a cluster of GPU accelerated servers were to increase as a result of an increase in the price of GPUs. This is partly due to cost reasons, but also, especially for universities and research centres based in the EEA, due to their funding schemes and the desire to not move their data on non-EEA cloud providers. See Replies to Questionnaire Q3 to End Customers, questions 13 and 13.1. Based on this, the Commission considers that discrete GPUs for datacentres are not part of the same product market as cloud-based solutions.

21 See agreed minutes of the conference call of 9 August 2019 with a competitor, paragraphs 7 and 10; agreed minutes of the conference call held with a major OEM on 12 August 2019, paragraphs 9 and 10.

22 Agreed minutes of the conference call of 11 October 2019 with a CSP, paragraph 11; agreed minutes of the conference call of 25 October 2019 with a CSP, paragraph 6.

23 Agreed minutes of the conference call of 11 October 2019 with a CSP, paragraph 11.

24 Agreed minutes of the conference call of 11 October 2019 with a CSP, paragraph 11; agreed minutes of the conference call of 25 October 2019 with a CSP, paragraph 7.

25 Replies to Questionnaire Q3 to End Customers, questions 6 to 7.1.

mainly support light acceleration tasks, while adding more CPUs entails adding expense and power consumption in comparison to using discrete GPUs.[26]

(31) Moreover, for demanding parallelised workloads, ASICs and FPGAs are not considered suitable alternatives.[27] This is because ASICs are typically designed for very specific workloads and only for a particular customer. Since they lack flexibility and built-in software environment, they are not a suitable option for customers with limited resources who need to perform a range of different HPC workloads.[28]

(32) As for FPGAs, the respondents to the market investigation considered that while they are highly flexible and can be programmed to run parallelized workloads, they are unsuitable for most HPC applications. This is because, when compared to GPUs, FPGAs are relatively inefficient in terms of energy consumption, cost more and are more difficult to program.[29]

(33) End customers confirmed that when they decided to acquire a GPU accelerated server in the past, no other type of accelerator was considered a suitable alternative.[30] All end customers that replied to the question confirmed that for all GPU accelerated datacentres they procured between 2017 and 2019, they procured the GPUs from NVIDIA and were either not open to any alternative accelerated processing solutions or considered only GPUs from AMD as an alternative.[31]

(34) Based on the above, the Commission considers that there is a separate product market for discrete GPUs for datacentres, which does not include other types of datacentre processing solutions.

4.2.3. Geographic market definition

(35) The Commission has not assessed the geographic scope of the market for discrete GPUs for datacentres in past decisions. However, it has concluded that the market for CPUs (and possible segments thereof)[32] as well as the market for FPGAs[33] is worldwide in scope.

(36) The Notifying Party submits that the relevant geographic market for datacentre processing should be defined as worldwide.[34] This should remain the case even for potentially narrower product markets.[35]

---

26 Replies to Questionnaire Q1 to Competitors, question 6.1.1; Questionnaire Q3 to End Customers, question 6.1.1.

27 See Replies to Questionnaire Q1 to Competitors, question 6.1.1; agreed minutes of the conference calls of 29 July and 12 August 2019 with a major OEM, paragraph 12; agreed minutes of the conference call of 12 August 2019 a major OEM, paragraph 10.

28 See Replies to Questionnaire Q1 to Competitors, question 6.1.1; agreed minutes of the conference call of 12 August 2019 with a major OEM, paragraph 10.

29 See Replies to Questionnaire Q1 to Competitors, question 6.1.1.

30 Replies to Questionnaire Q3 to End Customers, question 10.

31 Replies to Questionnaire Q3 to End Customers, question 10.

32 Commission decision of 14 October 2015 in case M.7688 – *Intel/Altera*, paragraph 25.

33 Commission decision of 14 October 2015 in case M.7688 – *Intel/Altera*, paragraph 57.

34 Form CO, paragraph 360.

35 Form CO, paragraph 367.

(37) The majority of competitors, OEMs and end customers that replied to the market investigation confirmed that accelerated processing solutions are supplied on a worldwide basis, irrespective of the location of the component vendor or the location of the end-customer.[36] Moreover, the majority of competitors and OEMs and all end customers that expressed a view confirmed that the conditions of competition do not differ depending on the location of the datacentre of the end customer.[37]

(38) In light of the results of the market investigation, for the purposes of this decision, the Commission considers that the geographic market for discrete GPUs for datacentres is worldwide in scope.

4.3. Datacentre network interconnects

4.3.1. Introduction

(39) Datacentre's network interconnects enable the transfer of data between servers or systems, *e.g.*, connecting multiple servers together or connecting a server in a datacentre to a storage appliance.

(40) Network interconnects are made up of the following main components: (i) NICs that are used in the server, enabling it to communicate with other devices on the network; (ii) switches and routers that manage communications between servers;[38] (iii) cables that connect devices together and carry the data signals between devices; and (iv) supporting software.

(41) Network interconnects can be based on a variety of protocols, some of which are based on open standards (*e.g.*, Ethernet, InfiniBand and Fibre Channel ("**FC**")), while others are custom or proprietary (*e.g.*, Cray's Aries/Gemini/Slingshot, Atos Bull's BXI, and Fujitsu's Tofu)[39]. The latter are currently not available to external server OEMs, they are only sold within their supplier's own systems.

(42) The technical parameters considered by customers when selecting an interconnect solution include *inter* alia bandwidth, latency, interoperability, congestion control, deployment. Based on the results of the market investigation, bandwidth and latency are particularly importance parameters for HPC customers.[40] Bandwidth is a measure of how much data can be sent and received at a time, which is a critical factor given it measures the capacity of a network. Latency measures the time required to transmit a packet across a network. In addition to these factors, customers would consider total cost of ownership, as well as other parameters such as the quality of service.

---

36 Replies to Questionnaire Q1 to Competitors, questions 12 and 12.1; Questionnaire Q2 to OEMs, questions 13 and 13.1; Questionnaire Q3 to End Customers, questions 16 and 16.1.

37 Replies to Questionnaire Q1 to Competitors, questions 13 and 13.1; Questionnaire Q2 to OEMs, questions 14 and 14.1; Questionnaire Q3 to End Customers, questions 17 and 17.1.

38 They are analogous to switchboard operators on old phone systems, connecting calls between devices.

39 FC is a standard that is used primarily in storage networks, whereas Ethernet, InfiniBand and other custom and proprietary network interconnects are mainly used to ensure the flow of data across the many servers composing the datacentre.

40 See for example agreed minutes of the conference call of 12 August 2019 with a large OEM, paragraph 6; agreed minutes of the conference call of 9 August 2019 with an end-customer providing HPC services to universities and research centres, paragraph 13; and Intel's submission of 10 September 2019 entitled "Intel response to case team's questions", question 4.b.

(43) Mellanox's products are based on the Ethernet and InfiniBand protocols. Mellanox offers network interconnect integrated systems and components, which include NICs, switches, cables, and related software.

4.3.2. Product market definition

4.3.2.1. Commission precedents

(44) There are no Commission precedents defining the relevant markets for the datacentre network interconnect products manufactured and supplied by Mellanox.

(45) Several Commission decisions however address transactions involving network interconnect switch suppliers. In *Broadcom/Brocade*,[41] the Commission distinguished between the two main components of Fibre Channel SAN networks, i.e. switches and adapters. As regards switches, the Commission has previously considered a segmentation based on the different protocols and network technologies that they support.

(46) In addition, in a previous decision,[42] the Commission stated that the majority of respondents to its market investigation were of the view that the SerDes (serializer/deserializer) intellectual property used in high-speed NICs should be placed in different markets "*for each standard speed (1G, 10G, 25G, 50G and the future 100G.)*".

4.3.2.2. Notifying Party's views

(47) The Notifying Party submits that the relevant product market is the market for datacentre network interconnect, which includes all types of datacentre network interconnect solutions and all individual components.

(48) First, the Notifying Party submits that, from a demand-side perspective, Mellanox's InfiniBand and Ethernet products are substitutable with datacentre interconnect products based on other protocols. This is based on the following main arguments:

- All of the competing protocols serve the same purpose and function and there are no technical differences between the major interconnects that would limit their substitutability. In particular, technically, Ethernet as well as other protocols are similar to InfiniBand on every relevant parameter, including, in particular, bandwidth and latency.

- Customers typically invite bids from different suppliers for different datacentre network interconnect solutions as illustrated by a number of examples of bids in which Mellanox's Ethernet- and InfiniBand-based solutions competed directly against network interconnect solutions based on other protocols.

- Customers can and do switch from one protocol to another for newer generation datacentres as illustrated by a number of examples of customers having switched away from InfiniBand. In addition, customers can

41 Commission decision of 12 April 2017 in case M.8314 – *Broadcom/Brocade*, paragraphs 25-46.

42 Commission decision of 23 November 2015 in case M.7686 – *Avago/Broadcom*, paragraph 60.

implement modules within a datacentre that use different network interconnects given that datacentre networking products comply with open standards and are interoperable.

(49) Second, the Notifying Party submits that there is a high degree of supply side substitutability for different interconnect solutions. Specifically, the Notifying Party explains that the barriers to enter another interconnect solution are quite low from the perspective of an existing supplier. In this respect, the Notifying Party estimates that suppliers can switch between the development and supply of different network interconnect solutions with an investment of less than USD […] and over a period of around […] years. The Notifying Party also describes examples of suppliers having abandoned InfiniBand and other interconnect technologies in favour of Ethernet, noting that these suppliers possess the technology and knowledge to readily switch back to these interconnect technologies.

(50) Third, the Notifying Party submits that, even though it may be conceivable from a demand-side perspective, it is not appropriate to define separate markets for each of the components that make up network interconnects given that competition among providers of datacentre components occurs at the "interconnect solutions" level. According to the Notifying Party this is also supported by the following supply-side substitutions considerations: (i) the know-how and technical skills for different components of interconnect solutions are similar and transferrable and (ii) some interconnect competitors (*e.g.*, Cray) price and sell their offerings at the system level, not the component level.

(51) Fourth, the Notifying Party submits that the market for datacentre network interconnect should not be further sub-segmented based on bandwidth. In particular, the Notifying party considers that it would be incorrect to define a market for Ethernet NICs that support bandwidths of 25 Gb/s or higher since both other network interconnect protocols and Ethernet NICs that support bandwidth of below 25 Gb/s exercise competitive constraints on 25 Gb/s+ Ethernet NICs.

(52) With respect to the competitive pressure exercised by Ethernet NICs that support bandwidth of below 25 Gb/s, the Notifying Party argues that the same hardware can be used for Ethernet NICs of different speeds and that a customer can therefore combine several 10 Gb/s NICs as an alternative to match the performance of a 25 Gb/s+ NIC. In addition, Mellanox reports that it has experienced situations where customers declared a strong preference for deploying 25 Gb/s+ Ethernet solutions in a datacentre node, but based on cost, convenience and other factors ultimately deployed 10 Gb/s solutions.

4.3.2.3. Commission's assessment

(53) The Commission has considered a potential segmentation of the market for datacentre network interconnects between the various protocols and network technologies that they support. In particular, the Commission has considered a distinction between high performance fabrics and Ethernet-based network interconnects. The Commission has also considered further potential sub-segmentations within high performance fabrics and Ethernet-based network interconnects.

(1) Distinction between protocols

(54) Among the different protocols used for running datacentre network interconnects, the Commission has considered a potential distinction between, on the one hand, high performance fabrics and, on the other hand, Ethernet based network interconnects.

(55) Ethernet is the most widely used protocol for network interconnect solution around the world. It is also the fastest growing interconnect protocol. Ethernet products are supplied by companies including Mellanox, Intel, Broadcom, Arista, Cisco, and Juniper.[43] Ethernet-based network interconnect solutions are available for a range of speed including speeds delivering high performance of above 25 Gb/s. In addition, Ethernet suppliers have developed new protocols to combine Remote Direct Memory Access ("**RDMA**") technology with Ethernet (namely RDMA over converged Ethernet ("**RoCE**") and iWARP protocols) allowing to reduce the latencies for Ethernet.

(56) Besides Ethernet-based network interconnects, a number of suppliers offer high performance fabrics.[44] High performance fabrics are integrated systems of custom hardware including NICs, switches, and cabling. They are designed to run on custom protocols and orchestrated by custom software. High-performance fabrics typically enable reliable high-speed communications across several hundreds or thousands of nodes. The different fabric components are designed to work together as part of the integrated fabric and are orchestrated by sophisticated software. High performance fabrics include Mellanox's InfiniBand and Intel's Omni-Path, as well as custom datacentre network interconnects based on protocols that are compatible with Ethernet such as Cray's Aries/Gemini/Slingshot, Bull's BXI and Fujitsu's Tofu.

(57) Based on its market investigation, the Commission identifies two distinct product markets for (i) high performance fabrics and (ii) Ethernet-based network interconnects. This conclusion is based on the limited substitutability between high performance fabrics and Ethernet-based network interconnects, both from a demand-side and supply-side perspective.

(58) First, from a demand-side perspective, while the performance of Ethernet-based systems has increased in the recent years, this performance remains significantly inferior to the performance achieved by high performance fabrics and in particular by Mellanox's InfiniBand on a number of key parameters. Based on the results of the market investigation, the main performance gap between high performance fabrics and Ethernet-based solutions concerns latency.[45] Indeed, a large number of

---

43 The Notifying Party also mentions smaller suppliers of Ethernet-based network interconnects including Extreme Networks, Solarflare (which is to be acquired by Xilinx), QLogic, Chelsio Communications, Myricom, Barefoot Networks (which is to be acquired by Intel), *etc*. In addition, some interconnect customers have developed their in-house interconnect. For example, Google uses its own Ethernet-based custom network interconnects.

44 Network interconnects can also be based on the Fibre Channel protocol. Fibre Channel is however primarily used for a different (complementary) function than Mellanox's InfiniBand and Ethernet-based network interconnects. Fibre Channel is used to connect storage servers in modules within datacentres. Today, Fibre Channel products are supplied by companies including Broadcom, Marvell, Cisco, IBM, HPE, *etc*.

45 Respondents also mentioned other parameters. For example, Intel explained that InfiniBand offers a number of advantages over Ethernet including lossless data transport (information or packets are not

customers and competitors explained that Ethernet-based solutions do not constitute a suitable alternative to high performance fabrics because of significantly lower performance in terms of latency. Certain respondents further explained that low latency requirements in customer's specifications *de facto* exclude Ethernet-based solutions from certain HPC tenders.[46]

(59) For example, a large OEM explained: "*for some applications, low latency is a necessity and so is a high performance fabric*".[47] With respect to the question whether Ethernet-based solution could be an alternative, this OEM further explained: "*as a general principle, Ethernet is unsuitable for low latency modules. Ethernet will therefore generally not be a possible substitute for applications requiring the level of performance which Mellanox's InfiniBand can deliver*".

(60) Mellanox's own data also confirms the gap in latency between high performance fabrics and Ethernet based solutions since these show that the latency of Mellanox's most advanced Ethernet NIC is […] the latency achieved by its InfiniBand products.[48]

(61) Second, respondents to the Commission's market investigation confirm that there is no demand-side substitutability between high performance fabrics and Ethernet-based solutions for a significant portion of HPC datacentres applications. All competitors, all OEMs and a majority of end-customers who expressed their views consider that there are applications and workloads for which their company would only consider high performance fabrics as suitable.[49] These are typically high-end HPC and AI deep learning training applications, which require large systems combining many hundreds or thousands of nodes and for which low latency and high bandwidth is particularly important.[50]

(62) A number of universities and research centres further explained that their server clusters must be capable of accommodating a broad range of HPC workloads, including workloads requiring low latency for which only high performance fabrics are suitable.[51] Since low latency is a requirement at least for certain workloads, these customers explained that they consider high performance fabrics as the only possible choice for equipping the HPC server clusters that they operate.

(63) In addition, two major OEMs provided data on the last ten GPU accelerated datacentres (or GPU accelerated server clusters within datacentres) connected with

---

dropped), higher throughput rates, and significantly lower latency. See Intel response to the case team's questions, 9 September 2019, pages 16-17.

46 Replies to Questionnaire Q2 to OEMs, question 21.1.

47 See agreed minutes of the conference calls of 28 July 2019 and 12 August 2019 with a large OEM, paragraphs 13-19. See also Replies to Questionnaire Q2 to OEMs, questions 21 and 21.1.

48 Form CO, Annex RFI 4 – 01, Table 5.

49 Replies to Questionnaire Q3 to End Customers, question 19; Questionnaire Q2 to OEMs, questions 21 and 21.1.; Questionnaire Q1 to Competitors, questions 15 and 15.1.

50 See for example agreed minutes of conference call of 12 August 2019 with a large OEM, paragraph 6; agreed minutes of conference call of 9 August 2019 with an end-customer providing HPC services to universities and research centres, paragraph 13; and Intel's submission of 10 September 2019 entitled "Intel response to case team's questions", question 4.b.

51 See, for example, agreed minutes of conference call of 9 August 2019 with an end-customer providing HPC services to universities and research centres, paragraph 14.

Mellanox's network interconnects that they have installed/equipped.[52] For all datacentres equipped with InfiniBand fabrics, these OEMs confirmed that they had no suitable network interconnect alternatives. In particular, they confirmed that they did not see Ethernet-based network interconnects as suitable alternatives.

(64) Third, several competitors and customers having indicated that there are certain HPC applications for which only high performance fabrics are suitable further specified that even RoCE-enabled Ethernet solutions with a speed of 25 Gb/s or higher do not currently achieve performance equivalent to InfiniBand in particular in terms of latency.[53] These respondents explained that while RoCE-enabled Ethernet solutions may be suitable for low-end HPC applications that do not demand the highest performance levels offered by high performance fabrics such as InfiniBand, these solutions do not currently constitute an alternative for a large number of other HPC applications and complex AI machine learning training work.

(65) Fourth, Mellanox's internal documents also [BUSINESS SECRETS – Information redacted regarding business strategy].[54] In addition, [BUSINESS SECRETS – Information redacted regarding business strategy].

(66) Fifth, high performance fabrics differ from Ethernet-based network interconnects because they are sold almost exclusively as integrated systems whereas Ethernet-based network interconnects are typically sold as individual components. While customers sourcing Ethernet-based network interconnects have the possibility to mix-and-match between several different suppliers (*e.g.*, buying NICS from one supplier and switches from another), this possibility is, in principle, not available for customers wanting to source high performance fabrics.

(67) This is confirmed by Mellanox which explains that: [BUSINESS SECRETS – Information redacted regarding business strategy].[55]

(68) In addition, the Commission considers that there is only limited supply-side substitutability between Ethernet-based solutions and high performance fabrics, as shown by the following elements.

(69) First, the results of the market investigation show that there is no credible prospect that a competitor of Mellanox with Ethernet-based network interconnects would develop and start offering a high performance fabric that would be able to compete successfully with Mellanox's latest generation InfiniBand fabric within less than a year.[56] Based on the replies from competitors, launching a competitive InfiniBand fabric or another type of high-performance fabric would take at least three years and would require costs significantly higher than the Notifying Party's estimate reproduced above. A number of competitors also mentioned Intel's Omni-Path

---

52 Replies to Questionnaire Q2 to OEMs, question 29.

53 See Replies to Questionnaire Q1 to Competitors, question 15.1; Questionnaire Q3 to End Customers, question 28.1; Intel's submission of 10 September 2019 entitled "Intel response to case team's questions", question 4.b.

54 For example, Mellanox internal document, […], slides 4, 49 and 72; Mellanox internal document, […], slides 4, 39 and 62.

55 Form CO, Annex RFI 3 – 0, paragraph 185. Mellanox only mentions […] standalone sales of InfiniBand products for repair purposes.

56 Replies to Questionnaire Q1 to Competitors, question 24.2.

failure as an example of the difficulty to enter the market for high performance fabrics.

(70) Second, while certain respondents envisage that further technological progress may allow Ethernet-based network interconnect systems to become suitable alternative to high performance fabrics in the future, they acknowledge that this is currently uncertain. In this respect, beyond Slingshot, which some end-customers considered as Ethernet-based, the majority of end-customers and competitors do not believe that a competing Ethernet fabric will emerge within the next 2-3 years that would be able to compete successfully with InfiniBand.[57]

(2) Segmentation within high performance fabrics

(71) Within high performance fabrics, the Commission has considered potential segmentations (i) between different protocols, (ii) between the various components composing high performance fabrics and (iii) based on bandwidth.

(a) Distinction between protocols within high performance fabrics

(72) The Commission has considered a further distinction within high performance fabrics based on protocols. In particular, the Commission has considered a potential separate relevant product market for InfiniBand high performance fabrics which could also potentially include Omni-Path given that this proprietary, high performance communication architecture developed and owned by Intel, is rooted from the InfiniBand technology Intel acquired from QLogic in 2012.

(73) Based on the results of the market investigation, there are indication that customers wanting to procure InfiniBand fabric would not consider any other alternative high performance fabrics (see above).

(74) However, for the purpose of this decision, the question whether high performance fabrics should be further distinguished based on protocols can be left open since it does not materially change the Commission's assessment.

(b) Distinction between individual components within high performance fabrics

(75) With respect to a potential distinction between individual components, as explained above, a key characteristic of high performance fabrics versus for example Ethernet-based network interconnect systems is that fabrics are integrated systems of custom hardware designed to run on custom protocols and orchestrated by custom software. There is therefore only very limited or no interoperability between individual components across the different custom protocols for high performance. In addition, […], the various components composing high performance fabrics are almost always sold together to final customers. This is also confirmed by several customers.[58]

(76) Based on these elements, the Commission considers that there is no need to further distinguish high performance fabrics based on the individual components composing the fabrics.

---

57 Replies to Questionnaire Q1 to Competitors, question 38; Questionnaire Q3 to End Customers, question 43.

58 Replies to Questionnaire Q3 to End Customers, question 27.1.; Questionnaire Q2 to OEMs, question 16.1.

(c) Distinction based on bandwidth within high performance fabrics

(77) With respect to a potential distinction based on bandwidth, Mellanox currently offers InfiniBand interconnect solutions with a bandwidth of 100 Gb/s and 200 Gb/s.[59] By contrast, other high performance fabrics currently only support speeds of 100 Gb/s.

(78) The Commission's market investigation confirm that bandwidth is an important parameter for the choice of network interconnect systems in particular in an HPC context.[60] One large OEM selling solutions equipped with InfiniBand fabrics to a large number of end customers considers that InfiniBand customers typically specify the exact bandwidth they require.[61] Another large OEM mentions that customers wanting to purchase InfiniBand typically expect to be offered the latest generation of this product with the highest bandwidth available.[62]

(79) However, for the purpose of this decision, the question whether high performance fabrics should be further distinguished based on bandwidth can be left open as it does not materially change the Commission's assessment.

(3) Segmentation within Ethernet-based network interconnects

(80) The Commission has assessed a potential segmentation between the various components composing Ethernet-based network interconnects. For each individual component, the Commission has also assessed a potential sub-segmentation based on bandwidth/speed.

(a) Distinction between individual components

(81) As explained above, in previous decisions, the Commission distinguished between the two main components of Fibre Channel SAN networks, i.e. switches and adapters (NICs). As regards switches, the Commission has previously considered a segmentation based on the different protocols and network technologies that they support.

(82) In line with these precedents, the Commission considers that each of the main individual components composing Ethernet-based network interconnects – i.e., NICs, switches, and cables[63] – constitute a separate product market. In particular, the Commission identifies a separate relevant product market for Ethernet NICs. This conclusion is based on the following elements.

---

59 In addition, in June 2019, Mellanox announced plans to launch a 400 Gb/s InfiniBand fabric (NDR) in 2020 and a 1,000 Gb/s InfiniBand fabric (XDR) sometime thereafter (see: Mellanox, InfiniBand In-Network Computing Technology and Roadmap, June 2019, available at http://www.mellanox.com/solutions/hpc/pdf/InfiniBnd_ISC19_BoF.pdf (at slide 7).

60 See for example agreed minutes of conference call of 12 August 2019 with a large OEM, paragraph 6; agreed minutes of conference call of 9 August 2019 with an end-customer providing HPC services to universities and research centres, paragraph 13; and Intel's submission of 10 September 2019 entitled "Intel response to case team's questions", question 4.b.

61 Replies to Questionnaire Q2 to OEMs, question 26.

62 Replies to Questionnaire Q2 to OEMs, question 26.

63 Network interconnect systems also incorporate a software component. This interconnect software is however typically not sold as a separate product but it is provided as part of the hardware required to deploy interconnect solution (see Form CO, paragraph 329). Therefore, there is no need to consider a separate market for interconnect software.

(83) First, as mentioned above, the Notifying Party acknowledges that a distinction between the components making up Ethernet-based network interconnects may be conceivable from a demand-side substitutability perspective.

(84) Second, the results of the Commission's market investigation confirm that the various components composing Ethernet network interconnect systems fulfil different functions and are not substitutable from the perspective of the customer. For example, a network interconnect supplier explains: "*Network interface Cards (NICs) are not a substitute for switches, and are not always sold together either"*.[64]

(85) Third, a large majority of end customers having expressed their views explained that they do not express a preference as to whether all components composing the network interconnect systems should be procured from one single supplier as a packaged system. Customers explain that they want to leave competition as open as possible.[65] They would therefore typically not require that OEMs procure full Ethernet systems from one single supplier as long as OEMs guarantee the interoperability of the system. A number of customers also explain that they want to be able to select the best individual components. For example, a customer explains that it evaluates each network component "*on its individual merit*".[66]

(86) Fourth, contrary to what is the case for components of high performance fabrics, OEMs also confirm that they source Ethernet network interconnect components both on a standalone basis and as integrated systems.[67]

(87) Fifth, from a supply-side perspective, there are clear limits to the transferability of know-how and technology between Ethernet switches and NICs. In the first place, while Mellanox is a technology leader for Ethernet NICs which translates in high market shares for Ethernet NICs with a bandwidth speed of 25 Gb/s or higher, it does not hold a similarly strong position in Ethernet switches where its market share is [0-5]%.[68] Conversely, although Cisco is the world's largest provider of Ethernet based network switches, based on the Notifying Party's estimate it only has a market share of around [5-10]% in Ethernet NICs with a bandwidth speed of 25 Gb/s or higher. This is indicative of different market conditions and/or different technology requirements between the two main Ethernet network interconnect components, i.e., NICs and switches.

(88) In the second place, a majority of competitors having expressed their view explain that a supplier of Ethernet NICs with a bandwidth speed below 25 Gb/s would need significant investment and time in order to develop and launch Ethernet NICs of 100 Gb/s with RoCE capability that would be able to compete successfully with Mellanox's ConnectX-6 100 Gb/s Ethernet NICs.[69] Based on this feedback, the barrier to entry for a supplier only active in Ethernet switches and/or cables but not supplying NICs would be even higher.

---

64 See agreed minutes of conference call of 25 July 2019 with Cisco.

65 Replies to Questionnaire Q3 to End Customers, questions 27.3 and 27.4.

66 Replies to Questionnaire Q3 to End Customers, questions 27.3 and 27.4.

67 Replies to Questionnaire Q2 to OEMs, questions 17 and 17.1.

68 Notifying Party's response to RFI 16.

69 Replies to Questionnaire Q1 to Competitors, question 25.

(b) Ethernet NICs of 25 Gb/s and higher vs. Ethernet NICs below 25 Gb/s

(89) The term NIC generally refers to a network interface controller, which is a hardware adapter that allows a computer to communicate with a network. In normal datacentre applications, Ethernet NICs are primarily differentiated by the speed at which they can transfer data (bandwidth), which is measured in gigabits per second. Ethernet NICs used in servers support speeds of 1, 10, 25, 50, 100 and now even 200 Gb/s.

(90) As explained above at paragraph 46, in a previous decision, the Commission has already considered a distinction according to bandwidth speed with respect to the SerDes (serializer/deserializer) intellectual property used in high-speed NICs.

(91) In line with this precedent, based on its market investigation, the Commission finds that there is a separate product market for Ethernet NICs with a bandwidth speed of 25 Gb/s or higher which is distinct from the market for Ethernet NICs with a bandwidth speed of less than 25 Gb/s. This conclusion is based on the following elements.

(92) First, from the demand-side perspective, a large majority of end customers and OEMs having expressed their views confirm that there are certain applications, performance needs, or mix of workloads for which customers would only consider Ethernet NICs of at least 25 Gb/s.[70] For example, several customers reported that the need for bandwidth speed of at least 25 Gb/s is particularly acute for hyperscale[71] customers, i.e. mainly cloud service providers.[72]

(93) With respect, specifically, to the possibility to use more 10 Gb/s NICs to achieve the performance of a 25 Gb/s, a customer explained that this would not be an optimal solution because: "*using less speed NIC for ex 10 Gbs will triple the amount of NICs wires and will entail use of more complex (port capacity) switches*".[73]

(94) Second, based on their recent procurement activity (last two years), all OEMs having expressed their views explained that when they delivered/installed GPU-accelerated servers connected with Mellanox Ethernet NICs of at least 25 Gb/s, the only other speeds that could also have been suitable were higher speeds.[74] One OEM notably explained that there may be some substitutability from the customer perspective between a given speed level and the level just below. However, this OEM further explained that this is only the case for speeds above 25 Gb/s. According to this OEM, while a customer willing to purchase 50 Gb/s NICs could potentially consider 40 Gb/s NICs as a suitable alternative, a customer willing to purchase 25 Gb/s NICs would not consider NICs with a speed of 10 Gb/s.

(95) Third, competitors also confirm that it is not credible to compete against Mellanox's Ethernet NICs with a speed of 25 Gb/s or more with Ethernet NICs with a speed of

---

70 Replies to Questionnaire Q3 to End Customers, questions 21 and 21.1; Questionnaire Q2 to OEMs, questions 22 and 22.1.

71 Hyperscale computing refers to the facilities and provisioning required in distributed computing environments to efficiently scale from a few servers to thousands of servers. Hyperscale computing is usually used in environments such as big data and cloud computing.

72 Replies to Questionnaire Q3 to End Customers, questions 21 and 21.1.

73 Replies to Questionnaire Q3 to End Customers, questions 21 and 21.1.

74 Replies to Questionnaire Q2 to OEMs, questions 25, 25.1, 25.2, 25.3 and 25.4.

10 Gb/s. For example, a competitor explains: "*If an end-user has requested an Ethernet NIC of 25 Gb/s or above,* [this competitor] *considers that it would not be possible to credibly compete with an Ethernet NIC of 10 Gb/s against any of the products listed under Questions 22.1 to 22.4.*[75] *The reason is that the end-user would have an identified need that generally cannot be met by an Ethernet NIC of 10 Gb/s*".[76]

(96) Fourth, Mellanox's recent internal documents presented to the board [BUSINESS SECRETS – Information redacted regarding business strategy].[77] This segment corresponds to NICs with a speed of 25 Gb/s or higher. Such a segment is also discussed separately [BUSINESS SECRETS – Information redacted regarding business strategy].[78]

(97) In addition, the Parties' internal documents [BUSINESS SECRETS – Information redacted regarding business strategy] describe and discuss what the Parties identified as [BUSINESS SECRETS – Information redacted regarding business strategy].[79] In particular, NVIDIA noted that this [BUSINESS SECRETS – Information redacted regarding business strategy].[80]

(98) Fifth, from a supply-side perspective, a majority of competitors consider that there are significant barriers in terms of time and costs for a supplier of Ethernet NICs with a bandwidth speed below 25 Gb/s to develop and launch Ethernet NICs of 100 Gb/s with RoCE capability that would be able to compete successfully with Mellanox's ConnectX-6 100 Gb/s Ethernet NICs.[81] Intel also explains that whereas there might be supply-side substitution between 25 Gb/s and higher bandwidth, "[t]*he 25 Gbps technology is foundational to higher Ethernet speeds, as the 50 Gbps and 100 Gbps Ethernet NICs achieve their speeds through multiple 25 Gbps lanes*".[82]

(c) Other Ethernet components (switches and cables)

(99) Ethernet switches and Ethernet cables could potentially be further segmented based on bandwidth between switches and cables of 25 Gb/s and higher on the one hand and switches and cables below 25 Gb/s.

(100) However, for the purpose of this decision the precise product market definition can be left open, as the Transaction does not raise serious doubts as to its compatibility with the internal market or the functioning of the EEA Agreement with regard to Ethernet switches and Ethernet cables, and any segments therein, under any plausible market definition.

---

75 Mellanox ConnectX Ethernet NICs of 25 Gb/s, Mellanox ConnectX Ethernet NICs of 50 Gb/s, Mellanox ConnectX Ethernet NICs of 100 Gb/s, and Mellanox ConnectX Ethernet NICs of 200 Gb/s.

76 Replies to Questionnaire Q1 to Competitors, question 22 and sub-questions.

77 E.g., Mellanox internal documents […], slides 7 and 8; […], slides 69-78; […], slides 49-50.

78 E.g., Form CO, Annex 5.4(ii) – 04, […], slides 3 and 5.

79 E.g., Mellanox internal document, […], pages 29-30; NVIDIA internal document, Form CO Annex 5.4(ii) – 03, […].

80 NVIDIA internal document, Form CO, Annex 5.4(ii) – 03, […].

81 Replies to Questionnaire Q1 to Competitors, question 25

82 Intel's non-confidential submission of 11 September 2019, entitled "Intel responses to case team's questions", page 18

(4) Conclusion

(101) The Commission considers that there are two distinct product markets for (i) high performance fabrics and (ii) Ethernet-based network interconnects.

(102) Within the market for high performance fabrics, the Commission does not consider a further segmentation based on (a) different components meaningful. The question whether high performance fabrics should be further segmented based on (b) protocols or (c) bandwidth/speed can be left open, since any such further segmentation would not materially change the Commission's assessment in this case.

(103) As for the Ethernet-based interconnects, the Commission considers that each of the main individual components composing Ethernet-based network interconnects – i.e., NICs, switches, and cables – constitute separate product markets. In particular, the Commission identifies a separate relevant product market for Ethernet NICs. The Commission considers that this market should be further segmented in a separate product market for Ethernet NICs with a bandwidth speed of 25 Gb/s or higher, which is distinct from the market for Ethernet NICs with a bandwidth speed of less than 25 Gb/s. Finally, for the purpose of this decision, the precise product market definition for Ethernet switches and Ethernet cables, as well as the question whether these markets should be further segmented, can be left open, as the Transaction does not raise serious doubts as to its compatibility with the internal market under any market that the Commission considers plausible.

4.3.3. Geographic market definition

(104) In previous decisions, the Commission has considered that the geographic market for all categories of network interconnect products (including IP/Ethernet switches and routers) to be either EEA-wide or worldwide in scope.[83] In *Broadcom/Brocade*,[84] respondents to the market investigation unanimously considered that the geographic market for IP/Ethernet switches and routers was worldwide given the global nature of both supply and demand.

(105) The Notifying Party submits that the relevant geographic market for datacentre network interconnects should be defined as worldwide.[85] This should remain the case even for potentially narrower product markets.[86]

(106) The majority of competitors, OEMs and end customers that replied to the market investigation confirmed that network interconnects are supplied on a worldwide basis, irrespective of the location of the component vendor or the location of the end-customer.[87] Moreover, the majority of competitors and the large majority of end

---

83 Commission decision of 12 April 2017 in case M.8314 – *Broadcom/Brocade*, paragraphs 58. Commission decision of 19 September 2008 in case M.5300 – *Gores Group/Siemens Enterprise Communications*, paragraph 14.

84 Commission decision of 12 April 2017 in case M.8314 – *Broadcom/Brocade*, paragraphs 59.

85 Form CO, paragraph 360.

86 Form CO, paragraph 367.

87 Replies to Questionnaire Q1 to Competitors, questions 28 and 28.1; Questionnaire Q2 to OEMs, questions 32 and 32.1; Questionnaire Q3 to End Customers, questions 30 and 30.1.

customers that expressed a view confirmed that the conditions of competition do not differ depending on the location of the datacentre of the end customer.[88]

(107) In light of the results of the market investigation, for the purposes of this decision, the Commission considers that the geographic scope of the various markets for network interconnect products is worldwide.

4.4. Datacentre servers

4.4.1. Introduction

(108) Datacentres are a collection of servers that are connected by a network and that work together to process/compute workloads. As such, servers are the computing power of datacentres.[89] They typically contain CPUs, network interconnects and optional accelerators together in a system.

(109) NVIDIA offers a family of server systems (DGX-1, DGX-2 and DGX-Station) that perform GPU-accelerated AI and deep learning ("**DL**") training and inference applications, among others. The main building block of the DGX servers is HGX, which combines a number of NVIDIA GPUs, connected with NVLINK and NVSwitches, enabling them to function as a single unified accelerator.[90] GDX-1 contains the HGX-1, a board with eight Tesla V100 GPUs, Core Intel Xeon CPUs, and four InfiniBand NICs. DGX-2 contains the HGX-2, which includes 16 NVIDIA Tesla V100 GPUs, dual-socket Intel Xeon CPUs, and eight InfiniBand NICs.[91]

(110) According to the Notifying Party, the DGX family is a "reference architecture" platform for NVIDIA to continue to innovate and demonstrate GPU innovations to server OEMs/ODMs, thereby generating demand for its GPUs. NVIDIA provides that innovation and the building blocks of its DGX servers to OEMs, ODMs and CSPs to use in their own server offerings. In addition, NVIDIA offers DGX servers for sale, but, according to the Notifying Party, these servers are not intended to displace any sales from NVIDIA's OEM/ODM partners.[92]

4.4.2. Product market definition

4.4.2.1. Commission precedents

(111) In past decisions, the Commission has considered a segmentation of datacentre servers by price band: (a) entry level (below USD 100 000), (b) mid-range (USD

---

88 Replies to Questionnaire Q1 to Competitors, questions 29 and 29; Questionnaire Q3 to End Customers, questions 31 and 31.1.

89 Commission Decision of 29 February 2016 in case M.7861 – *Dell/EMC*, paragraph 38.

90 Notifying Party's response to RFI 5, question 7.

91 Form CO, paragraph 179. See also https://www.nvidia.com/content/dam/en-zz/Solutions/Data-Center/dgx-1/dgx-1-print-infographic-738238-nvidia-web.pdf and https://www nvidia.com/content/dam/en-zz/Solutions/Data-Center/dgx-1/dgx-2-datasheet-us-nvidia-955420-r2-web-new.pdf.

92 Form CO, paragraphs 32 and 179; see also Notifying Party's response to RFI 5, question 7.

100 000 – USD 999 999), and (c) high-end (USD 1 million and above). The Commission ultimately left the product market decision open.[93]

(112) In *Dell/EMC*, the Commission noted that the market investigation did not provide a clear result as to a possible segmentation of datacentre servers by operating systems or by the applications they serve.[94]

4.4.2.2. Notifying Party's views

(113) The Notifying Party considers that the relevant product market should encompass all datacentre servers.[95] However, in this case, the precise product market definition can be left open.[96]

(114) First, according to the Notifying Party, all datacentre servers have the same functions: they are used to process data in the datacentre, which can be achieved with differently priced servers. However, if a further division of the market for datacentre market were necessary, the Notifying Party submits that a segmentation according to price band would be the most sensible, as it would be in line with the Commission's precedents. In that situation, the Notifying Party considers that NVIDIA's DGX servers would belong to the potential market for mid-range servers.[97]

(115) Second, the Notifying Party argues that the market for datacentre servers should not be segmented according to the particular workloads they serve because (i) suppliers do not know which applications end-customers will accelerate with their server and (ii) there is no workload that only a highly accelerated server (such as NVIDIA's DGX server) could handle.[98]

4.4.2.3.Commission's assessment

(116) The Commission has considered whether the market for datacentre servers could be further segmented according to price bands or to the applications/end uses for which the datacentre servers are designed or used. During the market investigation, these possible market segmentations were tested.

(117) First, the market investigation regarding a possible segmentation according to price bands provided mixed views. While a number of respondents considered that this segmentation was appropriate, others disagreed with it. For example, a few respondents indicated that such segmentation is "*typically used*" or is "*quite common*" and "*similar to what IDC provides*".[99] However, Tech Data Europe (OEM) explained that "*it is not necessary to define segments as narrowly as high-end, mid-range and low-end servers, as customers switch between servers of*

---

93 Commission decision of 29 February 2016 in case M.7861 – *Dell/EMC*, paragraphs 38-42 and 45; Commission decision of 21 January 2010 in case M.5529 – *Oracle/Sun Microsystems*, recital 941; Commission decision of 31 January 2002 in case M.2609 – *HP/Compaq*, paragraphs 20-22.

94 Commission decision of 29 February 2016 in case M.7861 – *Dell/EMC*, paragraph 42.

95 Form CO, paragraph 352.

96 Form CO, paragraph 359.

97 Form CO, paragraphs 404-407.

98 Form CO, paragraph 355.

99 Replies to Questionnaire Q3 to End Customers, questions 32.1 and 32.1.1.

different types, and distributors typically sell all types".[100] Similarly, HPE stated that "[t]*here are many different reasons that a server may be prices in a certain way and all servers in a similar price bands are not substitutes for one another*".[101]

(118) Second, the majority of the respondents to the market investigation that expressed a view did not consider that a segmentation of datacentre servers based on the applications/end use they serve is appropriate.[102] For example, an OEM explained that "*a segmentation by type of application is not appropriate as servers may be used across multiple applications and customers usually do not request significantly different types of servers based on the use-case.*"[103] Similarly, an end customer indicated that "*any size* [of servers] *can serve any applications*".[104] Moreover, some of the respondents that expressed a view listed a number of suppliers offering alternatives to NVIDIA's DGX servers, including, for example, HPE, IBM, Dell, Intel, Atos and Lenovo.[105] Finally, competitors such as Oracle and IBM confirmed that it would be relatively easy for them to start supplying datacentre servers offering the same level of performance or that would be suitable for the applications/end uses for which DGX servers are used.[106]

(119) In light of the above, the Commission considers that the market for datacentre servers should not be segmented according to the applications/end uses for which the datacentre servers are designed or used. In addition, the Commission considers that, for the purpose of this decision, it can be left open whether the market for datacentre servers should be further segmented according to price bands, as the Transaction does not raise serious doubts as to its compatibility with the internal market in any plausible product markets, even in a plausible mid-range server market where NVIDIA's position would be stronger.

4.4.3. Geographic market definition

(120) In past decisions, the Commission found the market for datacentre servers to be at least EEA-wide if not worldwide.[107]

(121) The Notifying Party submits that the relevant geographic market for datacentre servers should be defined as worldwide.[108] This should remain the case even for a potentially narrower product market for mid-range servers.[109]

---

100 Tech Data Europe GmbH's reply to Questionnaire Q2 to OEMs, question 34.1.1.

101 Replies to Questionnaires Q2 to OEMs, question 34.1.; Questionnaire Q3 to End Customers, question 32.1.

102 Replies to Questionnaires Q2 to OEMs, question 34.2.; Questionnaire Q3 to End Customers, question 32.2.

103 Replies to Questionnaire Q2 to OEMs, question 34.2.1.

104 Replies to Questionnaire Q3 to End Customers, question 32.2.1.

105 Replies to Questionnaire Q2 to OEMs, questions 36.1 and 36.2; Questionnaire Q3 to End Customers, question 33.3.

106 Replies to Questionnaire Q2 to OEMs, question 36.3.

107 Commission decision of 29 February 2016 in case M.7861 – *Dell/EMC*, paragraph 44; Commission decision of 21 January 2010 in case M.5529 – *Oracle/Sun Microsystems*, recitals 941 and 950; Commission decision of 31 January 2002 in case M.2609 – *HP/Compaq*, paragraph 23.

108 Form CO, paragraph 360.

109 Form CO, paragraph 367.

(122) The results of the market investigation indicate that the geographic scope of the market for datacentre servers is most likely worldwide.[110] A number of OEMs confirmed that they sell datacentre servers worldwide. Moreover, OEM respondents explained that "[t]*here are few country-specific reasons which would prevent us from looking at the market global*", "[t]*transport costs are low*" and "*prices are similar or identical and customers procure servers on an EEA-wide or even worldwide level*".[111] Similarly, end customers confirmed that "*clients are international*" and that the servers they use are made available worldwide by global suppliers.[112]

(123) In light of the results of the market investigation, for the purposes of this decision, the Commission considers that the geographic scope of the datacentre server markets (including a potential market for mid-range servers) is worldwide.

5. COMPETITIVE ASSESSMENT

5.1. Analytical framework

(124) Under Article 2(2) and (3) of the Merger Regulation and Annex XIV to the EEA Agreement, the Commission declares a proposed concentration incompatible with the internal market and with the functioning of the EEA Agreement if that concentration would significantly impede effective competition in the internal market or in a substantial part of it, in particular through the creation or strengthening of a dominant position.

(125) Under Article 57(1) of the EEA Agreement, the Commission declares a proposed concentration incompatible with the EEA Agreement if that transaction creates or strengthens a dominant position as a result of which effective competition would be significantly impeded within the territory covered by the EEA Agreement or a substantial part of it.

(126) In this respect, a merger may entail horizontal and/or non-horizontal effects. Horizontal effects are those deriving from a concentration where the undertakings concerned are actual or potential competitors of each other in one or more of the relevant markets concerned. Non-horizontal effects are those deriving from a concentration where the undertakings concerned are active in different relevant markets.

(127) The Commission appraises non-horizontal effects in accordance with the guidance set out in the relevant notice, that is to say the Non-Horizontal Merger Guidelines.[113]

(128) As regards non-horizontal mergers, two broad types of such mergers can be distinguished: vertical mergers and conglomerate mergers.[114] Vertical mergers

---

110 Replies to Questionnaires Q2 to OEMs, question 38; Questionnaire Q3 to End Customers, question 35.

111 Replies to Questionnaires Q2 to OEMs, questions 38 and 38.1

112 Replies to Questionnaire Q3 to End Customers, questions 35 and 35.1.

113 Guidelines on the assessment of non-horizontal mergers under the Council Regulation on the control of concentrations between undertakings ("Non-Horizontal Merger Guidelines"), OJ C 265, 18.10.2008, p. 6-25.

114 Non-Horizontal Merger Guidelines, paragraph 3.

involve companies operating at different levels of the supply chain.[115] Conglomerate mergers are mergers between firms that are in a relationship that is neither horizontal (as competitors in the same relevant market) nor vertical (as suppliers or customers).[116]

(129) In this particular case, the Transaction does not give rise to any horizontal overlaps between the Parties' activities, but results in a vertical and a conglomerate relationship. Accordingly, the Commission will only examine whether the Transaction is likely to give rise to non-horizontal effects. In particular, the Commission will assess potential conglomerate and vertical effects.

5.2. Conglomerate non-coordinated effects

(130) NVIDIA and Mellanox are active in closely related markets. They both supply components used in datacentres or server clusters, in particular those used for HPC. NVIDIA's discrete datacentre GPUs equip servers that constitute certain (parts of) datacentres (also referred to as GPU-accelerated server clusters) to accelerate a number of applications, typically computations that require massive parallel execution of relatively simple computational tasks. They are necessarily used side-by-side with CPUs, which are always present in datacentre servers. Servers within datacentres are connected to each other through network interconnect solutions, offered among others by Mellanox, and composed of network cables, connecting the NICs within servers to network switches. NVIDIA's discrete datacentre GPUs and Mellanox's network interconnect solutions are therefore complementary components. which can be purchased directly or indirectly via OEMs/ODMs by the same set of customers for the same end use (HPC datacentres).

(131) In this decision, the Commission carries out three assessments as regards the conglomerate relationships identified above.

(132) The first assessment consists in determining whether the Transaction would likely confer on the Merged Entity the ability and incentive to leverage Mellanox's potentially strong market position in both the market for high-performance fabric (with its InfiniBand fabric) and in the market for Ethernet NICs of at least 25 Gb/s into the discrete datacentre GPU market, and whether this would have a significant detrimental effect on competition in the discrete datacentre GPU market, thus causing harm to customers.

(133) The second assessment consists in determining whether the Transaction would likely confer on the Merged Entity the ability and incentive to leverage NVIDIA's potentially strong market position in the discrete datacentre GPU market into any possible network interconnect markets, and whether this would have a significant detrimental effect on competition in the network interconnect markets, thus causing harm to customers. This assessment is done overall, rather than for each potential network interconnect product market that could potentially be the target of the Merged Entity's leveraging strategy. This is because most of the competitive assessment is similar irrespective of the exact network interconnect product. The only difference relates to the assessment of the Merged Entity's incentive. The

115 Non-Horizontal Merger Guidelines, paragraph 4.

116 Non-Horizontal Merger Guidelines, paragraph 5.

Commission will consider the various interconnect products when assessing the Merged Entity's incentive.

(134) The third assessment consists in determining whether the Merged Entity would likely have the ability and incentive to misuse commercially sensitive information that it obtains from competing GPU and network interconnect suppliers (in the context of cooperation with these competitors to ensure interoperability of their respective products) to favour its own position on the discrete datacentre GPU and/or network interconnects relevant markets.

5.2.1. Legal framework

(135) According to the Non-Horizontal Merger Guidelines, in most circumstances, conglomerate mergers do not lead to any competition problems.[117]

(136) However, foreclosure effects may arise when the combination of products in related markets may confer on the merged entity the ability and incentive to leverage a strong market position from one market to another closely related market by means of tying or bundling or other exclusionary practices. [118]

(137) In assessing the likelihood of conglomerate effects, the Commission examines, first, whether the merged firm would have the ability to foreclose its rivals, second, whether it would have the economic incentive to do so and, third, whether a foreclosure strategy would have a significant detrimental effect on competition. In practice, these factors are often examined together as they are closely intertwined.[119]

(138) Mixed bundling refers to situations where the products are also available separately, but the sum of the stand-alone prices is higher than the bundled prices.[120] Tying refers to situations where customers that purchase one good (the tying good) are required also to purchase another good from the producer (the tied good). Tying can take place on a technical or contractual basis.[121] Tying and bundling as such are common practices that often have no anticompetitive consequences. Nevertheless, in certain circumstances, these practices may lead to a reduction in actual or potential rivals' ability or incentive to compete. Foreclosure may also take more subtle forms, such as the degradation of the quality of the standalone product.[122] This may reduce the competitive pressure on the Merged Entity allowing it to increase prices.[123]

(139) In order to be able to foreclose competitors, the merged entity must have a significant degree of market power, which does not necessarily amount to dominance, in one of the markets concerned. The effects of bundling or tying can only be expected to be substantial when at least one of the merging parties' products is viewed by many customers as particularly important and there are few relevant alternatives for that product.[124] Further, for foreclosure to be a potential concern, it

117 Non-Horizontal Merger Guidelines, paragraph 92.
118 Non-horizontal Merger Guidelines, paragraph 93.
119 Non-horizontal Merger Guidelines, paragraph 94.
120 Non-horizontal Merger Guidelines, paragraph 96.
121 Non-horizontal Merger Guidelines, paragraph 97.
122 Non-horizontal Merger Guidelines, paragraph 33.
123 Non-horizontal Merger Guidelines, paragraph 93.
124 Non-horizontal Merger Guidelines, paragraph 99.

must be the case that there is a large common pool of customers, which is more likely to be the case when the products are complementary.[125]

(140) The incentive to foreclose rivals through bundling or tying depends on the degree to which this strategy is profitable.[126] Bundling and tying may entail losses or foregone revenues for the merged entity.[127] It may also increase profits by gaining market power in the tied goods market, protecting market power in the tying good market, or a combination of the two.[128]

(141) It is only when a sufficiently large fraction of market output is affected by foreclosure resulting from the concentration that the concentration may significantly impede effective competition. If there remain effective single-product players in either market, competition is unlikely to deteriorate following a conglomerate concentration.[129] The effect on competition needs to be assessed in light of countervailing factors such as the presence of countervailing buyer power or the likelihood that entry would maintain effective competition in the closely related markets concerned.[130]

5.2.2. Affected markets

(142) NVIDIA is active in the discrete datacentre GPU market, while Mellanox is active in the various network interconnect markets (depending on the exact segmentation) which are neighbouring markets closely related to the discrete datacentre GPU markets.

(143) Table 1 below presents NVIDIA's and its competitors' market shares on a worldwide market for discrete datacentre GPUs from 2016 to 2018. It follows that NVIDIA has a market share of [90-100]%. Mellanox is active in various network interconnect markets that are neighbouring markets closely related to the discrete datacentre GPU market. As a result, the discrete datacentre GPU market as well as all possible relevant network interconnect markets are affected.

Table 1: Value market shares in discrete datacentre GPUs, worldwide, 2016-2018

Vendor	2016	2017	2018
NVIDIA	**[90-100]%**	**[90-100]%**	**[90-100]%**
AMD	[0-5]%	[0-5]%	[5-10]%
Total	**100%**	**100%**	**100%**

Source: IDC, NVIDIA actual sales data

(144) Currently, NVIDIA has only one competitor (AMD) in the discrete datacentre GPU market, which has so far only managed to gain [5-10]% market share in 2018 (see Table 1 above). However, according to the Notifying Party, the discrete datacentre

125 Non-horizontal Merger Guidelines, paragraph 100.
126 Non-horizontal Merger Guidelines, paragraph 105.
127 Non-horizontal Merger Guidelines, paragraph 106.
128 Non-horizontal Merger Guidelines, paragraph 108.
129 Non-horizontal Merger Guidelines, paragraph 113.
130 Non-horizontal Merger Guidelines, paragraph 114.

GPU market is characterised by two major, recent developments, that are radically transforming its competitive dynamics.

(145) The first is the rise of AMD. A long-time participant in the GPU segment for gaming GPUs, AMD is leveraging its gaming Radeon GPU architecture for datacentre GPUs. In November 2018, AMD launched a new discrete datacentre GPU model: **Radeon Instinct**, which competes directly with NVIDIA's datacentre GPUs. At the same time, AMD announced an ambitious product roadmap including the next Instinct generation as well as a commitment to on-going launches at a "*predictable cadence with generational performance gains.*"[131] AMD's emergence as a strong competitor is evidenced by its recent successes at the top-levels of HPC computing. For instance, AMD announced that the upcoming **Frontier datacentre** at Oak Ridge National Laboratory will rely on AMD's Radeon Instinct GPUs (and AMD EPYC CPUs). This will be the fastest datacentre in the world when it comes online in 2021 – and will run scientific, AI, and data analytics workloads. Beyond this success, AMD has also recently announced other top-end datacentre wins, including European datacentres.[132]

(146) The second is the entry of Intel. As explained by Intel, Intel is developing a new GPU to compete with NVIDIA's GPUs for computational workloads in datacentres. Intel intends to enter in two stages. It plans to release a discrete GPU for graphics rendering workloads on PCs in 2020. Intel expects that this product, which will be the first new entry into the GPU market in nearly two decades, will also have limited deployment in datacentres. Intel plans to follow that with the release in 2021 of a discrete datacentre GPU designed specifically for computational uses in servers (also known as a "GPGPU", short for general purpose GPU).[133] According to the Notifying Party, Intel's entry as a credible competitor of NVIDIA is evidenced by its recent win in the tender organised by the U.S. Department of Energy for the upcoming **Aurora datacentre** at Argonne National Laboratory. This will be one of the fastest datacentres in the world by 2021 when it comes online. The U.S. Department of Energy reportedly selected Intel Xe GPUs (and Intel Xeon CPUs).[134]

(147) Tables 2 and 3 below present Mellanox's and its competitors' market shares in the market for Ethernet network NICs with a bandwidth of at least 25 Gb/s and in the market for high-performance fabric (where Mellanox is active with its InfiniBand fabric).

---

131 See D. Wang, "AMD Next Horizon", available at https://www.amd.com/system/files/documents/next_horizon_david_wang_presentation.pdf.

132 Lawrence Livermore National Lab (USA), NERSC (USA), the High-Performance Computing Center of the University of Stuttgart (Germany), the CSC – IT Center for Science Ltd (Finland), as well as Eni S.p.a. (Italy). See the Notifying Party's response to RFI 1, paragraph 184; Replies to Questionnaire Q3 to end customers, question 50.1.1.

133 Intel's submission of 10 September 2019 entitled "Intel response to case team's query regarding foreclosure mechanism resulting from NVIDIA's acquisition of Mellanox", page 1.

134 Form CO, paragraph 500.

Table 2: Value market shares in datacentre Ethernet network NICs with a bandwidth of at least 25 Gb/s, worldwide, 2016-2018

Vendor	2016	2017	2018
Mellanox	**[70-80]%**	**[60-70]%**	**[60-70]%**
Intel	[5-10]%	[10-20]%	[10-20]%
Broadcom	[0-5]%	[5-10]%	[5-10]%
Cisco	[5-10]%	[5-10]%	[5-10]%
Solarflare	[0-5]%	[0-5]%	[0-5]%
QLogic	[0-5]%	[0-5]%	[0-5]%
Chelsio	[0-5]%	[0-5]%	[0-5]%
Total	**100%**	**100%**	**100%**

Source: CREHAN, Mellanox actual sales data

Table 3: Market shares in high performance fabric, worldwide, 2016-2018[135]

Vendor	2016	2017	2018
Mellanox	**[60-70]%**	**[60-70]%**	**[60-70]%**
Intel	[20-30]%	[20-30]%	[20-30]%
Cray	[5-10]%	[10-20]%	[10-20]%
Bull	[0-5]%	[0-5]%	[0-5]%
Other	[0-5]%	[0-5]%	[0-5]%
Total	**100%**	**100%**	**100%**

Source: TOP 500 lists (June 2017, June 2018, June 2019)

(148) As can be seen from Tables 2 and 3 above, Mellanox has a market share of [60-70]% in the market for high-performance fabric (with its InfiniBand fabric) and of [60-70]% in the market for Ethernet NICs of at least 25 Gb/s. In all other plausible network interconnect market segments, Mellanox's market share is low and in any event significantly lower than 30%.[136] As NVIDIA is active in the discrete

[135] The Parties were unable to provide value market shares on a high performance fabric market, because CREHAN does not report sales data on such a market. Therefore, the Parties provided estimated market shares based on a count of the number of supercomputers equipped with each interconnect protocol, excluding Ethernet, in the TOP 500 lists. The TOP 500 lists are produced bi-annually, in June and in November. The Parties originally provided the shares of supercomputers equipped with each interconnect non-Ethernet protocol in the November lists of 2016, 2017 and 2018, to reflect the share at the end of each corresponding year. However these shares represent the share of each non-Ethernet protocol in the installed base of each list rather than the shares of each non-Ethernet protocol in the number of installations made in a given year. Therefore, the Commission reiterated the calculation of market shares based on the TOP 500 lists focusing only on the installations of a given year. For each year t, in order to cover the entire year, the Commission used the June t+1 list (e.g. June 2019 list to calculate shares for the year 2018) and calculated the share of each non-Ethernet protocol among those supercomputers of this list that were installed in year t.

[136] Mellanox's market share at the worldwide level is [0-5]% in a possible market for Ethernet switches, [0-5]% in a possible market for Ethernet switches of at least 25 Gb/s, and [0-5]% in a possible market for Ethernet switches below 25 Gb/s.
No public data exist for the market for Ethernet cables. Therefore the Parties could not provide estimates of Mellanox's market share on Ethernet cable markets. However, the Parties submit that Mellanox's

datacentre GPU market, which is closely related to both the markets for high-performance fabric and for Ethernet NICs of at least 25 Gb/s, it can be concluded that both the discrete datacentre GPU market as well as the markets for high-performance fabric and for Ethernet NICs of at least 25 Gb/s are affected.

5.2.3. Leveraging the position of Mellanox in the markets for high-performance fabric and for Ethernet NICs of at least 25 Gb/s into the discrete datacentre GPU market where NVIDIA is active

5.2.3.1. Potential concern

(149) The Commission has assessed a potential competition concern whereby the Merged Entity would leverage Mellanox's potentially strong position in the plausible markets for high performance fabric and for Ethernet NICs of at least 25 Gb/s, into the market for discrete datacentre GPUs where NVIDIA is active and thereby foreclose competitors on the discrete datacentre GPUs market, thus causing harm to customers.

(150) The Commission has assessed in particular the ability and the incentive of the Merged Entity to engage in the following tying/bundling practices:

- technical tying: differentiating the degree of technical compatibility and therefore overall performance of the Merged Entity's joint solution compared to mix-and-match solutions involving only one of its products; and/or

- contractual tying: imposing the purchase of NVIDIA GPUs if the customer wants to purchase Mellanox's InfiniBand fabric and/or Ethernet NICs of at least 25 Gb/s; and/or

- mixed bundling: incentivising the joint purchase of the Merged Entity's own products by offering higher prices for mix-and-match solutions involving only one of its products as compared to the bundle.

(151) Both AMD[137] and Intel raised concerns that they may be foreclosed from the discrete datacentre GPU markets due to one or a combination of the three practices described above. In particular, Intel claims that, for demanding HPC and AI deep learning training server deployments, customers require their servers to be connected with a high performance fabric and/or Ethernet NICs of at least 25 Gb/s and that customers do not have credible alternatives to Mellanox.[138]

(152) The Commission has assessed specifically whether the Merged Entity would have the ability and incentive to foreclose enough discrete datacentre GPU market output to hinder Intel's effective long-term entry and AMD's expansion into the discrete datacentre GPU market.

---

market shares on a market for datacentre Ethernet cables, on a market for datacentre Ethernet cables of at least 25 Gb/s, and on a market for datacentre Ethernet cables below 25 Gb/s are well below 30%.
Finally, Mellanox's market share at the worldwide level in a plausible market for Ethernet NICs below 25 Gb/s is [0-5]%. See Notifying Party's response to RFI 16.

137 Agreed minutes of the conference call of 9 August 2019 with AMD.

138 Intel's submission of 10 September 2019 entitled "Intel response to case team's questions", question 7.

(153) The reason why the Commission has assessed the potential leveraging from two distinct markets together is that GPU-accelerated servers, depending on the requirements of the end-customers, are in practice connected with various types of interconnect solutions. In particular, some GPU-accelerated server clusters are connected with high-performance fabrics (including Mellanox's InfiniBand fabric), while others are connected with Ethernet network interconnect solutions composed among others of NICs of different speeds, including 25 Gb/s and above. NICs are particularly important as far as interoperability between GPUs and the overall network interconnect is concerned, because NICs are the piece of hardware allowing the various servers within a datacentre to communicate with each other. When servers are accelerated with GPUs, NICs may need to interact directly with GPUs.

5.2.3.2. Notifying Party's view

(154) The Notifying Party submits that the Merged Entity will not have the ability and incentive to leverage Mellanox's potentially strong market position in any plausible markets into the market for discrete datacentre GPUs where NVIDIA is active. In any event, the Notifying Party submits that any putative leveraging could not lead to anticompetitive foreclosure of NVIDIA's rivals. The reasons are the following.

(1) As regards ability

(a) As regards Mellanox's alleged market power

(155) First, the Notifying Party argues that the Merged Entity will not have the ability to anticompetitively leverage Mellanox's market position post-Transaction, because Mellanox lacks market power in the supply of network interconnect products. Even in the narrow market segments identified in Section 4.3.2. which are limited to (i) high performance interconnect fabrics, and (ii) Ethernet NICs with a data speed of 25 Gb/s or higher, the Notifying Party argues that Mellanox is subject to strong competitive constraints.

– As regards Mellanox's InfiniBand fabric

(156) In the first place, the Notifying Party argues that Mellanox's InfiniBand fabric is subject to significant competitive constraints from other high-performance fabrics, such as Intel's Omni-Path, Cray's Aries, Gemini and Slingshot, and Fujitsu's Tofu. In particular, the Notifying Party argues that Cray's Slingshot fabric will compete strongly with InfiniBand in the foreseeable future. Also, according to the Notifying Party, while Intel has discontinued the next generation of Omni-Path, it is maintaining support for the existing products and Mellanox expects that it will continue to exercise real constraint for at least the next two years.

(157) In the second place, the Notifying Party argues that Mellanox's InfiniBand fabric is subject to significant competitive constraints from Ethernet. According to the Notifying Party, InfiniBand is no longer protected by any material technical advantages, such as low latency, from Ethernet competition. Therefore, the Notifying Party claims that there is no application for which a particular type of interconnect solution, such as InfiniBand, would be the only option available to customers.

(158) According to the Notifying Party, Mellanox has actually contributed to this trend with the launch of RDMA over converged Ethernet ("RoCE"). RDMA provides direct memory access from the memory of one host to the memory of another host while reducing the burden on the Operating System and CPU. This boosts performance and reduces latency. With RoCE, Mellanox shared this InfiniBand advantage with Ethernet, thus accelerating Ethernet's uptake. In addition, and on account of customers' preferences, Mellanox made this technology open-source, allowing all Ethernet suppliers to take advantage of it.

– As regards Mellanox's Ethernet NICs of at least 25 Gb/s

(159) In the first place, the Notifying Party argues that even on a narrowly defined market for Ethernet NICs with a speed of 25 Gb/s and more, Mellanox faces significant and growing competitive constraints which are not yet fully visible in a backward-looking/static market share analysis. Mellanox was the first producer to launch an Ethernet NIC of 25 Gb/s and above. It started in 2012. Since then, new providers have entered the market, including Intel, Cisco, Broadcom, and Chelsio taking away market share from Mellanox.

(160) In the second place, the Notifying Party submits that some of Mellanox's largest historical customers are now building their own Ethernet NICs. This includes […]. These companies used to source Ethernet NICs from Mellanox and others but have since decided to build their own NICs in-house. Others can follow their lead. These "in-house" solutions also exert competitive pressure on Mellanox, which is not reflected at all on market share data.

(161) In the third place, the Notifying Party argues that [BUSINESS SECRETS – Information redacted regarding business strategy].

(b) As regards the Merged Entity's ability to engage in tying/bundling practices

(162) Second, as regards technical tying, the Notifying Party claims that there are no practicable means through which the Parties could degrade interoperability between Mellanox's network interconnect and the discrete datacentre GPUs of NVIDIA's competitors'.[139] Moreover, the Notifying Party claims that it is commercially necessary to continue promoting interoperability and compatibility with third parties, in particular with Intel and AMD, because the Parties are dependent on the CPU makers, who have the truly indispensable products that form the backbone of any system, and who could retaliate.[140] The Merged Entity would not have the technical ability to degrade its competitors' performance because datacentres use open standards and systems that the Parties do not control.

(163) Third, the Notifying Party argues that the procurement structure of this industry precludes the ability to leverage. In the first place, the Notifying Party argues that end-customers are large, sophisticated enterprises that exert considerable countervailing buyer power in a *bidding market* with credible alternatives. In the second place, the Notifying Party argues that the Parties mainly sell through intermediaries (OEMs and ODMs), limiting the possibility to engage in mixed

139 Form CO, paragraphs 561-580.

140 Form CO, paragraphs 556-560.

bundling strategies, given that OEMs could easily buy the "bundle" from the Parties but then sell the components separately to their customers. In the third place, customers often buy processing and network interconnect products in distinct transactions, not synchronously, again limiting the possibility to engage in mixed bundling strategies, as the combined entity would need to make separate offers for the different types of products, in order to conform with the customer's purchasing practices.[141]

As regards incentives

(164) According to the Notifying Party, the Merged Entity would also not have the incentive to degrade the interoperability of Mellanox's InfiniBand fabric and/or Ethernet NICs of at least 25 Gb/s with third parties' discrete datacentre GPUs or to raise Mellanox's products' relative price when combined with third party GPUs.

(165) First, the Notifying Party argues in general terms that the Parties have strong commercial incentives to continue interoperating with other datacentre component suppliers, including their rivals. According to the Notifying Party, by fostering compatibility, datacentre component suppliers contribute to growing both the overall market and their addressable share of it. This is particularly so because OEMs have a strong preference for using components that interoperate with other components (including the OEM's own). The prevalence of standardized interfaces, such as Peripheral Component Interconnect Express ("PCIe"), and protocols, like Ethernet, illustrates the necessity for suppliers to maintain interoperability for their products to be viable.[142]

(166) Second, the Notifying Party argues that the cost of foreclosing suppliers of rival processing solutions in terms of lost interconnect sales would outweigh any benefit from increased GPU sales. However, the only scenario considered by the Notifying Party is a scenario whereby the Merged Entity would refuse all together to sell Mellanox's interconnects unless a customer also buys NVIDIA's GPUs. The Notifying Party claims that this would be unprofitable because many Mellanox customers do not need GPUs.[143]

(167) Third, the Notifying Party argues that any leveraging strategy would lead to retaliation from Intel and AMD, which control the ecosystems attached to their CPUs. Moreover, the Notifying Party argues that any leveraging strategy (assuming the Merged Entity's dominance on some interconnect markets) would expose the Parties to antitrust scrutiny and possible follow-on litigation. This risk acts as a significant deterrent to carry out any putative anti-competitive foreclosure strategy.[144]

(168) Fourth, the Notifying Party explains that if Mellanox had the ability and incentive to leverage an alleged strong market position on a hypothetical market for Ethernet NICs of at least 25 Gb/s, it would already have done so into the adjacent hypothetical market for Ethernet switches of at least 25 Gb/s. NICs and switches are

141 Form CO, paragraphs 581-587.
142 Form CO, paragraphs 590-598.
143 Form CO, paragraphs 607-610.
144 Form CO, paragraphs 611-617.

perfect complementary products that compose interconnect systems, and switches sales are worth much more than NICs sales. Despite this, Mellanox has never sought to tie, bundle, or degrade interoperability between its NICs and rivals' switches.[145]

As regards effects

(169) According to the Notifying Party, even on the basis of narrow market segments where the Parties' current positions would be stronger, and assuming the Merged Entity was seeking to leverage its position in the putative markets for high-performance interconnect fabrics and Ethernet NICs of at least 25 Gb/s, the Transaction will not lead to anti-competitive foreclosure. This is for the following reasons.

(170) First, the Notifying Party argues that no foreclosure strategy based on Mellanox's position in these market segments would cover a fraction of the market output large enough in the related market for discrete datacentre GPUs to cause anti-competitive foreclosure. According to the Notifying Party, even assuming that both Mellanox Ethernet NICs of at least 25 Gb/s and Mellanox's InfiniBand fabric were must-have products, a large majority of GPU sales correspond to GPU-accelerated servers connected either with competing high-performance fabrics, competing Ethernet NICs of at least 25 Gb/s or Ethernet NICs below 25 Gb/s (from either Mellanox or competitors) for which it is clear that Mellanox does not hold a strong position. Only a small minority of GPU sales correspond to GPU-accelerated servers connected with Mellanox InfiniBand or Mellanox Ethernet NICs of at least 25 Gb/s. As a result, the Notifying Party argues that there would still be considerable GPU sales opportunities available to other GPU vendors. Mellanox could thus not foreclose a "*sufficiently large fraction*" of sales of discrete GPUs used in datacentres.[146]

(171) Second, the Notifying Party argues that Mellanox's position within the two putative markets is eroding due to the entry/expansion of competitors (see above). Therefore, not only the ability to leverage Mellanox's position becomes untenable but even assuming that the Merged Entity had such ability, the fraction of discrete datacentre GPU sales that would be affected would be even more limited.

(172) Third, the Notifying Party argues that the discrete datacentre GPU market is growing quickly, meaning that the discrete datacentre GPU sales not connected with Mellanox products will increase in the coming years, leaving even more GPU sales opportunities to competing GPU suppliers.

(173) Fourth, the Notifying Party argues that Intel and AMD are leveraging R&D from gaming to the datacentre. Therefore, Intel and AMD will have additional GPU scale beyond datacentres and beyond any potential reach of Mellanox's products.

(174) Fifth, the Notifying Party argues that Intel's decision not to buy Mellanox and leaving Mellanox to NVIDIA, while sunsetting its Omni-Path fabric and continuing investing in discrete datacentre GPUs is inconsistent with the idea that the Merged Entity could foreclose Intel from launching its discrete datacentre GPU.

---

145 Form CO, paragraphs 618-623.

146 Form CO, paragraphs 630-644 and Notifying Party's response to RFI 10, question 3.

(175) Sixth, the Notifying Party argues that Intel and AMD could team up with Cray's Slingshot to overcome any foreclosure strategy.

(176) Seventh, the Notifying Party explains that Mellanox and NVIDIA depend on the CPU suppliers, in particular, (i) to release roadmap information (which defines the standards that Mellanox and NVIDIA need to meet in their product designs) and (ii) to offer "Early Access Programs," which allow Mellanox and NVIDIA to test and validate their products with next-generation CPUs. According to the Parties, the threat of Intel and AMD withholding information and cutting off access to early releases of CPUs acts as a powerful disciplining constraint against the Parties trying to exploit Mellanox's market position through bundling/tying practices.

(177) Eighth, the Notifying Party argues that OEMs and at least some end-customers have considerable countervailing buyer power and that they could put pressure on the Merged Entity not to bundle/tie their products, either by threatening to remove the Parties from approved vendor lists (in the case of OEMs) or by switching to alternatives or creating in-house variants in particular in the case of large CSPs.

5.2.3.3. Commission's assessment

(1) As regards ability

(a) Assessment of Mellanox's potential market power

(178) **As regards Mellanox's potential market power on a high performance fabric market**, the Commission considers that Mellanox most likely has a sufficient degree of market power to leverage its position with its InfiniBand fabric in order to influence the choice of the GPU supplier. This is because the market investigation showed that, for Mellanox's InfiniBand fabric customers, there are and will continue to be within the next 2-3 years too few relevant alternatives. InfiniBand's features appear to be key to connect their GPU-accelerated servers given their workloads.

(179) In the first place, the market investigation has shown that as of today, end-customers connecting their GPU-accelerated servers with Mellanox InfiniBand fabrics do not have sufficient alternatives. The vast majority of end-customers procuring GPU-accelerated servers connected with Mellanox's InfiniBand fabric over the last two years declared that they did not consider any alternative as credible. These customers explained that for the clusters of GPU-accelerated servers they recently acquired Mellanox's InfiniBand fabric was the only credible choice.[147] Some of them explained that this is because InfiniBand has unique features that cannot be replicated to date, including specific optimisations for MPI[148] and very low latency.[149]

(180) This is confirmed by OEMs and competitors. Major OEMs responding to the Commission's market investigation explain that there are end-customers for which Mellanox's InfiniBand fabric is the only credible choice to connect their GPU-accelerated server clusters. According to them, there are not sufficient credible

147 Replies to Questionnaire Q3 to End Customers, questions 39 and 45.

148 Message Passing Interface (MPI), a message passing standard for parallel computing architectures, which is a fundamental requirement for networking devices for HPC and AI training workloads.

149 Replies to Questionnaire Q3 to End Customers, questions 38, 39 and 45.

alternatives to Mellanox's InfiniBand fabric to connect GPU-accelerated server clusters on the market.[150] The vast majority of competitors also confirm that there are specific end-customers for which Mellanox's InfiniBand fabric is the only credible choice to connect their GPU-accelerated server clusters.[151]

(181) In the second place, even looking forward to the next 2-3 years, the Commission considers that InfiniBand will continue to be key for many customers building GPU-accelerated server clusters and that the alternatives available will either not be sufficiently performant and/or will not be sufficiently broadly available.

(182) First, as regards Intel's Omni-Path fabric, the Commission considers that the first and only generation of Omni-Path (which is limited to 100 Gb/s bandwidth) is unlikely to constitute a credible alternative to InfiniBand going forward in case of bundling/tying strategies.

(183) This is supported by Mellanox's internal documents. In one of its presentations to the Board of Director, Mellanox considered already in October 2018, that InfiniBand had [BUSINESS SECRETS – Information redacted regarding business strategy].[152] [BUSINESS SECRETS – Information redacted regarding business strategy].[153]

(184) Since then Mellanox's InfiniBand HDR 200G has become a commercial success while Intel discontinued the development of Omni-Path. The latest Top500 list of supercomputers of November 2019 shows that 200 Gb/s HDR InfiniBand accelerates 31% of the new 2019 InfiniBand systems on this list.[154] If customers make the choice to go for a 200 Gb/s bandwidth, the Commission considers that Intel's Omni-Path cannot compete against Mellanox's InfiniBand HDR 200Gb/s. Moreover, given the absence of a roadmap for future faster fabrics, the Commission considers that even customers opting for a 100 Gb/s fabric may be even more reluctant than before to consider Omni-Path as a credible alternative. On this point, it should be noted that even for past opportunities, before the announcement by Intel that it would stop developing Omni-Path, the vast majority of customers acquiring GPU-accelerated server clusters connected with Mellanox InfiniBand fabric did not consider Omni-Path to be a suitable alternative.[155]

(185) Overall, given the strong adoption already in 2019 of the new generation of InfiniBand HDR 200Gb/s and the fact that Omni-Path was already considered inferior to the previous generation of Mellanox's InfiniBand fabric, the Commission considers that Omni-Path is unlikely to constitute a credible alternative to InfiniBand going forward in case of bundling/tying strategies.

(186) Second, as regards Cray Slingshot, the Commission considers that even if Slingshot may technically be a credible alternative to Mellanox's InfiniBand, it will most likely not be sufficiently broadly available for Intel and AMD to compete successfully against a bundle NVIDIA GPU – Mellanox InfiniBand fabric, with the

150 Replies to Questionnaire Q2 to OEMs, questions 39 and 42.

151 Replies to Questionnaire Q1 to Competitors, question 34.

152 Mellanox's internal document, […], slide 233.

153 Mellanox's internal document, […], slide 97.

154 https://www hpcwire.com/off-the-wire/mellanox-announces-200g-hdr-infiniband-accelerates-31-of-new-infiniband-systems/.

155 Replies to Questionnaire Q3 to End Customers, question 45.

exception of very specific large (exascale) supercomputer opportunities, as explained above.

(187) As explained above, both Intel and AMD discrete datacentre GPUs have won two recent opportunities, i.e. for the Aurora and Frontier supercomputers, which will sit atop the Top500 list of supercomputers. As the Notifying Party explains, they have done so with Cray's Slingshot fabric (as part of Cray servers systems), not Mellanox's InfiniBand. The Commission considers that these two examples strongly support the fact that Cray's Slingshot, at least from a technical point of view, is emerging as a credible alternative to Mellanox's InfiniBand in the short term – the Frontier and Aurora supercomputers are both expected to be deployed by 2021. These examples may suggest that access to Mellanox's InfiniBand fabric is not indispensable for Intel and AMD to win discrete datacentre GPU opportunities even for the most demanding HPC/AI applications. Both Intel and AMD could in principle team up with Cray (which was recently acquired by HPE) to compete with a bundle NVIDIA GPU – Mellanox InfiniBand fabric for a given opportunity.

(188) However, there remain doubts whether such a counter-strategy could be deployed sufficiently widely to defeat a bundling strategy by the Merged Entity. This is because Cray's Slingshot fabric is currently planned to be available as part of Cray's next generation Shasta supercomputer only, not on the merchant market for other OEMs/ODMs to deploy in other server clusters configurations.[156] The Commission considers that Cray's Shasta supercomputer is a unique platform for exascale supercomputer opportunities, as reflected by the two wins with AMD and Intel GPUs discussed above which both involve exascale supercomputers. In addition, Cray announced a third exascale supercomputer win, El Capitan, expected to come online in 2023. Together these supercomputer projects will be the three first exascale supercomputers built in the United States.[157] However, the Commission considers that Cray's Shasta supercomputer is unlikely to be a good fit for many customers seeking to acquire smaller scale GPU-accelerated servers as part of their datacentres. An Intersect 360 Research paper of October 2019 for instance explains that Cray *"has continued to be relatively weak outside of its powerhouse core segment"* of large supercomputer procurements by the US government. In particular, *"Cray's opportunities are sometimes limited by the company's lack of participation in the entry-level and midrange HPC server classes"*.[158]

(189) This is in line with end-customers feedback to the Commission's market investigation. According to some of them, in order to be suitable as an alternative, HPE would have to develop Slingshot NICs for non-Cray server systems.[159] On this point, the OEMs expressing a view on the question explained that even if HPE

---

156 Slingshot was created to act as a suitable network backbone, one that would offer a host of features to allow the Shasta system to comfortably straddle the supercomputing and datacentre worlds. https://www.cray.com/sites/default/files/Slingshot-The-Interconnect-for-the-Exascale-Era.pdf .

157 https://www.cray.com/blog/cray-announces-third-exascale-supercomputer-win/. https://www.cray.com/customers/argonne-national-laboratory. https://www.cray.com/company/customers/oak-ridge-national-laboratory.

158 Mellanox's internal document, […], page 6.

159 Replies to Questionnaire Q3 to End Customers, question 42.

decided to extend Slingshot to other server systems, it would take more than 2-3 years to make this technology generally available.[160]

(190) Third, as regards Atos-Bull's BXI fabric, Atos is confident that its new (Bull's) BXI fabric, although it is still in development, will be a suitable alternative to Mellanox's InfiniBand by 2020, even in case of low latency requirements.[161] However, so far, Atos has only sold its network interconnects as part of its own servers, and there are no indications that this will change post-Transaction. The Commission considers that Atos' niche presence in the server markets means that Bull's BXI will not be sufficiently broadly available for Intel and AMD to compete successfully against a bundle NVIDIA GPU – Mellanox InfiniBand fabric.[162] Finally, even if Atos decided to extend BXI to other server systems, OEMs consider that it would take more than 2-3 years to make this technology generally available.[163]

(191) Fourth, the vast majority of customers, OEMs and competitors consider that no other proprietary high performance fabric could become a suitable alternative to Mellanox's InfiniBand fabric within the next 2-3 years.[164] The majority of OEMs, end-customers and competitors expressing a position on the question also do not believe that a competing Ethernet fabric will emerge within the next 2-3 years that would be able to compete successfully with InfiniBand.[165] On this point, the Commission notes that contrary to the Notifying Party's view, the results of the market investigation indicate that InfiniBand maintains an advantage in terms of latency over Ethernet.[166] This is supported by the Parties' own data. These data show that Mellanox's most advanced Ethernet NICs (ConnectX-6 EN) have a latency […] compared to Mellanox's most advanced InfiniBand NICs (ConnectX-6 VPI).[167]

(192) Therefore, the Commission considers that Mellanox most likely has a sufficient degree of market power to leverage its position in the market for high performance fabric in order to influence the choice of the GPU supplier. *A fortiori*, this conclusion holds in a possible narrower market for InfiniBand fabric, in which Mellanox holds a market share of 100%. This conclusion also holds if the high-performance fabric market were to be segmented according to the bandwidth. Mellanox currently offers InfiniBand fabrics with three different bandwidth – i.e. 56Gb/s, 100Gb/s and 200Gb/s. The Commission considers that Mellanox most likely has a sufficient degree of market power to leverage its position to influence the choice of the GPU supplier irrespective of the bandwidth considered.

(193) **As regards Mellanox's potential market power on a market for Ethernet NICs of at least 25 Gb/s**, on balance, the Commission considers that Mellanox most likely

---

160 Replies to Questionnaire Q2 to OEMs, question 45.

161 Replies to Questionnaire Q1 to Competitors, question 17.

162 According to Hyperion Research, Atos has a market share of 1.1% in the HPC server market.

163 Replies to Questionnaire Q2 to OEMs, question 45.

164 Replies to Questionnaire Q2 to OEMs, question 45; Questionnaire Q3 to End Customers, question 42; and Questionnaire Q1 to Competitors, question 37.

165 Replies to Questionnaire Q2 to OEMs, question 46; Questionnaire Q3 to End Customers, question 43; and Questionnaire Q1 to Competitors, question 38.

166 Replies to Questionnaire Q2 to OEMs, question 46; Questionnaire Q3 to End Customers, question 43; and Questionnaire Q1 to Competitors, question 38.

167 Notifying Party's response to RFI 4, question 24.

does not have sufficient degree of market power to leverage its position in order to influence the choice of the GPU supplier. This is for the following reasons.

(194) First, overall, OEMs' and end-customers' responses to the Commission's market investigation suggest that Mellanox does not have sufficient degree of market power today to leverage its position in order to influence the choice of the GPU supplier. For instance, the vast majority of end-customers and OEMs expressing a view on the question consider that there are sufficient credible alternatives to Mellanox's Ethernet NICs with a speed of at least 25 Gb/s to connect GPU-accelerated server clusters on the market, including Broadcom, Intel, Marvell, Cisco and Chelsio who all offer Ethernet NICs of at least 25 Gb/s.[168] This does not mean that today all these NICs offer exactly the same performance levels according to all metrics and all environments. But at least, from the point of view of end-customers and OEMs, these alternatives are already sufficiently credible today to prevent the Merged Entity from leveraging its position in Ethernet NICs of at least 25 Gb/s to impose an NVIDIA GPU on them.

(195) Second, going forward, even if in the recent past there may have been a performance gap between Mellanox and some of its competitors as regards their Ethernet NICs of at least 25 Gb/s, the Commission considers that Mellanox's competitors are developing new lines of more performant Ethernet NICs and these new products will compete more strongly with Mellanox's Ethernet NICs than is the case today. As explained by the Notifying Party, Mellanox was the first mover in the market for Ethernet NICs of at least 25 Gb/s in 2012. Since then, competitors, such as Broadcom, Intel, Marvell, Cisco and Chelsio have entered, improved their products and expanded – from 0% market share in 2012 competitors have reached a combined market share of [30-40]% in 2018. For example, Broadcom was the first to sample Ethernet NICs of 200Gb/s in August 2018, only later followed by Mellanox ConnectX-6 Ethernet NICs. According to a Linley Group Report, the two competitors offer similar features for mainstream datacentre applications.[169]

(196) The vast majority of customers and OEMs expressing an opinion on the question, consider that, within the next 2-3 years, to the extent that it was not the case already, competing Ethernet NICs suppliers will be able to offer a competitive Ethernet NIC able to compete successfully with Mellanox's latest generation of high-speed Ethernet NICs.[170]

(197) Of particular interest is Chelsio. According to some customers, Chelsio is a particularly good alternative technically. However, according to one specific

---

168 Replies to Questionnaire Q3 to End Customers, question 37; Questionnaire Q2 to OEMs, question 40.

169 https://www.globenewswire.com/news-release/2018/09/12/1569930/0/en/Broadcom-Samples-Thor-World-s-First-200G-Ethernet-Controller-with-50G-PAM-4-and-PCIe-4-0 html; https://www.linleygroup.com/newsletters/newsletter_detail.php?num=5910&year=2018&tag=3.

170 Replies to Questionnaire Q3 to End Customers, question 44; Questionnaire Q2 to OEMs, question 47. For instance, to the question "Do you consider that within the next 2-3 years, competing Ethernet NICs suppliers (Broadcom, Intel, Marvell, etc.) will be able to offer a competitive Ethernet NIC able to compete successfully with Mellanox's latest generation of high-speed Ethernet NICs (considering also Mellanox's new products release roadmap), an end customer explained "Broadcom, Intel, and Marvell will be able to compete successfully and reach similar levels of performance due to functional similarities in their products", another lists "Intel, Broadcom, Netronome", while a third one says "Other Ethernet NICs suppliers such as Intel and Broadcom already offer similar products".

customer, Chelsio lacks the support from big OEMs.[171] Assuming the Merged Entity were to engage in a bundling or tying strategy, the Commission considers that major OEMs would most likely increase their support for Chelsio, who currently supplies Ethernet NICs of up to 100Gb/s, in order to defeat the bundle and allow Intel and AMD to team up with a good alternative to Mellanox's Ethernet NICs of at least 25 Gb/s.

(198) The Commission also considers that Broadcom and Intel will most likely close the gap to a sufficient extent with Mellanox's Ethernet NICs to be considered a credible alternative in case of tying/bundling strategies. In a presentation to the Board of Directors, Mellanox explains that [BUSINESS SECRETS – Information redacted regarding business strategy]. In particular, a Tolly Report commissioned by Mellanox dated September 2019, shows that Mellanox ConnectX-5 Ethernet NIC delivers better performance than Broadcom NetXtreme E NIC.[172] However, in the same presentation, Mellanox explains that [BUSINESS SECRETS – Information redacted regarding business strategy].[173]

(199) Third, following the State of Play meeting of December 5, 2019, the Parties produced an analysis of Mellanox profit margins (in dollar) for its Ethernet NICs of at least 25 Gb/s, distinguishing between previous generations of NICs and ConnectX-5. This analysis shows a [BUSINESS SECRETS – Information redacted regarding profit margins] in profit margins (in dollar) both for previous generations of NICs [BUSINESS SECRETS – Information redacted regarding profit margins] and for ConnectX-5 [BUSINESS SECRETS – Information redacted regarding profit margins].[174] This is consistent with the increased competitive pressure exerted by competing Ethernet NICs suppliers like Broadcom, Chelsio, Intel, Marvell and Cisco, who only entered the market for Ethernet NICs of at least 25 Gb/s at a later stage and who are still expanding.

(200) Fourth, in further support of the fact that Mellanox most likely does not have sufficient market power with its Ethernet NICs of at least 25 Gb/s to leverage into any other markets, the Commission notes that Mellanox's strong market position in Ethernet NICs of at least 25 Gb/s has apparently not placed it in a position to leverage this position into the adjacent hypothetical market for 25 Gb/s+ Ethernet switches and meaningfully grow its sales of Ethernet switches of at least 25 Gb/s. According to the Notifying Party, Mellanox's share of sales for Ethernet switches of at least 25 Gb/s was around [0-5]% in 2018.[175] This is despite the fact that, as

---

171 Replies to Questionnaire Q3 to End Customers, question 44.

172 According to this study, Mellanox ConnectX-5 Ethernet NIC delivers up to twice the throughput of the Broadcom NetXtreme E adapter for various environments and workloads common in cloud, enterprise, and flash-storage deployments. In addition, the testing showed that Mellanox can handle more connections, prevents packet loss, and consumes fewer host CPU cycles per packet. See Mellanox's press release available at https://ir mellanox.com/news-releases/news-release-details/mellanox-releases-independent-report-demonstrating-connectx; and Tolly Enterprises, Mellanox ConnectX®-25GbE Ethernet Adapter – Adapter performance v. Broadcom NetXtreme® E, p. 1 (Sep. 2019), available at https://www mellanox.com/reports/tolly/.

173 Mellanox's internal document, […], slide 49.

174 [BUSINESS SECRETS – Information redacted regarding profit margins]. However, as shipments of ConnectX-5 only started late October 2016, the Commission considers that profit margins generated in 2016 might not be representative. The Commission therefore only considered profit values from 2017 onwards. See Notifying Party's Supplementary Submission of 9 December 2019, paragraph 33.

175 Form CO, paragraph 520.

explained by the Notifying Party, Ethernet NICs and switches are much closer complementary products than Ethernet NICs and discrete datacentre GPUs. NICs and switches are components of the same interconnect systems. When datacentre end-customers procure Ethernet NICs, directly or via OEMs/ODMs, to connect their servers, they also need to buy Ethernet switches. Moreover, according to the Notifying Party, in an average integrated Mellanox Ethernet network interconnect system of at least 25 Gb/s, the gross profit (in dollar) made by Mellanox from the sales of Ethernet switches is […] the gross profit (in dollar) made by Mellanox from the sales of Ethernet NICs.[176]

(b) Assessment of the Merged Entity's ability to engage in the various tying/bundling practices

(201) **As regards the Merged Entity's technical ability to selectively degrade interoperability of competing GPUs with Mellanox's InfiniBand fabric and/or Ethernet NICs of at least 25 Gb/s**, the Commission considers that the Merged Entity would not have such an ability, for the following reasons.

(202) First, Mellanox's network adapters (whether InfiniBand or Ethernet) and NVIDIA's GPUs currently use the open standard PCIe to communicate with other components in the datacentre. PCIe is an industry standard used to connect the CPU host to the peripherals within a server – for example, CPU-to-network cards, CPU-to-memory, CPU-to-hard drive, CPU-to-accelerator.[177] PCIe is an open standard solution,[178] available to everyone on FRAND terms, and it is the *de facto* standard for interconnecting systems within a server.

(203) As a matter of principle, every major datacentre component, including NICs and GPUs, must interoperate with the CPU, which is the central component always present in any kind of servers.[179] Intel and AMD, the two main CPU suppliers, support PCIe. Therefore, it is currently a requirement for both Mellanox's network interconnects and NVIDIA's GPUs to support PCIe in order to communicate with the CPUs. It follows that NVIDIA and Mellanox both use the published PCIe standard to design their products.[180]

(204) In addition, the Parties explained that, even though Mellanox's "PeerDirect" protocol enables direct data transfers between Mellanox NICs and for example, GPUs without having to go through the CPU, this direct communication still takes place via the PCIe bus. In this respect, the Parties explained that any PCIe-enabled device that implement the publicly available PCIe peer-to-peer standard can interoperate with Mellanox's NICs.[181]

---

176 Notifying Party's Supplementary Submission of 9 December 2019, paragraph 34; Notifying Party's response to RFI 13, question 2; and Notifying Party's response to RFI 16, question 4.

177 Commission's decision of 14 October 2015 in case M.7688 – *Intel/Altera*, paragraph 125.

178 Developed though cooperative standard-setting under the auspices of the electronics industry consortium PCI-SIG (https://pcisig.com/).

179 The Parties confirmed this in the Notifying Party's Supplementary Submission of 9 December 2019, paragraphs 58-64.

180 Form CO, paragraph 541.

181 Notifying Party's Supplementary Submission of 9 December 2019, paragraph 45.

(205) Second, Mellanox seems to be generally committed to open source software as a means to ensure the broadest interoperability possible for its products. In particular, Mellanox contests the claim by certain competitors of NVIDIA[182] that some of the features it developed contain proprietary elements. In this respect, Mellanox, states that its "PeerDirect" protocol is "*standards based and open source*".[183] While Mellanox has worked with both NVIDIA and AMD in the past in order to allow them to implement "PeerDirect" for their respective GPU products, as also explained in the above paragraph, Mellanox considers that this PCI-e-based feature is now open source and available to any GPU supplier.[184] Moreover, Mellanox explains that, given that the software is open-source, any exchange of information between a GPU provider and Mellanox regarding its implementation might be convenient, but is not necessary.[185]

(206) Mellanox also provided information with respect to a number of specific application programming interfaces ("APIs") in order to support its claim that these APIs are open-source.[186]

(207) Third, there are a number of technical limits to the Merged Entity's ability to selectively degrade interoperability with Mellanox's network interconnects when they are combined with third parties GPUs. In the first place, in a classic CPU-centric system, the network interconnect communicates with the CPU via PCIe, and the CPU passes on any commands to the GPU. Hence, the network interconnect is "blind" to what processor is accelerating any given task and it is technically not possible to degrade interoperability only for selected accelerators. In the second place, given that Mellanox exclusively relies on open source software, any accelerator-specific code, which would degrade the listed features when running with competing GPUs, would be detected and rejected by Linux, Microsoft or VMWare kernels.[187]

(208) Fourth, contrary to the claim of competitors[188], the Merged Entity would likely not have the ability to replace the PCIe standard with a proprietary network interface. This is because, as apparent from NVIDIA's internal documents,[189] [BUSINESS SECRETS – Information redacted regarding business plans][190] [BUSINESS SECRETS – Information redacted regarding business plans].[191] In addition, the Notifying Party notes that developing a new interface would be successful only if it

---

182 Intel's response to RFI 1, page 2; AMD's response to RFI 1, question 4 a).

183 Notifying Party's Supplementary Submission of 9 December 2019, paragraph 45.

184 Notifying Party's response to RFI 14, paragraph 50.

185 Notifying Party's response to RFI 14, paragraph 39.

186 Including the Mellanox Messaging Accelerator ("MXM"), which Mellanox contributed to the open-source Unified Communication X ("UCX") framework, as well as Mellanox's OFED distribution ("MLNX_OFED").

187 Notifying Party's Supplementary Submission of 9 December 2019, paragraphs 58-60.

188 AMD's response to RFI 1, questions 2.a) and 2.b).

189 Notifying Party's Supplementary Submission of 9 December 2019, annex 4.

190 Notifying Party's Supplementary Submission of 9 December 2019, paragraph 53.

191 Notifying Party's response to RFI 14, paragraphs 53-54. In particular, NVLink is a memory fabric interconnect, allowing connected GPUs to share memory over a simple and specific protocol. It is purpose-designed to connect identical GPUs to each other and in near proximity. In contrast to this, InfiniBand and Ethernet adapters are IO peripherals and connect to a general-purpose IO bus, such as PCIe. They have rich networking protocols necessary to connect systems that have completely different capabilities throughout the datacentre, and therefore require a versatile and complex bus like PCIe.

is widely available and can interoperate with all different components of the datacentre.[192] In particular, adapters need to interoperate with Intel and AMD's CPUs which rely on PCIe.

(209) The Commission notes that, in the current setting, the existence of PCIe, an industry standard used by all major suppliers of datacentre components including Mellanox and NVIDIA and Mellanox's strong commitment to open source, leave little room for any selective degradation of interoperability. While a change of these current circumstances cannot be excluded, the Commission considers that it does not need to take a position on the question whether the Merged Entity will have the technical ability to engage in selective interoperability degradation because, in any case, the Commission considers that the Merged Entity will not have any incentive to do so as explained below at paragraphs 221 to 223.

(210) **As regards the Merged Entity's ability to engage in contractual tying whether with Mellanox's InfiniBand fabric and/or Ethernet NICs of at least 25 Gb/s**, the Commission considers that it is unclear whether the Merged Entity would have such ability. As explained by the Notifying Party, most of Mellanox's network interconnects and NVIDIA's GPUs are sold through OEMs and ODMs. The question is whether these intermediates would have the ability and incentive to defeat a contractual tying strategy. The Notifying Party claims that OEMs/ODMs could defeat such practice by buying the "bundle" from the Parties and then selling the components separately to their customers. However, OEMs explained that to date, they only rarely store components of the Parties for future projects. Instead the vast majority of OEMs' and ODMs' purchases of NVIDIA GPUs and Mellanox network interconnect appear to take place in the context of specific projects tendered out by end customers.[193] This is among others because GPUs and some network interconnect components are expensive; technology evolves quickly; and there are many (hundreds if not more than a thousand) different individual network interconnect products depending on the exact requirements of customers.[194] Based on the evidence available at the date of drafting this decision, the Commission considers that there is not enough evidence to conclude whether OEMs and ODMs would likely expose themselves to the risk of storing vast amount of bundled products to be able to resell them unbundled, in order to defeat a contractual tying strategy. The Commission therefore concludes that it cannot exclude that the Merged Entity would have the ability to engage in contractual tying.

(211) **As regards the Merged Entity's ability to engage in mixed bundling whether with Mellanox's InfiniBand fabric and/or Ethernet NICs of at least 25 Gb/s**, all competitors and OEMs expressing a view on the question consider that the Merged Entity would have such ability.[195] However, the Commission considers that by engaging in such conduct, the Merged Entity would not have the ability to leverage Mellanox's position to significantly steer end-customers' choice towards NVIDIA's GPUs. This is because the Commission considers that AMD and Intel would still have the ability to compete on price especially considering (1) the low relative price

192 Notifying Party's response to RFI 14, paragraph 57.

193 Replies to Questionnaire Q2 to OEMs, questions 11, 19-20, 24 and 28.

194 Mellanox's Opportunity data shows that Mellanox has sold more than 1000 different products, based on the "product name" covering more than 100 "product types".

195 Replies to Questionnaire Q1 to Competitors, question 41; Questionnaire Q2 to OEMs, questions 49.

of network interconnect products as compared to GPUs, implying that a discount on the interconnect products would not affect the price of the bundle significantly,[196] and (2) the possibility for AMD and Intel to provide a discount on a bundle including their CPUs and their GPUs to counter a reduced price by the Merged Entity on a bundle GPU-interconnect when compared to the sum of standalone prices.[197] Therefore, in the rest of the decision, the Commission does not cover the mixed bundling scenario any longer.

(c) Assessment of the Merged Entity's ability to foreclose

(212) **As regards the Merged Entity's ability to leverage its position in the high performance fabric market to foreclose its discrete datacentre GPU rivals**, the Commission considers that the Merged Entity would not have such an ability. This is because, as will be explained further below in paragraph 214, the vast majority of discrete datacentre GPU customers do not equip their GPU-accelerated server clusters with a Mellanox InfiniBand fabric. As a result, irrespective of the tying/bundling strategy considered above, there would remain sufficient demand for GPUs corresponding to servers systems that do not use Mellanox's InfiniBand fabric. As a result, the Commission considers that, even if the Merged Entity could fully leverage its alleged market power with its InfiniBand fabric (i.e. assuming that no customer would switch to a competing interconnect to be able to use a competing GPU), the Merged Entity would still not be able to foreclose enough market output to hinder Intel's effective entry and AMD's expansion into the datacentre GPU market.

(213) This result relies on an extensive analysis carried out by the Parties of their respective transaction data. The Parties carried out this analysis in two major steps. First, they matched their respective transaction databases to determine which NVIDIA discrete datacentre GPU customers also purchased a Mellanox InfiniBand fabric in 2018. Second, the Parties estimated for each NVIDIA customer who also purchased Mellanox's InfiniBand fabrics, the share of NVIDIA GPU sales that were actually used in servers connected with the InfiniBand fabrics purchased.

(214) The Parties found that only [0-30]% of NVIDIA's 2018 datacentre GPU revenue was made from sales into servers that also use Mellanox InfiniBand.[198] Considering that NVIDIA with its [90-100]% share of the overall discrete datacentre GPU market is representative of the overall market, this would mean that the Merged Entity could at most foreclose [0-30]% of the discrete datacentre GPU market (in terms of value) by leveraging its position with InfiniBand. This result relies however on a number of assumptions, some of which are conservative, while some others may not be. The

196 The 2018 GPU dollar profit per server was around […] the 2018 average InfiniBand fabric dollar profit per server. See Notifying Party's response to RFI 15.Considering a broadly similar gross margin ratio for NVIDIA GPUs as for InfiniBand fabric, this means that the ratio of sales value per server would be around […] as well. When considering only Ethernet NICs, the ratio would be closer to […] (see below).

197 Together, AMD and Intel account for 97% of the server CPU market (94% Intel and 3% AMD). See Notifying Party's response to RFI 19.

198 The rest of NVIDIA's GPU sales correspond to GPU-accelerated server clusters using other types of network interconnects, i.e. competing high performance fabrics (Intel's Omni-Path, Cray's Aries or Gemini (and in the future Slingshot), Bull's BXI, Fujitsu's Tofu, *etc*.), Mellanox's Ethernet NICs of at least 25 Gb/s, competing Ethernet NICs of at least 25 Gb/s, and Ethernet NICs below 25 Gb/s (where Mellanox has a market share of [0-5]%).

Commission replicated the analysis of the Parties, keeping the conservative assumptions but relaxing to the maximum the non-conservative ones. By doing so, the Commission found that at most [0-30]% of NVIDIA's 2018 datacentre GPU revenue could have been made from sales into servers that also used Mellanox InfiniBand. This would leave at least [70-100]% of the discrete datacentre GPU market (in terms of value) unaffected by any bundling/tying practices considered above.

(215) This result is particularly conservative considering that it ignores that fact that even if the Merged Entity were to engage in contractual tying or degradation of interoperability involving Mellanox's InfiniBand fabric, AMD and Intel could most likely in the future capture some of the GPU sales corresponding to servers for which customers have a preference for Mellanox InfiniBand fabric. This would be the case in particular when the customer procures or has the possibility to procure the server cluster from HPE (which recently acquired Cray) or Atos, acting as OEM. As explained above, both these OEMs offer proprietary high-performance fabrics which technically may be credible alternatives to InfiniBand going forward. If the Merged Entity were to try and leverage the customer's preference for InfiniBand to force the purchase of NVIDIA GPUs against the end-customers' will, the Commission considers that these OEMs would likely replace Mellanox's InfiniBand by their own proprietary high-performance fabric in the server system offered to the end-customer.

(216) **As regards the Merged Entity's ability to leverage its position both in the high performance fabric market and in the market for Ethernet NICs of at least 25 Gb/s to foreclose its discrete datacentre GPU rivals,** even assuming that Mellanox also holds significant market power in Ethernet NICs of at least 25 Gb/s, which the Commission considers is not the case, and that the Merged Entity would have the ability to fully leverage this alleged market power, the Parties found that only [0-30]% of NVIDIA's 2018 datacentre GPU revenue was made from sales into servers that also use Mellanox Ethernet NICs of at least 25 Gb/s. After lifting to the maximum the potentially non-conservative assumptions, the Commission found that the overlap in terms of GPU sales would be at most [0-30]%. Therefore, in the worst-case scenario, assuming that the Merged Entity would have the ability to fully leverage its position both with InfiniBand and with its Ethernet NICs of at least 25 GB/s, and only using conservative assumptions, the Merged Entity could at most foreclose [0-40]% of the discrete datacentre GPU market. Again, this result is particularly conservative, as it assumes that the Merged Entity would have the ability to foreclose GPU rivals from selling GPUs in all systems involving Mellanox InfiniBand and/or Ethernet NICs of at least 25 Gb/s. This is unlikely to be the case for the reasons explained above.

(217) In 2018, according to the Notifying Party, the putative discrete datacentre GPU market was worth EUR [1 000-3 000] million.[199] Even if the Merged Entity could foreclose every opportunity linked to its Ethernet NICs of at least 25 Gb/s and InfiniBand fabric, considering the most conservative assumptions, that would leave at the very least EUR [1 000-3 000] million of open opportunities.[200] As the Commission considers that Mellanox most likely does not have sufficient degree of

199 Notifying Party's response to RFI 10, question 3.

200 [60-100]% of EUR [1 000-3 000] million.

market power in the market for Ethernet NICs of at least 25 GB/s to leverage its position in order to influence the choice of the GPU supplier, it is more reasonable to consider that the Merged Entity would at most engage in a tying practice involving its InfiniBand fabric. Even if the Merged Entity could foreclose every opportunity linked to its InfiniBand fabric, considering the most conservative assumptions, that would leave at the very least EUR [1 000-3 000] million of open opportunities.[201]

(218) These figures are based on the size of the GPU market in 2018. In addition, the datacentre market is growing. AMD predicted in May 2019 in an investor presentation that as early as 2021, the value of opportunities for GPUs in datacentres would be USD 12 000 million (EUR 11 000 million).[202] If this forecast is correct, this would mean that even if the Merged Entity could foreclose every opportunity linked to its InfiniBand fabric,[203] this would leave at least EUR [5 000-10 000] million of open opportunities addressable by AMD and Intel (EUR [5 000-10 000] million in case of full leverage of Mellanox's position also in the market for Ethernet NICs of at least 25 Gb/s). This is several times higher than the sales value of NVIDIA in 2018.[204] This means that in all likelihood, AMD and Intel will be able to reach their minimum viable scale, even if foreclosed from the segment of the market linked to Mellanox's products.

(219) Therefore, the Commission concludes that the Merged Entity would not have the ability to foreclose enough market output to hinder Intel's effective long-term entry and AMD's expansion into the discrete datacentre GPU market. This conclusion also holds under the assumption of a possible narrower market for InfiniBand fabric or if the high-performance fabric market were to be further segmented by bandwidth ranges. Irrespective of the exact market delineation from which the Merged Entity would attempt to leverage its position, Intel and AMD would still be able to address most of the growing market for discrete datacentre GPUs, i.e. the part of the market which is unrelated to Mellanox's network interconnect products.

(2) As regards incentive

(220) The incentive to degrade interoperability and/or engage in contractual tying depends on the degree to which such a strategy would be profitable.[205] When considering whether or not to engage in such practices, the Merged Entity faces a trade-off between sales of network interconnect products foregone and GPU sales retained that could otherwise have gone to Intel and/or AMD.

(221) As regards **degradation of interoperability**, as explained above, the only way the Merged Entity would potentially be able to degrade interoperability would be to develop a proprietary interface for its NICs to which competing GPU suppliers would not have access. However, NICs need to communicate not only with GPUs but also with CPUs and other datacentre components. As long as the majority of devices communicates via PCIe, the Merged Entity would have no incentive to

201 [70-100]% of EUR [1 000-3 000] million.

202 AMD Investor Presentation, May 2019, Slide 15, available at: http://ir.amd.com/static-files/9c985e84-bbb6-4e23-99bd-dcbb21f18592.

203 Assuming the proportion of opportunities linked to InfiniBand remains constant.

204 In FY2019 (ending 27 January 2019), NVIDIA turnover in Datacentre GPUs was around EUR [1 000-3 000] million, see Form CO, paragraph 27.

205 Non-Horizontal Merger Guidelines, paragraph 105.

depart from that standard. If the Merged Entity were to develop new Ethernet NICs of at least 25 Gb/s or InfiniBand NICs[206] with a proprietary interface, the Commission considers that these would not be accepted by OEMs and end-customers in datacentres because these new proprietary NICs and high performance fabrics would not be interoperable with the vast majority of CPUs[207] and other components.[208]

(222) Additionally, the Merged Entity will rely, as the Parties currently do, on OEMs as its largest and most important go-to-market channel. OEMs demand interoperability throughout the datacentre and could delist the Merged Entity's products if the Merged Entity would degrade their interoperability.[209] This would put at risk all of the Merged Entity's network interconnect sales without protecting NVIDIA's GPU sales.

(223) Overall, therefore, the Commission considers that the Merged Entity would not have the incentive to engage in degradation of interoperability of its network interconnects with competing GPUs. This conclusion also holds under the assumption of a possible narrower market for InfiniBand fabric or if the high-performance fabric market were to be further segmented by bandwidth ranges.

(224) As regards **contractual tying**, the Commission considers that the Merged Entity would have the ability to focus its contractual tying on situations where end-customers require GPU-acceleration and not try to force GPUs on customers who do not need GPUs. This is because most OEMs and competitors not only indicated that GPU suppliers typically offer their products in the context of specific projects, but also, in relation to these projects, these GPU suppliers generally have good information on end-customers' specifications.[210] By targeting the contractual tying to GPU customers only, the Merged Entity could limit the foregone sales on the network interconnect side while increasing the probability that the tie will trigger additional GPU sales (as compared to the counterfactual).

---

206 As part of an InfiniBand fabric.

207 The Commission is aware that IBM and NVIDIA developed NVLINK as a high-speed connection between the IBM POWER 8 CPU and the NVIDIA Tesla P100 GPU (See https://www.ibm.com/blogs/systems/ibm-nvidia-present-nvlink-server-youve-waiting/). However, Intel and AMD still rely exclusively on and continue to develop new generations of PCIe. Together they represent 97% of server CPU sales. As such, it is crucial for Mellanox's NICs to support that standard.

208 According to the Notifying Party, all NICs must support the CPU, the root of every computer system. A NIC is an I/O peripheral with an I/O register address space. PCIe is designed to communicate with such an I/O peripheral. Every NIC is not only capable of connecting to PCIe—it must connect to the CPU using PCIe. In addition, Mellanox NICs also support the PCIe standard's peer-to-peer functionality through open-source software. Any PCIe device that supports the peer-to-peer functionality can communicate with NICs, including CPUs, memory devices, GPUs, FPGAs, NNPs, IPUs, ASICs, and others. Moreover, the Notifying Party explains that the use of NVLINK between Tesla P100 GPU and IBM Power 8 CPU is not a NIC interface, it is a memory interface to directly connect GPU memory to CPU memory. All IBM Power 8 servers include a PCIe NIC for network communications, illustrating that NVLINK is not a suitable interconnect for a NIC. In fact, NVLINK is not suitable to replace any PCIe connection. Even with NVLINK, IBM Power systems must also include PCIe connections between the CPU and the GPU. See Notifying Party's response to RFI 18.

209 Form CO, paragraphs 544-545.

210 Replies to Questionnaire Q1 to Competitors, questions 18-19; Questionnaire Q2 to OEMs, questions 19-20, 24 and 28.

(225) The Notifying Party indicated in the Form CO that the average GPU profit made by NVIDIA per server is around […] times higher than the average NIC profit made by Mellanox per server.[211] This means that, the Merged Entity may have an incentive to engage in contractual tying leveraging its position in Ethernet NICs of at least 25 Gb/s if it manages to retain at least 1 GPU opportunity for every […] lost Ethernet NIC opportunity. However, as explained above, the Commission considers that Mellanox most likely does not have sufficient degree of market power in Ethernet NICs of at least 25 Gb/s to leverage its position in order to influence the choice of the GPU supplier. This means that engaging in contractual tying would potentially expose the Merged Entity to massive switching from customers. It is therefore unclear whether the Merged Entity would have the incentive to engage in contractual tying involving its Ethernet NICs of at least 25 Gb/s.

(226) As regards InfiniBand, the Notifying Party submits that the average 2018 GPU dollar profit per server was around […] times the 2018 average InfiniBand fabric dollar profit per server.[212] The Commission considers that, in all likelihood, the gained sales (or retained sales) on the GPU side of imposing a contractual tie involving its InfiniBand fabric would therefore more than compensate any lost sales on the network interconnect side.

(227) However, the Commission considers that the Merged Entity would not have the incentive to engage in contractual tying.

(228) In fact, as regards both **degradation of interoperability and contractual tying,** the Commission considers that the Merged Entity's incentive to engage in such practices should also be assessed taking into account Intel's and AMD's potential counter-strategies. On this point, the Notifying Party explained that Mellanox and NVIDIA absolutely depend on access to Intel's and AMD's CPU roadmaps, product prototypes, and other early-release information in order for NVIDIA and Mellanox to align their roadmaps and to be able to offer solutions that support Intel's and AMD's CPUs at the time those products launch.[213] This is key, given that CPUs are at the heart of every system with which the Parties have to interoperate, and Intel and AMD together account for the vast majority of CPU sales.[214] According to the Notifying Party, if Intel and AMD were to withhold that information, it would have an immediate and durable impact on Mellanox's and NVIDIA's ability to bring their products to market and to compete in a timely way. This would have a disciplining effect on the Merged Entity, eliminating all incentives it may have to engage in contractual tying or degradation of interoperability.[215]

(229) Intel, however, claimed that NVIDIA is not dependent on Intel roadmap information or product samples and has not been an active participant in Intel's roadmap programs because of its reluctance to share roadmaps with Intel. According to Intel, the PCIe bus facilitates communications between CPUs and other devices and facilitates interoperability with minimal information sharing. As regards the

---

211 Form CO, paragraph 601.

212 Notifying Party's response to RFI 15.

213 Notifying Party's Supplementary Submission of 9 December 2019, paragraph 47.

214 Together, AMD and Intel account for 97% of the server CPU market (94% Intel and 3% AMD). See Notifying Party's response to RFI 19.

215 Notifying Party's Supplementary Submission of 9 December 2019, paragraph 47.

Mellanox's side of the Merged Entity, Intel explains that it would not have the incentive to withhold information considering the strong competition from AMD and ARM and the strong position of Mellanox on the network interconnect markets.[216]

(230) As regards the dependence of NVIDIA on Intel's roadmap, the Commission considers that NVIDIA needs access to information from Intel to ensure interoperability with Intel's newest CPUs. This is because, as explained by the Notifying Party, while it is true that NVIDIA does not need the entirety of Intel's CPU roadmap, NVIDIA needs Intel's PCIe roadmap to ensure that NVIDIA's products, which are peripherals to Intel CPUs, are matched in capabilities and will interoperate with Intel's newest CPUs. Given the lead-time to design GPUs, NVIDIA and Mellanox need that information *at least* [BUSINESS SECRETS – Information redacted regarding business strategy] in advance. In practice, since PCIe came into existence, NVIDIA has received and relied on Intel's PCIe roadmap, which reveals the PCIe generational level and production timing for Intel's forthcoming CPUs. According to the Notifying Party, this gives Intel massive power over NVIDIA.[217] The same is true for AMD's roadmaps: NVIDIA needs access to AMD's PCIe roadmap to be able to interoperate with AMD CPUs.[218]

(231) As regards product samples, the Commission considers that NVIDIA relies on Intel to get advance access to Intel's CPU product samples in order to perform testing and validation of NVIDIA's GPUs with Intel's CPUs. This is because, as explained by the Notifying Party [BUSINESS SECRETS – Information redacted regarding NVIDIA's advance access to Intel's CPUs].[219] The same is true for AMD's product samples: NVIDIA needs access to AMD's CPU product samples in order to perform testing and validation of its GPUs with AMD's CPUs. [220]

(232) According to the Notifying Party, access to information and product samples from Intel and AMD is important because the CPU is the root complex of PCIe; it controls every computer. Intel and AMD alone decide when to implement new levels and generations of PCIe. Intel and AMD decide what PCIe timing to use. Intel's and AMD's interpretations of PCIe for their respective CPUs control and override any contrary view. Intel's and AMD's physical implementations of PCIe control every Intel and AMD CPU-based computer.[221]

(233) According to the Notifying Party, the CPU is the one device in the ecosystem that indisputably rules all. A peripheral that does not operate seamlessly with Intel's and/or AMD's CPU—or fails to keep up with Intel's and/or AMD's PCIe CPU roadmap—is simply not marketable to the Intel and/or AMD ecosystems.[222]

(234) Considering that Intel has a market share of 94% in the server CPU market[223] and that CPUs equip all datacentres, the Commission considers that interoperability with Intel's CPUs is crucial for both NVIDIA's GPUs and Mellanox's network

216 Intel's response to RFI 1, question 8.
217 Notifying Party's response to RFI 14, question 10.
218 Notifying Party's response to RFI 21.
219 Notifying Party's response to RFI 14, question 10.
220 Notifying Party's response to RFI 21.
221 Notifying Party's response to RFI 14, question 10; Notifying Party's response to RFI 21.
222 Notifying Party's response to RFI 14, question 10; Notifying Party's response to RFI 21.
223 Notifying Party's response to RFI 19.

interconnect products. Furthermore, despite the 3% market share of AMD on the server CPU market, the Commission considers that interoperability with AMD's CPUs is crucial for both NVIDIA's GPUs and Mellanox's network interconnect products. This is for the following reasons.

(235) First, if NVIDIA's GPUs and Mellanox's network interconnect products would not interoperate with AMD's CPUs, this would make NVIDIA's datacentre business almost entirely dependent on Intel's PCIe roadmap and Intel CPU price, availability, features, quality,

(236) Second, the Commission considers that today, ensuring that NVIDIA GPUs and Mellanox' network interconnect products interoperate with AMD CPUs is more important than ever because AMD's Second-Generation EPYC CPUs (codenamed Rome) have rapidly emerged as the preferred CPU for high performance computing. AMD's Rome CPUs are winning high-profile business throughout the datacenter ecosystem, including the Frontier supercomputer at Oak Ridge National Labs (US)[224], the Archer supercomputer at the University of Edinburgh's supercomputing center (EU)[225], and many more.[226]

(237) Rome's rapid success reflects several concrete advantages, which the Commission considers make it very important for NVIDIA and Mellanox to interoperate with AMD.

(238) First, AMD leap-frogged Intel's CPU PCIe roadmap with Rome, the first commercial X86 CPU that supports PCIe Gen4.[227] As explained above, NVIDIA must have access to AMD's Gen4 CPUs, if it wants to develop and test NVIDIA GPUs with PCIe Gen4 and other PCIe Gen4 compatible peripherals including NICs, storage controllers, FPGAs, and others. In other words, NVIDIA cannot market PCIe Gen4 products without help and access to AMD's Rome CPUs. The Commission would consider it irrational if NVIDIA were to take the risk of falling behind the competition when it comes to supporting PCIe Gen4.

(239) Second, the Commission considers that AMD EPYC's impact on GPU opportunities vastly exceeds and will exceed AMD's past market share. As explained by the Notifying Party, AMD EPYC is the first PCIe Gen4 CPU on the market, and each AMD EPYC CPU supports 128 Gen4 PCIe lanes.[228] In contrast, Intel's top end datacentre processor, the 2nd Gen Intel Xeon Scalable Platinum series, supports only 48 Gen3 PCIe lanes. Even setting aside AMD's lead in PCIe Generations, each AMD CPU provides 2.66 (i.e. 128/48) times as many PCIe lanes as Intel's CPUs.

---

224 See https://www hpcwire.com/2019/05/07/cray-amd-exascale-frontier-at-oak-ridge/.

225 See https://www.nextplatform.com/2019/10/18/amd-cpus-will-power-uks-next-generation-archer2-supercomputer/.

226 See https://www.amd.com/en/press-releases/2019-11-18-amd-delivers-best-class-performance-supercomputers-to-hpc-the-cloud-sc19.

227 Johan De Gelas, AMD Rome Second Generation EPYC Review: 2x 64-core Benchmarked, AnandTech (Aug. 7, 2019) ("As the first commerical x86 server CPU supporting PCIe 4.0, the I/O capabilities of second generation EPYC servers are top of the class."); AMD EPYC 7002 Series Processors ("AMD EPYC™ is the first and only current x86-architecture server processor supporting PCIe 4.06."), available at https://www.amd.com/en/processors/epyc-7002-series.

228 *See* AMD EPYC 7002 Series Processors ("All-in feature set"), available at https://www.amd.com/en/processors/epyc-7002-series.

AMD's advantage in PCIe lanes will translate directly into more opportunities for PCIe peripherals, including GPUs.[229]

(240) Third, as explained by the Notifying Party, in addition to its advantages in PCIe, AMD Rome has other competitive advantages over Intel Xeon, including (1) twice the memory capacity, (2) greater memory frequency, (3) lower power, (4) lower licensing costs for per socket software licenses, and (5) lower list price. As a result, AMD Rome is rapidly gaining server CPU market share and is far more important than past market share implies.

(241) The importance for NVIDIA to guarantee interoperability with AMD is also reflected in NVIDIA's choice to [BUSINESS SECRETS – Information redacted regarding new product development].[230]

(242) Overall, the Commission considers therefore that the Merged Entity's requirement to get access to Intel's and AMD's PCIe roadmaps and advance product samples to interoperate with their respective CPUs will most likely eliminate all incentives by the Merged Entity to degrade interoperability or engage in contractual tying at the expense of Intel and AMD. In particular, the Commission considers that Intel and AMD are more pivotal than Mellanox and NVIDIA in datacentres. The vast majority of the server CPU sales are indeed not dependent on access to Mellanox's InfiniBand fabric (the vast majority of servers are not connected with InfiniBand). On the other hand, the vast majority NVDIA's GPUs and Mellanox's network interconnect products sales depend on being interoperable with Intel's and AMD's CPUs. This conclusion also holds under the assumption of a possible narrower market for InfiniBand fabric or if the high-performance fabric market were to be further segmented by bandwidth ranges.

(3) As regards effects

(243) The Commission considers that, even if (1) the Merged Entity had significant market power in a high performance fabric market, (2) it had the ability to fully leverage such market power into the discrete datacentre GPU market through contractual tying or degradation of interoperability practices, and (3) it had the incentive to engage in such practices, the reduction in GPU sales prospects faced by AMD and Intel would be so limited that it would not lead to a reduction in Intel's and AMD's ability or incentive to compete. As explained above (in the sub-section on ability), at least [70-100]% of the discrete datacentre GPU market would be unaffected by any bundling/tying practices involving InfiniBand.

(244) Even if the Merged Entity could also fully leverage its position (which the Commission considers unlikely) in the market for Ethernet NICs of at least 25 Gb/s, at least [60-100]% of the discrete datacentre GPU market would be unaffected by any bundling/tying practices involving Mellanox's InfiniBand and Ethernet NICs of at least 25 Gb/s.

(245) Given these proportions, the current size of the GPU market (EUR [1 000-3 000] million) and its expected growth (see above in the sub-section on ability), the

229 Notifying Party's response to RFI 21.

230 Notifying Party's response to RFI 21.

Commission concludes that AMD and Intel would still have access to sufficient discrete datacentre GPU sales, and that their ability and incentive to compete would therefore remain unaffected. The Commission considers that in all likelihood, AMD and Intel will be able to reach their minimum viable scale, even if foreclosed from the segment of the market linked to Mellanox's products (whether just InfiniBand or also Ethernet NICs of at least 25 Gb/s).

(246) Therefore, the Commission concludes that Intel's effective entry and AMD's expansion into the discrete datacentre GPU market will in all likelihood not be hindered and that therefore competition is very unlikely to deteriorate.[231] The Commission therefore considers that the Transaction is very unlikely to harm consumers, even if the Merged Entity were to engage in any of the tying practices considered.[232] This conclusion also holds under the assumption of a possible narrower market for InfiniBand fabric or if the high-performance fabric market were to be further segmented by bandwidth ranges. Irrespective of the exact market delineation[233] from which the Merged Entity would attempt to leverage its position, Intel and AMD would still be able to address most of the growing market for discrete datacentre GPUs, i.e. the part of the market which is unrelated to Mellanox's network interconnect products.

(247) The conclusion that there will not be any appreciable negative impact on customers is furthermore confirmed by the results of the market investigation. The vast majority of end customers and OEMs expressing a view on the question consider that the impact of the Transaction on their company would be positive or neutral and that the impact of the Transaction on the intensity of competition in the discrete datacentre GPU market would be positive or neutral. Moreover, the vast majority of

---

231 As explained above Intel and AMD are already competing as of today in the market for discrete datacentre GPUs. AMD has been present in the market for a number of years and is expanding with new products. Intel has announced the launch of its Xe GPU for datacentre by 2021. But it is already starting to compete today by participating in tenders, as evidenced by its recent win in the tender organised by the U.S. Department of Energy for the upcoming Aurora datacentre at Argonne National Laboratory.

232 The Non-Horizontal Merger Guidelines (paragraph 113) state that "[i]*t is only when a sufficiently large fraction of market output is affected by foreclosure resulting from the merger that the merger may significantly impede effective competition. If there remain effective single-product players in either market, competition is unlikely to deteriorate following a conglomerate merger*". For completeness, the Commission however explains below why it considers that even customers in the affected segment (the part of the GPU sales that corresponds to GPU-accelerated servers connected with Mellanox's interconnect products for which Mellanox has market power) would not be harmed. This is because, the Commission considers that the Merged Entity will not have the ability to raise GPU prices or force unwanted GPUs on customers solely because it will now also control Mellanox interconnect products. Indeed, it will generally be more profitable for Mellanox' owners (both pre- and post-merger) to exploit whatever market power Mellanox may possess over its customers by directly raising interconnect product prices to the highest level the market will bear, rather than imposing unwanted products on customers. A strategy of forcing unwanted products on customers would instead reduce customers' willingness to pay for GPU/interconnect combinations, and would thus diminish the rents the Merged Entity can hope to extract. In the absence of a realistic prospect of hampering rivals' ability and incentive to compete, it is therefore unlikely that the Merged Entity could profitably impose competitive damage on customers. This is confirmed by the results of the market investigation, as the vast majority of end-customers declaring that they recently procured a cluster of GPU-accelerated servers for which Mellanox's InfiniBand fabric was the only credible choice as a connection between the servers, were not concerned that the Transaction may impact them negatively or that the Transaction would decrease the intensity of competition in the discrete datacentre GPU market. Questionnaire Q3 to End-Customers, questions 39 and 56-57

233 High performance fabric, InfiniBand fabric, high performance fabric considering any bandwidth range possible, or even InfiniBand fabric considering any bandwidth range possible.

end-customers declaring that they recently procured a cluster of GPU-accelerated servers for which Mellanox's InfiniBand fabric was the only credible choice as a connection between the servers, were not concerned that the Transaction may impact them negatively or that the Transaction would decrease the intensity of competition in the discrete datacentre GPU market.[234]

5.2.4. Leveraging the position of NVIDIA in the market for discrete datacentre GPUs into the various network interconnect markets in which Mellanox is active

5.2.4.1. Potential concern

(248) The Commission has assessed a potential competition concern whereby the Merged Entity would leverage NVIDIA's strong position in the plausible market for discrete datacentre GPU into any network interconnect markets where Mellanox is active, and whether this would have a significant detrimental effect on competition in these network interconnect markets, thus causing harm to datacentre end customers. The Commission has carried out this assessment overall, rather than for each potential network interconnect product market that could be the target of the Merged Entity's potential leveraging strategy.

(249) The Commission has assessed in particular the ability and the incentive of the Merged Entity to engage in one or both of the following tying/bundling practices:

- Differentiating the degree of technical compatibility and therefore overall performance of its joint solution compared to mix-and-match solutions involving only one of its products ("technical tying"); and/or

- Incentivising the joint purchase of its own products by offering higher prices for mix-and-match solutions involving only one of its products as compared to the bundle ("mixed bundling").

5.2.4.2. Notifying Party's view

(250) The Notifying Party submits that the Merged Entity will not have the ability and incentive to leverage NVIDIA's potentially strong position in the plausible market for discrete datacentre GPU into any network interconnect markets where Mellanox is active. In any event, the Notifying Party submits that any putative leveraging could not lead to anticompetitive foreclosure of Mellanox's rivals. The reasons are the following.

As regards ability

(251) First, the Notifying Party argues that the Merged Entity will not have the ability to leverage NVIDIA's market position post-Transaction, because NVIDIA lacks market power in the supply of discrete datacentre GPUs.

---

[234] Replies to Questionnaire Q2 to OEMs, questions 65-66 ; Questionnaire Q3 to End-Customers, questions 39 and 56-57. For example, an end-customer "*sees the proposed acquisition of Mellanox by Nvidia as beneficial in terms of technology competition in the supply chain* […]*; Nvidia with Mellanox significantly broadens an important ecosystem*".

(252) In the first place, the Notifying party argues that NVIDIA faces strong competition from suppliers of other types of processing solutions, including CPUs, CPU-based accelerators such as Intel's Xeon Phi, CPUs that integrate acceleration capabilities such as Intel's Xeon Scalable CPUs, ASICs, FPGAs, as well as in-house options developed by CSPs such as those developed by Google (TPU) and Amazon (Inferentia) for their respective cloud services.[235]

(253) In the second place, even on the narrow market of discrete datacentre GPUs identified in Section 4, the Notifying Party argues that NVIDIA will face intense and increasing dynamic competition. The Notifying Party argues that past market shares are not and cannot be a good proxy for measuring current and future market power, in particular given the rise of AMD and the entry of Intel. According to the Notifying Party, together, AMD's and Intel's recent launches have created increasing and growing competitive pressure on NVIDIA that is not captured by NVIDIA's market shares in a hypothetical GPU-only market. This pressure is exacerbated by AMD's and Intel's ability to supply and market both CPUs and GPUs. [236]

(254) In the third place, the Notifying Party argues that NVIDIA's software stack does not in any way protect it from these new entrants (as AMD's market share growth proves). According to the Notifying Party, NVIDIA's API, called CUDA, does not give NVIDIA market power and does not raise barriers to switching for customer and therefore barriers to entry for AMD and Intel.[237]

(255) In the fourth place, the Notifying Party argues that NVIDIA's market behaviour will be constrained by the practice of the industry. Large, sophisticated datacentre customers seek bids from large, sophisticated OEMs/ODMs, both of which, in addition to organising and disciplining the actual bidding and proposal process, exert considerable countervailing buyer power.[238]

(256) Second, as regards technical tying, the Notifying Party claims that there are no practicable means through which the Parties could degrade interoperability.

(257) Third, as explained above (see Section 5.2.3.2.), the Notifying Party argues that the procurement structure of this industry precludes the ability to leverage.

As regards incentives

(258) According to the Notifying Party, the Merged Entity would also not have the incentive to degrade the interoperability of NVIDIA's discrete datacentre GPUs with third parties' network interconnect products or to raise NVIDIA's GPUs' relative price when combined with third party network interconnect products.

235 Form CO, paragraphs 486-494.
236 Form CO, paragraphs 495-502.
237 Form CO, paragraphs 503-504.
238 Form CO, paragraphs 504-505.

(259) First, as explained above, the Notifying Party argues in general terms that the Parties have strong commercial incentives to continue interoperating with other datacentre component suppliers, including their competitors.[239]

(260) Second, the Notifying Party argues that the cost of foreclosing suppliers of network interconnects in terms of lost GPU sales would outweigh any benefit from increased interconnect sales. According to the Notifying Party, this is for two main reasons. In the first place, following a bundling strategy, a majority customers would opt to turn to a competing accelerator processing solution in order to keep their preferred network interconnect solution. In the second place, the reduction in GPU profits from losing a GPU customer far exceeds the increase in network interconnect profits from gaining an interconnect customer because the GPU profit per server is significantly higher than the network interconnect profit per server. As noted above, the average 2018 GPU dollar profit per server was around […] the average NIC profit made by Mellanox per server[240] and around […] InfiniBand fabric dollar profit per server.[241] This is reinforced by the fact that OEMs can exert countervailing buyer power by threatening to remove NVIDIA and Mellanox from their approved vendor lists.[242]

(261) Third, the Notifying Party argues that any leveraging strategy would lead to retaliation from Intel and AMD, which control the ecosystems attached to their CPUs. Moreover, the Notifying Party argues that any leveraging strategy (assuming the Merged Entity's dominance on the market for discrete datacentre GPUs) would expose the Parties to antitrust scrutiny and possible follow-on litigation. This risk acts as a significant deterrent to carry out any putative anti-competitive foreclosure strategy.[243]

As regards effects

(262) According to the Notifying Party, even on the basis of narrow market segments, the Transaction will not lead to anti-competitive foreclosure. First, the Notifying Party argues that the vast majority of datacentre interconnect sales are made to customers that do not buy NVIDIA GPUs.[244] Second, the Notifying Party argues that NVIDIA's position within the putative market for discrete datacentre GPUs is eroding due to the entry/expansion of Intel and AMD. Therefore, not only the ability to leverage NVIDIA's position becomes untenable but even assuming that the Merged Entity had such ability, the fraction of network interconnect products sales that would be affected would be even more limited.[245]

5.2.4.3. Commission's assessment

(263) The Commission investigated whether the Merged Entity would have the ability and the incentive to leverage its market power in the plausible market for discrete datacentre GPUs into any markets for Mellanox network interconnects by engaging

[239] Form CO, paragraphs 590-598.
[240] Notifying Party's response to RFI 13, question 2.
[241] Notifying Party's response to RFI 15.
[242] Form CO, paragraphs 599-606.
[243] Form CO, paragraphs 611-617.
[244] Form CO, paragraphs 668-672.
[245] Form CO, paragraphs 673-675

in technical tying and/or in mixed bundling practices with a view to foreclose its competitors. In summary, the Commission considers that the Merged Entity would neither have the ability nor the incentive to foreclose Mellanox's competitors, and that in any event such a strategy would not have anti-competitive effects.

As regards ability

(264) The Commission considers that the Merged Entity would not have sufficient **market power** to leverage its position in the market for discrete datacentre GPUs into the markets for network interconnects. While the Merged Entity has a high market share of [90-100]% in the market for discrete datacentre GPUs today, this position is being challenged by AMD's expansion and Intel's announced entry into the market. As shown in Table 1 above, AMD has been able to almost triple its market share in the market for discrete datacentre GPUs from [0-5]% to [5-10]% between 2016 and 2018. Moreover, as noted above in Section 5.2.2., both AMD and Intel have managed to win tenders for the most performant supercomputers with their GPU offerings. This shows that the GPUs developed by both companies are already considered suitable alternatives to the ones provided by NVIDIA. Additionally, both Intel and AMD also supply server CPUs[246] and are able to offer both CPUs and GPUs to datacentre customers. This could further increase the attractiveness of their GPU offerings.[247]

(265) Moreover, the majority of end customers that expressed a view consider that Intel's and AMD's discrete datacentre GPUs will probably be credible alternatives in the near future (i.e. in the next 2-3 years).[248] The Commission therefore considers that, as long as the Merged Entity will not foreclose Intel and AMD by leveraging Mellanox's position in network interconnects into the GPU market,[249] their entry/expansion in the market for discrete datacentre GPUs will likely significantly reduce NVIDIA's market power.

(266) The Commission also does not consider that NVIDIA's CUDA software will allow the Merged Entity to maintain market power in the market for discrete datacentre GPUs. CUDA is a common API for all of NVIDA's chips. It allows engineers to port applications to run, in part, on NVIDIA's GPUs. NVIDIA develops CUDA libraries specifically for accelerating HPC and AI workloads. The software is not sold by NVIDIA, but available for download.[250] The majority of respondents to the market investigation considered that currently, the difficulty of migrating software written and run in the CUDA environment to other platforms constitutes a significant barrier to switching from NVIDIA's discrete datacentre GPUs to competing discrete datacentre GPUs.[251] However, competitors have developed tools to translate or replace CUDA: AMD has for instance developed the Heterogeneous compute

---

246 In 2018, Intel had a market share of 94.1% and AMD a market share of 3% in server CPUs, see Notifying Party's response to RFI 19.

247 For instance, AMD and Microsoft have announced an AMD datacentre win for Microsoft's Azure cloud, in which Microsoft will deploy AMD CPUs and GPUs together: https://azure.microsoft.com/en-us/blog/announcing-new-amd-epyc-based-azure-virtual-machines/.

248 Replies to Questionnaire Q3 to End Customers, questions 50 to 50.2.

249 See the Commission's assessment in Section 5.2.3.3.

250 Notifying Party's response to RFI 4, question 14.

251 Replies to Questionnaire Q1 to Competitors, questions 42 and 42.1; Questionnaire Q2 to OEMs, questions 53 and 53.1; Questionnaire Q3 to End Customers, questions 48 and 48.1.

Interface for Portability (**"HIP"**), an API that allows developers to create portable applications that can run on AMD's accelerators as well as CUDA devices.[252] Intel on the other hand is developing One API, a software that supports direct programming and API programming and will deliver a unified language and libraries that offer full native code performance across a range of hardware, including CPUs, GPUs, FPGAs and AI accelerators.[253]

(267) Therefore, the Commission considers that the ability of the Merged Entity to leverage its market power in the market for discrete datacentre GPUs would likely be limited.

(268) Moreover, the Commission considers that even if the Merged Entity had market power in the market for discrete datacentre GPUs, it would not have the **ability to engage in technical tying and/or in mixed bundling practices** with a view to foreclose its competitors. The respondents to the market investigation stated that the Merged Entity would have the technical ability to degrade the interoperability of rival network interconnects with NVIDIA's discrete datacentre GPUs[254], as well as the ability to offer commercial bundles including both types of components at a lower price than the sum of the individual components.[255] However, for the reasons explained above in Section 5.2.3.3., mainly that GPUs and network interconnects need to use open source interfaces to communicate with each other as well as other datacentre components, and that the Parties mainly sell their products through OEMs, the Commission considers that the Merged Entity would not have the ability to engage in such practices in order to foreclose its competitors in the markets for network interconnects.

As regards incentive

(269) The Commission considers that the Merged Entity would not have the incentive to engage in foreclosure practices towards Mellanox's competitors. First, doing so would lead to a loss of GPU sales that would not be compensated by sales of network interconnects. This is because, as stated above in Section 5.2.3.3., most of the profit per server results from GPU sales and not from network interconnect sales. The average 2018 GPU dollar profit per server was around […] than the average NIC profit made by Mellanox per server and around […] InfiniBand fabric dollar profit per server.

(270) Second, Mellanox's most significant competitors for high-performance fabric both belong to OEMs, who would have the ability and incentive to counter any foreclosure attempt designed to benefit Mellanox. In particular, Cray Slingshot was recently acquired by HPE, who is also one of NVIDIA's largest OEM customers. HPE could delist NVIDIA's GPUs from its approved vendor list, which would

---

252 Damon McDougall et al., "Introduction to AMD GPUs programming with HIP", 6 July 2019, available at: https://www.exascaleproject.org/wp-content/uploads/2017/05/ORNL_HIP_webinar_20190606_final.pdf.

253 Intel, "Intel's 'One API' Project Delivers Unified Programming Model Across Diverse Architectures", 19 June 2019, available at: https://newsroom.intel.com/news/intels-one-api-project-delivers-unified-programming-model-across-diverse-architectures/#gs.6o2cny.

254 Replies to Questionnaire Q1 to Competitors, questions 47 to 47.4.2; Questionnaire Q2 to OEMs, questions 58 to 58.4.1.

255 Replies to Questionnaire Q1 to Competitors, questions 48 to 48.4.2; Questionnaire Q2 to OEMs, questions 59 to 59.4.2.

represent a loss of at least USD […] based on HPE's total purchases of NVIDIA's Tesla GPUs in 2018. If HPE also delisted Mellanox as part of this retaliation, the Merged Entity would lose an additional USD […] of sales from interconnects to this customer.[256] Accordingly, if it were to engage in such a foreclosure strategy against Cray Slingshot, the Merged Entity would stand to lose a significant amount of sales for both GPUs and network interconnects through HPE.

As regards effects

(271) As for the effects, the Commission notes that the majority of sales in the markets for datacentre network interconnects are made to datacentres that do not buy NVIDIA's GPUs. Based on the Top500 list of June 2019, only 42% ([0-40]% based on the opportunity data of the Parties) of the datacentres connected with Mellanox InfiniBand also use GPUs.[257] Similarly, only 25% ([0-30]% based on the opportunity data of the Parties) of the datacentres connected with Mellanox Ethernet NICs of at least 25 Gb/s also use GPUs.[258] This was confirmed by the results of the market investigation. The OEMs and end customers that provided data on the total number of NICs that their company installed as part of server clusters in 2018-2019 confirmed that the vast majority of sales in the markets for datacentre network interconnects are made to datacentres that do not buy NVIDIA's GPUs and often are not accelerated at all.[259] Moreover, given the entry/expansion of Intel and AMD in the market for discrete datacentres GPUs, the share of network interconnect product sales that would be affected will likely be even more limited.

(272) Moreover, the Commission notes that all major competing suppliers of network interconnects (except Intel) considered the Transaction to have a neutral impact on their business[260] as well as on the markets for network interconnects[261]. A competitor for network interconnects submitted that *"we hope that it will enable us to sell our […] products"*[262] and that *"as of today, the products sold by NVIDIA and MELLANOX are the best on the market irrespective of this transaction"*[263]. In addition, the large majority of end customers that expressed a view, including a number of universities and research centres, considered the Transaction to have a positive or neutral impact on their business[264] as well as on the markets for network interconnects[265]. A European research centre submitted that "*this will boost the network industry and perhaps will boost the EU investments inside sovereign solution as network like processors or memories are key technologies to master.*"[266]

(273) The Commission therefore considers that any foreclosure strategy leveraging NVIDIA's position in the market for discrete datacentre GPUs into the network

256 Form CO, paragraph 606.

257 The other datacentres of the Top500 list connected with Mellanox InfiniBand are either accelerated with co-processors (*i.e.*, CPUs added in addition to the baseline CPUs), or not accelerated at all.

258 Form CO, paragraphs 669-672.

259 Replies to Questionnaire Q2 to OEMs, question 60; Questionnaire Q3 to End Customers, question 53.

260 Replies to Questionnaire Q1 to Competitors, questions 51 and 51.1.

261 Replies to Questionnaire Q1 to Competitors, questions 52.2, 52.4 and 53.

262 Replies to Questionnaire Q1 to Competitors, question 51.1.

263 Replies to Questionnaire Q1 to Competitors, question 52.4.

264 Replies to Questionnaire Q3 to End customers, questions 56 and 56.1.

265 Replies to Questionnaire Q3 to End customers, questions 57.2, 57.4 and 58.

266 Replies to Questionnaire Q3 to End customers, question 57.4.

interconnect markets would likely have no significant effect on competition in these markets.

5.2.5. Possible leakage of commercially sensitive information

5.2.5.1. Potential concern

(274) Market participants explained that suppliers of datacentre components may enter into partnerships with suppliers of different components in order to ensure the best interoperability between their respective products and the best level of joint performance for their customers. These arrangements are generally not formalised and likely vary depending on the companies and/or products involved. The companies involved may exchange different types of information such as their respective roadmaps and product pans, they may set up joint validation-processes, they may collaborate when issues arise, *etc*.[267]

(275) Currently, NVIDIA and Mellanox receive information from suppliers of datacentre components in the context of such cooperation arrangements enabling them to ensure interoperability between their respective products. Of most relevant to the present case, Mellanox receives information from GPU suppliers (AMD and Intel) in order to enable interoperability of their GPUs with Mellanox's network interconnects, and NVIDIA may receive information from network interconnect suppliers in order to enable interoperability of their network interconnects with NVIDIA's GPUs.

(276) In this context, the Commission has assessed a potential concern that the Merged Entity would receive commercially sensitive information from competing GPU and/or network interconnect providers that the Merged Entity could use to favour its own position on the GPU and/or network interconnects relevant markets.

5.2.5.2. Notifying Party's view

(277) The Notifying Party submits that the Merged Entity will not be able to use information received by Mellanox from NVIDIA's competitors to favour its own position in the GPU market to the detriment of AMD and Intel, or to use information received by NVIDIA from Mellanox's competitors to favour its own position in the network interconnect markets to the detriment of rival network interconnect suppliers.

(278) First, according to the Notifying Party, while Mellanox receives substantial confidential information from AMD and Intel relating to their CPUs, it does not receive similar confidential/commercially sensitive information from relating to their GPUs.[268] The Notifying Party acknowledges that Mellanox may receive some information from AMD and Intel about their GPUs (*e.g.*, general roadmap or timeline information) but that information is not "competitively sensitive".[269]

---

267 See, for example, […].

268 Notifying Party's response to RFI 12, question 1, paragraphs 1-2.

269 Notifying Party's response to RFI 12, question 1, paragraph 4.

(279) In relation to network interconnect information, the Notifying party indicates that "*NVIDIA does not obtain any confidential information from Mellanox or any other network interconnect supplier for the development of its GPUs*."[270]

(280) Second, the Notifying Party argues that, in the event that confidential information were shared, GPU suppliers can protect their information via non-disclosure agreements ("NDAs") restricting Mellanox's employees from disclosing confidential information to NVIDIA's employees working on GPUs.[271] Moreover, the Notifying Party argues that, given their level of sophistication, AMD and Intel would be able to negotiate broader/more restrictive NDAs with the Merged Entity if necessary.[272]

(281) Finally, the Notifying Party considers that rival GPU suppliers have different means to retaliate if the Merged Entity were to share their commercially sensitive information internally. First, they could use the legal grounds provided for by the NDAs concluded with the Parties (*e.g.*, fast-track dispute mechanisms). Second, since information about CPUs is critical for Mellanox, AMD and Intel could stop (or threaten to stop) providing such information. Finally, the Notifying Party notes that, if Mellanox were to share information with the NVIDIA side of the Merged Entity in violation of an NDA, this would greatly affect the Merged Entity's credibility in the market.[273]

5.2.5.3. Commission's assessment

(1) Possible leakage of GPU competitors' commercially sensitive information

(282) The results of the market investigation revealed a concern relating to the possible leakage within the Merged Entity of commercially sensitive information shared by competing GPU vendors. In particular, post-Transaction, the Merged Entity's networking division could make this information accessible to the Merged Entity's GPU division, which could potentially misuse it to favour the Merged Entity's own GPUs to the detriment of competing GPU suppliers.[274]

(283) Therefore, the Commission has considered whether the Merged Entity could leak and misuse the commercially sensitive information that the Mellanox side of the business may receive from GPU suppliers to favour the NVIDIA side of the business on the market for discrete GPUs for datacentre.

(284) **As regards ability to foreclose**, first, the confidentiality concern, as expressed by some third parties, is based on the premise that GPU suppliers willing to cooperate with Mellanox to ensure the interoperability of their GPUs with Mellanox's network interconnects must share information on their GPUs that is confidential or commercially sensitive. They fear that this information could be used to their disadvantage by the GPU side of the Merged Entity. Therefore, the Commission first

[270] Notifying Party's response to RFI 13, question 1.

[271] Notifying Party's response to RFI 12, question 1, paragraphs 14-15.

[272] Notifying Party's response to RFI 12, question 1, paragraph 20.

[273] Notifying Party's response to RFI 12, question 1, paragraphs 22-29.

[274] Replies to Questionnaire Q1 to Competitors, questions 49.3 and 49.3.1. See also agreed minutes of the conference call of 9 August 2019 with AMD, paragraph 35, and Intel's submission of 10 September 2019 entitled "Intel response to case team's query regarding foreclosure mechanism resulting from NVIDIA's acquisition of Mellanox", p. 5.

assessed whether, based on the information currently exchanged, the Merged Entity could have access to competitors' commercially sensitive information received in the context of cooperation arrangements with the latter.

(285) A majority of the respondents to the market investigation that expressed a view suggested that, when cooperating with Mellanox to ensure interoperability of their products with the latter's network interconnects, they provide commercially sensitive information (*e.g.*, product roadmaps) about their company's acceleration products.[275]

(286) More specifically, AMD explained that, under its current partnership with Mellanox, both companies […].[276] […].[277] Similarly, Intel is concerned that "[t]*o the extent that the Merged Entity's networking division cooperates with Intel and other suppliers of acceleration products, it will gain access to confidential information regarding their products' designs and roadmaps*."[278]

(287) However, the Notifying Party refutes these claims and argues that Mellanox only receives limited information from GPU suppliers,[279] but not "*any meaningful competitively sensitive information about AMD and Intel's product development plans, roadmaps, or product specifications for their GPUs, nor does it expect to going forward*".[280]

(288) Contrary to competitors' claims, the Notifying Party explains that "*as a rule* [AMD and Intel] *do not share GPU information with Mellanox that might be competitively sensitive as it relates to NVIDIA.*" The information that they may share (such as general roadmap or timeline information) is "*equivalent to what they already disclose at industry conferences and* product announcements".[281] This information is "*not competitively sensitive, and access to* [it] *would not confer an advantage to NVIDIA.*"[282]

(289) For example, Mellanox does not receive detailed, pre-release product information on GPUs, and AMD and Intel do not have "early access programmes" (through which they would provide Mellanox with prototype products or detailed product specifications) or "roadmap alignment" discussions with Mellanox about their GPUs.[283] Similarly, as regards information shared during the product testing and validation phase, it concerns "*products that are already available, not confidential, pre-release products from those suppliers*"[284], which, according to the Notifying Party, does not constitute "*competitively sensitive information regarding future or*

---

275 Replies to Questionnaire Q1 to Competitors, questions 49 and 49.1.

276 See agreed minutes of the conference call of 9 August 2019 with AMD, paragraph 27.

277 […].

278 Intel's submission of 10 September 2019 entitled "Intel response to case team's query regarding foreclosure mechanism resulting from NVIDIA's acquisition of Mellanox", p. 5.

279 Notifying Party's response to RFI 12, question 1, paragraph 13.

280 Notifying Party's response to RFI 12, question 1, paragraph 12. Similarly, the Notifying Party argues that "*there is no history of AMD or Intel sharing sensitive information with Mellanox about their GPU products that actually would be competitively meaningful to NVIDIA*", see Notifying Party's response to RFI 12, question 1, paragraph 9.

281 Notifying Party's response to RFI 12, question 1, paragraph 4.

282 Notifying Party's response to RFI 12, question 1, paragraph 1.

283 Notifying Party's response to RFI 12, question 1, paragraph 4.

284 Notifying Party's response to RFI 12, question 1, paragraph 8.

current GPUs."[285] In fact, "*AMD or Intel could do the testing and validation themselves, without sharing any information with Mellanox at all*".[286]

(290) The Parties have illustrated their arguments by analysing historic information shared between them relating to GPUs. This analysis revealed that [BUSINESS SECRETS – Information redacted regarding information that is not shared between NVIDIA and Mellanox]. Any discussions regarding design and interoperability issues with respect to particular customers between the Parties, [BUSINESS SECRETS – Information redacted regarding information that is not shared between NVIDIA and Mellanox].[287]

(291) Furthermore, the Notifying Party explains that, in any event, AMD and Intel do not need to provide such information to Mellanox.[288] While access to this type of information is "critical" for Mellanox and other datacentre component providers in relation to CPUs (given their central role in datacentre equipment), "*Mellanox does not need similar advanced information from GPU suppliers*" because its "*own product development plans and roadmaps do not depend on technical details from the GPU suppliers.*"[289] This is so because Mellanox's network interconnects operate with GPUs via PCIe such that "*there is no real need for AMD and Intel to share pre-release product information about their GPUs* [...] *for Mellanox to interoperate with those products*".[290]

(292) Based on the Notifying Party's submissions and the market investigation, the Commission has established that, on balance, it is unlikely that the Merged Entity will be in a position to obtain commercially sensitive information from rival GPU suppliers that it could use to favour its own GPUs' position to the detriment of AMD and Intel. In any event, as explained below, even if the merged company had access to such strategically important confidential information, AMD and Intel would have ways to protect their information.

(293) Second, based on the Notifying Party's submissions and the market investigation, the Commission has established that, in the industry concerned in the present case, companies cooperating with each other generally rely on certain safeguards to protect their confidential and commercially sensitive information, notably though NDAs.

(294) However, according to a few market participants, the currently applicable NDAs would likely not be sufficient to prevent Mellanox's business employees from sharing their commercially sensitive information with NVIDIA's business post-Transaction.[291] Intel submits that "*Intel has standard NDAs in place with Mellanox that place little controls on how the information could be shared with Nvidia*."[292]

---

285 Notifying Party's response to RFI 12, question 1, paragraph 1.
286 Notifying Party's response to RFI 12, question 1, paragraph 8.
287 Notifying Party's response to RFI 12, question 1, paragraphs 6-7.
288 Notifying Party's response to RFI 12, question 1, paragraph 13.
289 Notifying Party's response to RFI 12, question 1, paragraph 5.
290 Notifying Party's response to RFI 12, question 1, paragraph 5.
291 Replies to Questionnaire Q1 to Competitors, question 49.2.
292 Intel's reply to Questionnaire Q1 to Competitors, question 49.2.1.

Similarly, AMD considers that "*AMD's non-disclosure agreements with Mellanox* […]."[293]

(295) The Notifying Party contests these findings and submits that GPU suppliers can protect their information through industry standard NDAs. According to the Notifying Party, the NDAs currently in place between Mellanox and AMD on the one hand, and Mellanox and Intel on the other, "*at a minimum* […] *would prevent Mellanox from misusing any information that these suppliers provide*".[294] Therefore, even if GPU suppliers were to provide commercially sensitive information to the Merged Entity, this information "*could never be provided to NVIDIA's GPU engineers as it would be covered by industry-standard NDAs*."[295]

(296) One large OEM confirmed the view that the safeguards currently in place are sufficient, as many players in this industry are "*vertically integrated or conglomerated in ways that require proper protections and firewalls on information use*."[296]

(297) The Notifying Party explains that, under the NDAs entered into by Mellanox with various datacentre component suppliers, "[…]" "*Mellanox takes these obligations seriously. Disclosure of confidential information would not only result in Mellanox breaching its legal obligation, but it would have serious repercussions on Mellanox's business relationship with its partners.*"[297]

(298) In particular, the Notifying Party argues that, currently, NVIDIA and Mellanox rely on such safeguards to prevent Intel or AMD to use the information the former provide to the latter in order to ensure interoperability between NVIDIA's GPUs and Mellanox's network interconnects on the one hand and AMD's and Intel's CPUs on the other. As they offer products competing with the Parties' products, AMD and Intel could use the information they received via their CPU business in order to favour their own network interconnect or GPU products (similarly to what AMD and Intel argue the Merged Entity would be able to do post-Transaction).[298]

(299) To illustrate their arguments, the Parties provided to the Commission NDAs entered into by Intel in relation to information received from the Parties, which they argue is relevant because Intel supplies GPUs and network interconnects and, therefore, competes with the Parties. These NDAs are therefore likely to contain provisions ensuring the protection of information provided by the Parties about products with which Intel competes.[299] For example, the Parties refer to the […]. This NDA states that the receiving party must, among other things, "[…]".[300] This suggests that Intel considers this type of NDA to be sufficient to protect this type of confidential information today.[301]

---

293 AMD's reply to Questionnaire Q1 to Competitors, question 49.2.1.
294 Notifying Party's response to RFI 12, question 1, paragraph 1.
295 Notifying Party's response to RFI 12, question 1, paragraph 14.
296 Reply from a large competitor to Questionnaire Q1 to Competitors, question 49.2.1.
297 Notifying Party's response to RFI 12, question 2.
298 Notifying Party's response to RFI 12, question 1, paragraphs 17-18.
299 Notifying Party's response to RFI 12, question 2.
300 Mellanox's internal document, […], clauses 3.1 and 3.2 (Annex RFI 10-3.16).
301 Notifying Party's response to RFI 12, question 2.

(300) In any event, the Commission considers that, if AMD and Intel believe that the currently applicable NDAs do not sufficiently safeguard their information post-Transaction, they could negotiate more restrictive NDAs with the Merged Entity. […].[302]

(301) The Commission considers that AMD's and Intel's negotiation power is credible because they are both large, sophisticated companies that already cooperate – and enter into this type of agreements – with a large number of companies, including NVIDIA and Mellanox, notably in relation to their CPUs.[303] They are used not only to using NDAs as a tool to protect their own information, but also as using NDAs in situation where they could receive information regarding products with which they compete (*e.g.*, if they receive information from Mellanox on network interconnect in the context of interoperability with CPUs, while they also supply network interconnect themselves).

(302) In light of these considerations, the Commission considers that, on balance, the Merged Entity will likely not have the ability to obtain commercially sensitive information from GPU suppliers. Even if the Merged Entity had access to such information, any foreclosure strategy would be limited to its InfiniBand fabric.

(303) **As regards incentives to foreclose**, the Commission has assessed whether the Merged Entity would have an incentive to leak (and misuse) the potentially commercially sensitive information received by the network interconnect side of the business from rival GPU suppliers to favour its own position on the discrete datacentre GPU market.

(304) First, the Commission considers that the Merged Entity's incentive to engage in such practices should be assessed taking into account Intel's and AMD's potential counter-strategies. On this point, the Notifying Party submits that, in the event that Mellanox received commercially sensitive information from rival GPU suppliers and shared such information with NVIDIA, AMD and Intel could withhold CPU information. As explained in more details in paragraphs 228 to 233, the Notifying Party submits this would have an immediate and durable impact on Mellanox's ability to bring its products to market and to compete in a timely way because Mellanox absolutely depends on access to Intel's and AMD's CPU roadmaps, product prototypes, and other early-release information in order for Mellanox to align its roadmaps and to be able to offer solutions that support Intel and AMD's CPUs at the time those products launch.[304] As the Parties put it, "[t]*his dependency gives AMD and Intel a natural disciplining force how Mellanox and NVIDIA handle their confidential information*."[305]

(305) On the contrary, Intel claims that it would not be able to adopt such a counter-strategy, as this "*would be damaging to Intel's own interest*".[306] This is due in particular to the fact that Mellanox's network interconnects are used in numerous CPU-based platforms (for which Intel is a leading provider), and that failing to

302 […].
303 Notifying Party's response to RFI 12, question 1.
304 Notifying Party's Supplementary Submission of 9 December 2019, paragraph 47.
305 Notifying Party's response to RFI 12, question 1, paragraph 25.
306 Intel's response to RFI 1, question 8.

enable Mellanox while rival CPU vendors support its solutions would cause Intel to lose CPU sales.[307]

(306) However, as discussed in Section 5.2.3.3, overall, the Commission considers that the Merged Entity's requirement to get access to PCIe roadmaps and advance product samples of CPU suppliers – in particular Intel and AMD – to interoperate with their CPUs will most likely eliminate all incentives by the Merged Entity to engage into the potential leaking of AMD's and Intel's commercially sensitive information about their GPUs to favour its own position on the discrete datacentre GPU market, at the expense of Intel and AMD.

(307) Second, Intel and AMD could also decide to stop cooperating with – and therefore providing information about their GPUs to – Mellanox. As explained above, there are a sufficient number of alternative suppliers of Ethernet NICS of at least 25 Gb/s with whom GPU suppliers could cooperate instead of Mellanox. As for InfiniBand, both Intel and AMD could in principle team up with Cray, which, with its Slingshot fabric, at least from a technical point of view, is emerging as a credible alternative to Mellanox's InfiniBand in the short term.[308]

(308) Finally, assuming that GPU suppliers do – and will continue to – provide commercially sensitive information to Mellanox, the Parties explained that "[i]*f Mellanox provided this information to NVIDIA in violation of an NDA and this fact became known, this would greatly impact Mellanox's credibility in the market*". This would, for example, discourage datacentre component vendors (such as switch suppliers) from continuing to provide information that they currently provide to Mellanox for interoperability purposes.[309]

(309) Therefore, based on the Notifying Party's submissions, the Commission considers that any leak of confidential or commercially sensitive information from Mellanox to NVIDIA, post-Transaction, would severely damage Mellanox's relationship with other market players and undermine its reputation in the market. As a result, the Merged Entity can be expected to put in place safeguards preventing such leak.

(310) In light of these considerations, the Commission considers that the Merged Entity will likely not have the incentive to leak (and misuse) GPU suppliers' potentially commercially sensitive information received by Mellanox (if any) to favour its own position on the discrete datacentre GPU market.

(311) **As regards effects**, the Commission has assessed whether a strategy whereby the Merged Entity would refuse to enter into enhanced NDAs with AMD and Intel or leak (and misuse) their GPU information would have any effect on competition in the discrete datacentre GPU market.

(312) The Commission considers that, even if the Merged Entity (i) had access to AMD's and Intel's commercially sensitive information through their cooperation with Mellanox, (ii) passed on such information to the NVIDIA side of the business and (iii) misused this information in order to favour its own position in the discrete

[307] Intel's response to RFI 1, question 8.

[308] See paragraph 187.

[309] Notifying Party's response to RFI 12, question 1, paragraph 28.

datacentre GPU market, the reduction in GPU sales prospects faced by AMD and Intel would be so limited that it would not lead to a reduction in Intel's and AMD's ability or incentive to compete. As explained in Section 5.2.3.3, at least [70-100]% of the discrete datacentre GPU market (in terms of value) would be unaffected by any foreclosing practices involving InfiniBand.

(313) Furthermore, even if the Merged Entity could also fully leverage its position (which the Commission considers unlikely) in the market for Ethernet NICs of at least 25 Gb/s, at least [60-100]% of the discrete datacentre GPU market would be unaffected by any such strategy involving Mellanox's InfiniBand and Ethernet NICs of at least 25 Gb/s.

(314) Consequently, the Commission considers that in all likelihood, AMD and Intel will be able, even if foreclosed from the segment of the market linked to Mellanox's products (whether just InfiniBand or also Ethernet NICs of at least 25 Gb/s), to reach their minimum viable scale.

(315) The Commission concludes that Intel's effective long-term entry and AMD's expansion into the discrete datacentre GPU market will in all likelihood not be hindered and that therefore competition is very unlikely to deteriorate. The Commission therefore considers that the Transaction is very unlikely to harm consumers, even if the Merged Entity were to engage in the practice considered.

(316) **Conclusion**. In light of the above considerations, the Commission considers that the Transaction does not raise serious doubts as to its compatibility with the internal market or the functioning of the EEA Agreement with respect to the risks of possible leakage (and misuse) by the Merged Entity of rival GPU suppliers' potentially commercially sensitive information provided to the Merged Entity.

(2) Possible leakage of network interconnect competitors' commercially sensitive information

(317) The Commission has also considered whether the Merged Entity could leak and misuse the information that NVIDIA may receive from network interconnect suppliers to favour Mellanox on the relevant network interconnect markets.

(318) The Commission has established that the Merged Entity would have neither the ability nor the incentive to adopt such a strategy, for the following reasons.

(319) First, the Notifying Party submits that "*NVIDIA does not obtain any confidential information from Mellanox or any other network interconnect supplier for the development of its GPUs*." In fact, "*NVIDIA does not need such information and has always developed its GPUs without it. Therefore, there is no information NVIDIA receives that the Merged Entity could use 'to advantage Mellanox's network interconnect products over competing products*'".[310]

(320) The results of the market investigation did not contradict this fact. While a few competitors indicated that they share commercially sensitive information with

[310] Notifying Party's response to RFI 13, question 1.

NVIDIA in the context of cooperation arrangements, their answers appear to relate to information about processors rather than network interconnects.[311]

(321) Second, the market investigation provided mixed views as to whether there are sufficient safeguards in place to ensure preservation of the confidentiality of the information provided by network interconnect suppliers to NVIDIA.[312] In any event, the considerations mentioned in paragraphs 293 to 301 relating to the use of NDAs between GPU suppliers and Mellanox, also apply in relation to the information shared between network interconnect suppliers and NVIDIA.

(322) Third, in Section 5.2.4.3, the Commission considers that the Merged Entity would not have enough market power to leverage its position in the market for discrete datacentre GPU into the markets for network interconnects.

(323) Similarly, the Commission takes the position that the Merged Entity would not have sufficient market power to engage into a strategy regarding the potential leakage of network interconnect competitors' commercially sensitive information to favour its own network interconnects.[313] As explained in more details in paragraph 264, while the Merged Entity has a high market share of [90-100]% in the market for discrete datacentre GPUs today, AMD's expansion and Intel's announced entry into the market challenge this position. Intel and AMD are already competing as of today in the market for discrete datacentre GPUs and have recently won tenders for the most performant supercomputers with their GPU offerings. Their GPUs already are considered as suitable alternatives to the ones provided by NVIDIA. Therefore, the ability of the Merged Entity to leverage its market power in the market for discrete datacentre GPUs would likely be limited.

(324) Since AMD's and Intel's GPUs constitute alternatives to NVIDIA's GPUs, network interconnect suppliers such as Broadcom, Intel, Marvell and Chelsio will have the possibility to team up with AMD and/or Intel, instead of the Merged Entity. It follows that if these network interconnect suppliers consider that the NDAs in place or those offered by the Merged Entity do not sufficiently safeguard their information, they will likely be in a sufficiently strong position to negotiate more restrictive NDAs, notably by threatening to shift away from NVIDIA's GPUs. Further, the Commission considers that, even if they were not able to negotiate more protecting NDAs, they could decide to stop cooperating with the Merged Entity and focus on cooperating with Intel and AMD as regards the interoperability of their network interconnects with GPUs.

(325) Fourth, the Commission considers that such a strategy would likely not have significant effect on competition in the markets for network interconnects. As explained in paragraph 262, the Notifying Party argues that the vast majority of interconnect sales are made to datacentres that do not buy NVIDIA GPUs.[314] The market investigation confirmed that the vast majority of interconnect sales are made

311 Replies to Questionnaire Q1 to Competitors, questions 50 and 50.1.

312 Replies to Questionnaire Q1 to Competitors, questions 50.2 and 50.2.1.

313 Non-Horizontal Merger Guidelines, paragraph 99.

314 Based on the Top500 list of June 2019, only 42% ([0-40]% based on the opportunity data of the Parties) of the datacentres connected with Mellanox InfiniBand also use GPUs. Similarly, only 25% ([0-30]% based on the opportunity data of the Parties) of the datacentres connected with Mellanox Ethernet NICs of at least 25 Gb/s also use GPUs. See Form CO, paragraphs 666-672.

to datacentres that do not buy NVIDIA's GPUs and often are not accelerated at all.[315] Moreover, given the imminent entry/expansion of Intel and AMD in the market for discrete datacentres GPUs, the share of network interconnect product sales that would be affected will likely be even more limited (see paragraph 271).

(326) Finally, during the market investigation, no competitor raised clear substantiated concerns regarding the potential misuse of their confidential information relating to their network interconnects to impair their efforts to compete with Mellanox. [316]

(327) In light of the above considerations, the Commission considers that the Transaction does not raise serious doubts as to its compatibility with the internal market or the functioning of the EEA Agreement with respect to the risks of misuse by Mellanox of competing network interconnect suppliers' potentially commercially sensitive information provided to the Merged Entity.

5.3. Vertical non-coordinated effects

5.3.1. Legal framework

(328) The Non-horizontal Merger Guidelines recognise that non-horizontal concentrations are generally less likely to significantly impede effective competition than horizontal concentrations.[317]

(329) Vertical non-coordinated effects may principally arise when non-horizontal concentrations give rise to foreclosure,[318] which occurs where actual or potential rivals' access to supplies or markets is hampered or eliminated as a result of the merger, thereby reducing these companies' ability and/or incentive to compete. Such foreclosure may discourage entry or expansion of rivals or encourage their exit. Such foreclosure is regarded as anti-competitive where the Merged Entity — and, possibly, some of its competitors as well — are as a result able to profitably increase the price charged to consumers.[319]

(330) The Non-Horizontal Merger Guidelines distinguish between two forms of foreclosure. Input foreclosure occurs where the merger is likely to raise the costs of downstream competitors by restricting their access to an important input. Customer foreclosure occurs where the merger is likely to foreclose upstream competitors by restricting their access to a sufficient customer base.[320]

(331) In assessing the likelihood of an anticompetitive foreclosure scenario, the Commission examines, first, whether the Merged Entity would have, post-merger, the ability to substantially foreclose access to inputs or customers, second, whether it would have the incentive to do so, and third, whether a foreclosure strategy would have a significant detrimental effect on competition.[321]

---

315 Replies to Questionnaire Q2 to OEMs, question 60; Questionnaire Q3 to End Customers, question 53.
316 Replies to Questionnaire Q1 to Competitors, questions 50.3 and 50.3.1.
317 Non-Horizontal Merger Guidelines, paragraph 11.
318 Non-Horizontal Merger Guidelines, paragraph 18.
319 Non-Horizontal Merger Guidelines, paragraph 29.
320 Non-Horizontal Merger Guidelines, paragraph 30.
321 Non-Horizontal Merger Guidelines, paragraph 32.

(332) As regards ability to foreclose, input foreclosure may lead to competition problems if the upstream input is important for the downstream product and if the vertically integrated Merged Entity has a significant degree of market power in the upstream market. It is only in those circumstances that the Merged Entity can be expected to have significant influence on the conditions of competition in the upstream market and thus, possibly, on prices and supply conditions in the downstream market.[322]

(333) As for customer foreclosure, it is a concern when it involves a company which is an important customer with a significant degree of market power in the downstream market. If, on the contrary, there is a sufficiently large customer base, at present or in the future, that is likely to turn to independent suppliers, the Commission is unlikely to raise competition concerns on that ground.[323]

(334) With respect to incentives to foreclose, the incentive of the Merged Entity to foreclose depends on the degree to which foreclosure would be profitable. In relation to input foreclosure, the Merged Entity faces a trade-off between the profit lost in the upstream market due to a reduction of input sales to rivals and the profit gain, in the short or longer term, from expanding sales downstream or, as the case may be, being able to raise prices to consumers.[324] In relation to customer foreclosure, the trade-off is between the possible costs associated with not procuring products from upstream rivals and the possible gains from doing so, for instance, because it allows the Merged Entity to raise price in the upstream or downstream markets.[325]

(335) As regards the effects on competition, input foreclosure raises competition concerns when it leads to increased prices on the downstream market.[326] If there remain sufficient credible downstream competitors whose costs are not likely to be raised, competition from those firms may constitute a sufficient constraint on the Merged Entity and therefore prevent output prices from rising above pre-merger levels.[327]

(336) By denying competitive access to a significant customer base for the foreclosed rivals' (upstream) products, the merger may reduce their ability to compete in the foreseeable future. As a result, rivals downstream are likely to be put at a competitive disadvantage (*e.g.*, raised input costs), which may allow the Merged Entity to profitably raise prices or reduce the overall output on the downstream market.[328] If there remain a number of upstream competitors that are not affected, competition from those firms may be sufficient to prevent prices from rising in the upstream market and, consequently, in the downstream market.[329]

5.3.2. Affected markets

(337) Pre-Transaction, all NVIDIA's DGX servers include Mellanox InfiniBand interconnects. Therefore, there is a vertical relationship between NVIDIA's

322 Non-Horizontal Merger Guidelines, paragraphs 34-35.
323 Non-Horizontal Merger Guidelines, paragraph 61.
324 Non-Horizontal Merger Guidelines, paragraph 40.
325 Non-Horizontal Merger Guidelines, paragraph 68.
326 Non-Horizontal Merger Guidelines, paragraphs 47-49.
327 Non-Horizontal Merger Guidelines, paragraph 50.
328 Non-Horizontal Merger Guidelines, paragraph 72.
329 Non-Horizontal Merger Guidelines, paragraph 74.

downstream presence in datacentre servers and Mellanox's upstream supply of datacentre network interconnects.

(338) NVIDIA is active in the datacentre server market, while Mellanox is active in the various network interconnect markets (depending on the exact segmentation), which are vertically related markets. A vertically affected market therefore arises in relation to the supply of network interconnects (upstream), which are an input for datacentre servers (downstream).

(339) Tables 4 and 5 below present NVIDIA's and its competitors' market shares in the potential market for datacenter servers and in the potentially narrower market for mid-range servers.

Table 4: Value market shares in datacentre servers, worldwide, 2016-2018[330]

Vendor	2016	2017	2018
Dell EMC	[10-20]%	[10-20]%	[20-30]%
HPE	[20-30]%	[20-30]%	[10-20]%
Self Build/ODM	[5-10]%	[5-10]%	[10-20]%
Inspur Electronics	[0-5]%	[5-10]%	[5-10]%
IBM	[5-10]%	[5-10]%	[5-10]%
Lenovo	[5-10]%	[5-10]%	[5-10]%
Huawei	[5-10]%	[5-10]%	[5-10]%
Cisco	[5-10]%	[5-10]%	[5-10]%
Fujitsu	[0-5]%	[0-5]%	[0-5]%
Oracle	[0-5]%	[0-5]%	[0-5]%
Sugon	[0-5]%	[0-5]%	[0-5]%
H3C	[0-5]%	[0-5]%	[0-5]%
NEC	[0-5]%	[0-5]%	[0-5]%
NVIDIA	**[0-5]%**	**[0-5]%**	**[0-5]%**
Others	[5-10]%	[5-10]%	[5-10]%
Total	**100%**	**100%**	**100%**

Source: Gartner, NVIDIA actual sales data

330 Form CO, Table 6.

Table 5: Value market shares in mid-range servers, worldwide, 2016-2018[331]

Vendor	2016	2017	2018
IBM	[40-50]%	[40-50]%	[30-40]%
HPE	[10-20]%	[20-30]%	[20-30]%
NVIDIA	**[0-5]%**	**[0-5]%**	**[10-20]%**
Oracle	[10-20]%	[10-20]%	[10-20]%
Fujitsu	[0-5]%	[0-5]%	[5-10]%
NEC	[0-5]%	[0-5]%	[0-5]%
Huawei	[0-5]%	[0-5]%	[0-5]%
Inspur Electronics	[0-5]%	[0-5]%	[0-5]%
Others	[10-20]%	[10-20]%	[0-5]%
Total	**100%**	**100%**	**100%**

Source: Gartner, NVIDIA actual sales data

(340) This vertical link is affected because of the market shares of Mellanox, which are above 30% in the upstream markets for high performance fabric ([60-70]% in 2018) and Ethernet NICs of at least 25 Gb/s ([60-70]% in 2018), as illustrated in Tables 2 and 3 (Section 5.2.2 above). In all other plausible network interconnect market segments, Mellanox's market share is significantly lower than 30%.[332]

(341) As NVIDIA is active in the market for datacentre servers, which is vertically related to both the markets for high-performance fabric and for Ethernet NICs of at least 25 Gb/s, it can be concluded that both the datacentre server market as well as the markets for high-performance fabric and for Ethernet NICs of at least 25 Gb/s are affected.

(342) In this decision, the Commission assesses whether the Transaction would likely confer on the Merged Entity the ability and incentive to implement an input foreclosure (Section 5.3.3) and/or a customer foreclosure (Section 5.3.4) strategy with regard to Mellanox's high-performance fabric and Ethernet NICs of at least 25 Gb/s.

5.3.3. Input foreclosure

5.3.3.1. Potential concern

(343) The Commission has assessed the ability and incentive of the Merged Entity to foreclose competing datacentre server suppliers by restricting their access to potentially critical inputs, *i.e.*, Mellanox's high performance fabric and Ethernet NICs of at least 25 Gb/s. In particular, the Commission has investigated whether the Merged Entity could engage in one or both of the following input foreclosure strategies:

- Raising the price of Mellanox's high performance fabric and Ethernet NICs of at least 25 Gb/s when sold to OEMs/end-customers to be incorporated into third party datacentre servers, compared to when they are sold as part of

331 Form CO, Table 21.

332 See paragraph 148 and footnote 136.

NVIDIA's DGX servers (thereby incentivising the purchase of the latter); and/or

- Degrading the quality of Mellanox's high performance fabric and Ethernet NICs of at least 25 Gb/s offered to third parties to be incorporated in third party datacentre servers and/or their compatibility with non-DGX servers.

5.3.3.2. The Notifying Party's views

(344) The Notifying Party claims that the Merged Entity will lack both the ability and the incentive to foreclose access to Mellanox's network interconnects to competing downstream datacentre server suppliers by refusing to supply Mellanox's products to them or by degrading their performance with non-NVIDIA datacentre servers.

As regards ability

(345) First, the Notifying Party submits that the Merged Entity will not be in a position to target and discriminate against rival datacentre server suppliers when supplying network interconnects because sales of datacentre products are largely carried out through intermediaries. Therefore, suppliers generally do not know who the end customers are.[333]

(346) Second, the Notifying Party argues that Mellanox lacks market power in any upstream market for network interconnects.[334]

(347) Third, according to the Notifying Party, Mellanox faces – and will continue to face – strong competitive constraints on any plausible markets for network interconnects. Moreover, only about [10-20]% of all datacentre servers use Mellanox's network interconnects.[335]

(348) Finally, the Notifying Party submits that datacentre server customers' countervailing bargaining power would enable them to divert GPU purchases away from NVIDIA as a response to any hypothetical input foreclosure strategy.[336]

As regards incentives

(349) The Notifying Party submits that the Merged Entity will lack the incentive to foreclose sales of Mellanox's network interconnects when sold to be incorporated in third party datacentre servers in order to increase sales of its own DGX servers.

(350) Since NVIDIA uses its DGX servers as a reference architecture to demonstrate and promote GPUs for use in third party datacentre servers, post-Transaction, NVIDIA will keep relying on OEMs/ODMs as its primary sales channel for GPUs. Moreover, these OEM/ODM customers are also rival datacentre suppliers such that antagonising them would harm NVIDIA's GPU business more than it would help its DGX business.

---

333 Form CO, paragraphs 415-416.

334 Form CO, paragraphs 419-421.

335 Form CO, paragraphs 422-430. See also Form CO, Annex 6.5 – 10, […].

336 Form CO, paragraphs 431-435.

(351) The Notifying Party also claims that the vast majority of Mellanox's network interconnects today are used in third party datacentre servers, customers can easily switch to alternative network interconnects, and NVIDIA is a very small player in the downstream market for datacentre servers. In any event, the Merged Entity would have sufficient sales force to grow significantly its datacentre server business.[337]

As regards effects

(352) According to the Notifying Party, any potential input foreclosure strategy would have no effects on competition, as the Merged Entity would be unable to raise the costs of competing datacentre server suppliers. First, network interconnects represent a very small portion of the cost of datacentre servers. Moreover, in case of a price increase, the small percentage of datacentre server suppliers using Mellanox's network interconnects could easily switch to alternative solutions. The gains from DGX servers would be too small to make the strategy profitable.

(353) The Notifying Party also considers that, post-Transaction, the entry barriers on the plausible network interconnect markets will remain low.[338]

5.3.3.3. Commission's assessment

(354) For the reasons set out below, the Commission considers that the Merged Entity will not have the ability or the incentive to engage in an input foreclosure strategy post-Transaction. Moreover, any such strategy would likely not have any material effect on competition.

As regards ability

(355) The Commission considers that the Merged Entity will likely not have the ability to engage in input foreclosure.

(356) According to paragraphs 34 and 35 of the Non-Horizontal Merger Guidelines, for input foreclosure to raise competition concerns (i) the Merged Entity must have a significant degree of market power in the upstream market and (ii) it must concern an important input for the downstream product.

(357) **As regards the Merged Entity's market power upstream**. As explained in Section 5.2.3.3, on balance, the Commission considers that Mellanox most likely does not have a sufficient degree of market power in the market for Ethernet NICs of at least 25 GB/s to leverage its position in order to influence the choice of the GPU supplier. This is so in particular because there are sufficient credible alternatives to Mellanox's NICs of at least 25 Gb/s (such as Broadcom, Intel, Marvell and Chelsio), and competitors will continue to develop new products that will compete more strongly with Mellanox's products.

(358) The Commission considers that, for the same reasons, Mellanox does not have a sufficient degree of market power in the market for Ethernet NICs of at least 25 GB/s to have a significant influence on the conditions of competition in the upstream

[337] Form CO, paragraphs 437-451.

[338] Form CO, paragraphs 452-457.

market and thus, possibly, on prices and supply conditions in the downstream market.[339] In line with this finding, a few respondents to the market investigation listed alternative suppliers to Mellanox regarding Ethernet NICs of at least 25 Gb/s for their datacentre servers, including, *inter alia*, Broadcom, Marvell and Intel.[340]

(359) In contrast, in Section 5.2.3.3, the Commission has established that Mellanox most likely has a sufficient degree of market power on a high performance fabric market to leverage its position with its InfiniBand fabric in order to influence the choice of the GPU supplier. Similarly, the Commission considers that Mellanox would likely have a sufficient degree of market power on a high performance fabric market to have a significant influence on the conditions of competition in the upstream market and thus, possibly, on prices and supply conditions in the downstream market.[341] This is confirmed by the results of the market investigation. As explained by a large OEM, "*InfiniBand is the #1 must have commercially available fabric on the market.*"[342] Further, end customers noted that "*for performance and scalability reasons we need to deploy high speed and low latency interconnects (for both HPC and AI)*" and that "[c]*urrently there is no real vendor independent alternative with respect to performance*".[343]

(360) Therefore, the Commission considers that the Merged Entity would at most be able to foreclose competing datacentre server suppliers in relation to Mellanox's InfiniBand, but not in relation to its Ethernet NICs of at least 25 Gb/s.

(361) **As regards whether Mellanox's high performance fabric and Ethernet NICs of at least 25 Gb/s constitute important inputs for the downstream market for datacentre servers**. The Commission investigated how important access to Mellanox's InfiniBand and Ethernet NICs of at least 25 Gb/s is for OEMs and end customers when deciding whether to supply/purchase datacentre servers designed for applications/end uses for which DGX servers are used.[344] The majority of respondents to the market investigation that expressed a view indicated that such access was "*essential*" or "*important*" in their decision to supply/purchase datacentre servers.[345]

(362) However, the Commission takes the view that it cannot be concluded from the results of the market investigation that Mellanox's InfiniBand and Ethernet NICs of at least 25 Gb/s are important inputs for the overall downstream market for datacentre servers.

(363) First, when respondents have indicated that access to Mellanox's network interconnects are important, this was in most cases limited to InfiniBand. For example, a large OEM explained that "*InfiniBand is the #1 must have commercially available fabric on the market. With regards to NICs, access to Mellanox's NICs may be essential in some use cases but not all.*"[346] The same OEM noted that, while

339 Non-Horizontal Merger Guidelines, paragraph 35.
340 Replies to Questionnaire Q2 to OEMs, question 63.2.; Questionnaire Q3 to End Customers, question 54.2.
341 Non-Horizontal Merger Guidelines, paragraph 35.
342 Replies to Questionnaire Q2 to OEMs, question 63.1.
343 Replies to Questionnaire Q3 to End Customers, question 54.1.
344 Replies to Questionnaire Q2 to OEMs, question 63; Questionnaire Q3 to End Customers, question 54.
345 Replies to Questionnaire Q2 to OEMs, question 63; Questionnaire Q3 to End Customers, question 54.
346 Replies to Questionnaire Q2 to OEMs, question 63.1.

"[f]*or InfiniBand, there are generally no competitive alternatives commercially available*", alternatives for NICs include "*Broadcom, Marvell, and Intel*".[347]

(364) This is in line with the Commission's findings that while the Merged Entity likely has market power on the market for high performance fabric, it likely does not have market power on the market for Ethernet NICS of at least 25 Gb/s.

(365) Second, the results of the market investigation concern DGX servers and datacentre servers that display the same level of performance and/or are suitable for the same workloads as DGX servers. The importance of InfiniBand for DGX servers is illustrated by the fact that InfiniBand – which equips all DGX servers[348] – appears to represent a significant source of product differentiation for the downstream market for datacentre servers.[349] For example, when launching DGX-2, the vice-president and general manager of Deep Learning Systems at NVIDIA explained that "*Mellanox InfiniBand and Ethernet solutions enable us to give maximum flexibility and performance to customers who build out large-scale clusters of DGX-2 systems*".[350] Similarly, the vice president and general manager of Accelerated Computing at NVIDIA stated that "[t]*ogether, we offer solutions that ensure the most demanding AI applications in the data center benefit from cutting-edge performance and scaling efficiency*."[351]

(366) Similarly, in a document describing IBM's Spectrum Storage for AI, which integrates NVIDIA DGX Systems, IBM highlights the advantages of InfiniBand as part of DGX servers: "[f]*or this reference architecture, the IBM Spectrum Scale on NVMe storage is attached to the DGX-1 or DGX-2 systems by a Mellanox EDR InfiniBand network to provide the most efficient scalability of the GPU workloads and datasets beyond a single DGX system while providing the inter-node communications between DGX systems*."[352]

(367) These considerations appear to suggest that InfiniBand may be considered as an important input for DGX servers. However, in paragraph 119, the Commission established that the market for datacentre servers should not be segmented according to the applications/end uses for which the datacentre servers are designed or used. Therefore, it is not sufficient that InfiniBand would be an important input for part of the market, but not for the majority of the datacentre server market.

(368) Third, most competing suppliers (including Dell and HPE) offer datacentre servers that do not rely on InfiniBand. This is illustrated by the fact that, according to the June 2019 Top500 list, the vast majority of supercomputers (the most powerful computer systems in the world) do not rely on InfiniBand as their network interconnect.[353] As explained in Section 4.3.2.3, high performance fabrics (including InfiniBand) are designed in such a way as to achieve the highest performance possible in very large systems combining several hundreds or thousands of nodes.

---

347 Replies to Questionnaire Q2 to OEMs, question 63.2.

348 Form CO, paragraph 412.

349 Non-Horizontal Merger Guidelines, paragraph 34.

350 https://www mellanox.com/page/press_release_item?id=2040.

351 https://apnews.com/d83fcac978944bb7920d35e19f51ec3b.

352 https://www.ibm.com/downloads/cas/MNEQGQVP.

353 Based on the June 2019 Top500 list, 377 out of the top 500 supercomputers in the world do not rely on InfiniBand, see https://www.top500.org/lists/2019/06/.

Therefore, if the majority of supercomputers do not need InfiniBand, it is likely that most datacentre servers in the world do not need InfiniBand either.

(369) In addition, Cray would continue to offer datacentre servers including its own high performance fabric, without the need to rely on Mellanox's InfiniBand. As explained in paragraph 187, the Commission considers that Cray Slingshot, at least from a technical point of view, is emerging as a credible alternative to Mellanox's InfiniBand in the short term. The recent Frontier and Aurora exascale supercomputers (which Intel and AMD discrete datacentre GPUs have won)[354] support the fact that both Intel and AMD could in principle team up with Cray to compete with a bundle NVIDIA GPU – Mellanox InfiniBand fabric for a given opportunity. This suggests that access to Mellanox's InfiniBand fabric is not essential for Intel and AMD to win discrete datacentre GPU opportunities even for the most demanding HPC/AI applications.

(370) Based on the elements gathered during its investigation, the Commission takes the view that InfiniBand is likely not a critical component without which datacentre servers could not be effectively sold on the downstream market, within the meaning of paragraph 34 of the Non-Horizontal Merger Guidelines.

(371) **As regards the Merged Entity's ability to foreclose rival datacentre server suppliers.** The Commission considers that the Merged Entity would not have such ability. This is because the vast majority of datacentre servers are not equipped with Mellanox's network interconnects. The Parties estimate that only about [10-20]% of all datacentre servers use Mellanox's network interconnects.[355] As a result, irrespective of a potential input foreclosure strategy, there would remain sufficient alternative network interconnect suppliers for the downstream market for datacentre servers (and its possible sub-segments).

(372) In light of the above, the Commission concludes that the Merged Entity will likely not have the ability to foreclose competitors on the market for datacentre servers (and its possible sub-segments) by restricting access to potentially critical inputs.

As regards incentives

(373) The Commission considers that the Merged Entity will lack the incentive to engage in input foreclosure because it would not be profitable to do so.

(374) The Merged Entity would face a trade-off between the potential loss of profit in the upstream market due to a reduction of input sales to downstream competitors and the potential profit gain from expanding sales downstream or raising prices to consumers.[356]

(375) First, while the Merged Entity may gain additional profits on the downstream market due to additional sales of its DGX servers resulting from an input foreclosure strategy, the Commission considers that these profits are likely to be small for the following reasons.

---

354 In addition, Cray announced a third exascale supercomputer win, El Capitan, expected to come online in 2023, see https://www.cray.com/blog/cray-announces-third-exascale-supercomputer-win/.

355 Form CO, paragraph 423 and Annex 6.5 – 10, […].

356 Non-Horizontal Merger Guidelines, paragraph 40.

(376) According to paragraph 42 of the Non-Horizontal Merger Guidelines, the incentive for the Merged Entity to engage in input foreclosure depends on the extent to which downstream demand is likely to be diverted away from foreclosed rivals and the share of that diverted demand that the downstream division of the Merged Entity can capture.

(377) Currently, NVIDIA's sales of DGX servers are small. In 2018, NVIDIA shipped approximately […] DGX servers, amounting to EUR […],[357] compared to a total market size of EUR 65 739 million in the worldwide market for datacentre servers[358] and EUR […] for a plausible narrower worldwide market for mid-range servers.[359]

(378) The Commission considers that the Merged Entity's sales of DGX servers will most likely remain limited post-Transaction. Indeed, according to the Notifying Party, the DGX server is a "reference architecture" platform for NVIDIA to continue to innovate and demonstrate GPU innovations to server OEMs/ODMs, thereby generating demand for its GPUs. NVIDIA provides that innovation and the building blocks of its DGX servers to OEMs, ODMs and CSPs to use in their own server offerings.[360] As explained by the Notifying Party, "*NVIDIA's DGX products are first and foremost reference design architecture platforms, and do not have a substantial footprint in the server market*."[361] A number of NVIDIA internal documents support this strategy.[362] In particular, an NIVDIA presentation to its Board of Directors of [BUSINESS SECRETS – Information redacted regarding business plans].[363]

(379) Furthermore, the Notifying Party confirmed that the "*Transaction will change nothing about NVIDIA's incentive to employ that business model*"[364] and that the Merged Entity "*does not have a plan to take market share from NVIDIA's OEM/ODM server customers following the Mellanox acquisition*".[365] In fact, the Merged Entity "*plans to grow the entire server market by creating improved GPUs, interconnects, and server architectures that it will share with its OEM/ODM customers*".[366] Moreover, the Commission has not found any evidence in NVIDIA's internal documents of a potential change of strategy post-Transaction.

(380) A large OEM confirmed that competition with NVIDIA's DGX servers "*remains rather limited since, so far, Nvidia has only sold a small number of individual DGX servers*" and that it "*does not believe that increasing the sales of Nvidia's DGX servers is a key rationale for acquiring Mellanox. This is because the value of servers is relatively low versus the value of the components sold by Nvidia (e.g., accelerators). Nvidia's motivation for selling cheap servers with their own*

---

357 Form CO, paragraph 180. According to the Notifying Party, this figures covers both direct sales to end-customers and sales to OEMs for integration into their own servers. The Notifying Party notes that in FY 2019, almost […] of DGX sales were made to end-customers and distributors.

358 Form CO, Table 6.

359 Form CO, Tale 21.

360 Form CO, paragraphs 32 and 179; see also Notifying Party's response to RFI 5, question 7.

361 Form CO, paragraph 729.

362 For example, see NVIDIA's internal document, [BUSINESS SECRETS – Information redacted regarding business plans].

363 NVIDIA's internal document, […], slide 6 (Annex RFI 2 – 12).

364 Form CO, paragraph 438.

365 Form CO, paragraph 448

366 Form CO, paragraph 448.

accelerators included therefore mainly relies in the opportunity to increase its sales of accelerators."[367]

(381) Therefore, it is likely that only a small portion of the demand for datacentre servers may potentially be diverted away from rivals. This significantly limits the Merged Entity's incentive to engage in input foreclosure.

(382) Second, the Commission believes that any potential loss of sales on the network interconnect side may be compensated by the additional (although small) sales on the DGX server side. While profit margins for network interconnects are relatively high,[368] these products only represent a small proportion of the total price of datacentre servers. Therefore, even if the Merged Entity increased its sales of network interconnects as a result of an input foreclosure strategy, the profits resulting from these additional sales would likely be smaller than the profits resulting from additional sales of datacentre severs.

(383) In any event, when analysing to what extent foreclosure would be profitable, the Merged Entity would in addition need to take into account potential counterstrategies implemented by customers/OEM. This is particularly relevant as regards the GPU side of the market, where OEMs/ODMs are not only competitors of NVIDIA in the datacentre server market (and, as such, the potential victims of an input foreclosure strategy) but also key customers for NVIDIA's GPUs (and Mellanox's network interconnects).[369]

(384) OEMs/ODMs are the main channel to market for NVIDIA's GPUs.[370] According to the Notifying Party, in 2018, […] of NVIDIA's GPU sales were made to OEMs/ODMs.[371] These include, *inter alia*, […]. The Notifying Party confirmed that [BUSINESS SECRETS – Information redacted regarding business strategy].[372] The Commission has not found any NVIDIA internal document suggesting that this distribution strategy would change post-Transaction.

(385) A restriction of access to Mellanox's InfiniBand by the Merged Entity would mean a loss of flexibility for OEMs/ODMs since it would become uneconomical to equip third party datacentre servers with InfiniBand or, at least, such datacentre servers might not be as performant as required by customers. Consequently, and since they are major customers for NVIDIA's GPUs, these OEMs/ODMs would likely retaliate by shifting (or threating to shift) a significant share of their GPU purchases away from NVIDIA to AMD or Intel. This would most likely result in a substantial loss of sales of GPUs for the Merged Entity.[373]

(386) The reduction of profits on the GPU side would exceed any potential increase in profits resulting from an input foreclosure strategy. This is so because GPUs are a

---

367 See agreed minutes of a conference call of 29 July 2019 with a large OEM, paragraphs 6 and 8.

368 According to the data provided by the Notifying Party, in 2018, Mellanox's gross profit margins were of […] on InfiniBand (including adapters and switches) and […] on Ethernet NICs of at least 25 Gb/s. See Notifying Party's response to RFI 5, question 1, Table 2.

369 Form CO, paragraph 440.

370 Form CO, paragraph 544.

371 Form CO, paragraph 582.

372 Form CO, paragraph 450.

373 Form CO, paragraphs 431-432 and 440.

much more valuable component than other parts of datacentre servers (including network interconnects). According to data provided by the Notifying Party, in fiscal year 2019, NVIDIA's gross profit margin resulting from the sale of GPUs as part of its DGX 1V servers amounted to USD […], corresponding to an […] gross profit margin on GPUs, and representing approximately […] of the total gross profit generated by the sales of DGX 1V. Similarly, NVIDIA's gross profit margin resulting from the sale of GPUs as part of its DGX-2 servers amounted to USD […], corresponding to an […] gross profit margin on GPUs, and representing approximately […] of the total gross profit generated by the sales of DGX-2. The remaining […] and […] of the total gross profit for DGX 1V and DGX-2, respectively, correspond to the gross profit generated by the sale of all the other components of the DGX servers, including the network interconnects.[374]

(387) Finally, the majority of the respondents to the market investigation who expressed an opinion consider that the Merged Entity would not have an incentive to restrict access to Mellanox's interconnects to server OEMs in order to boost its sales of DGX servers.[375] For example, a large OEM explained that "[i]*t's doubtful. It is probably more profitable for Nvidia to sell Mellanox freely than to try to restrict sales to benefit the narrow set of use cases where its DGX product competes.*"[376] Similarly, end customers indicated that "*I doubt it would, the DGX is just a fraction of NVIDIA's HPC sales*", "*Should they do, the market would be very unhappy and it would be a non sense*" and "*I really don't expect that otherwise it's a major mistake from NVIDIA*".[377]

(388) In light of the findings of this section and of the outcome of the market investigation, the Commission concludes that the Merged Entity would not have the incentive to engage into input foreclosure.

As regards effects

(389) The Commission considers that a potential input foreclosure strategy for network interconnects would likely have no material effect on competition. This is so because any potential input foreclosure strategy implemented by the Merged Entity would not result in higher prices for consumers and/or higher barriers to entry for potential competitors in the downstream market for datacentre servers.[378]

(390) First, given the fact that, according to the Notifying Party, only approximately [10-20]% of all datacentre servers use Mellanox's network interconnects,[379] only a limited proportion of rival datacentre server suppliers could potentially be foreclosed. This makes it less likely that the Transaction could be expected to result in a significant price increase – and therefore to significantly impede competition – in the downstream market.[380]

---

374 Notifying Party's response to RFI 5, question 10 and Table 5.

375 Replies to Questionnaire Q2 to OEMs, question 64; Questionnaire Q3 to End Customers, question 55.

376 Replies to Questionnaire Q2 to OEMs, question 64.1.

377 Replies to Questionnaire Q3 to End Customers, question 55.1.

378 Non-Horizontal Merger Guidelines, paragraphs 47-49.

379 Form CO, paragraph 423.

380 Non-Horizontal Merger Guidelines, paragraph 48.

(391) Second, there will remain sufficient credible alternative datacentre server suppliers in the downstream markets that will not be affected by a potential input foreclosure strategy including those suppliers that do not rely on Mellanox's InfiniBand as well as those that are vertically integrated such as Cray.[381] Furthermore, the Notifying Party explains that there are no barriers for customers to switch between different datacentre server suppliers. In fact, customers often "mix and match" different servers within their datacentre.[382]

(392) In light of the above, the Commission considers that a decision post-Transaction to restrict access to Mellanox's InfiniBand would have no material impact on rival datacentre server suppliers.

Conclusion

(393) In light of the above considerations and in view of the outcome of the market investigation, the Commission considers that the Transaction does not raise serious doubts as to its compatibility with the internal market or the functioning of the EEA Agreement with respect to potential input foreclosure.

5.3.4. Customer foreclosure

5.3.4.1. Potential concern

(394) All NVIDIA DGX servers are equipped with network interconnects. Therefore, as post-Transaction the Merged Entity will be vertically integrated, it could decide to source internally the entire quantity of network interconnects it needs for its DGX servers, thereby potentially foreclosing Mellanox's rivals on the upstream plausible network interconnect markets.

5.3.4.2. The Notifying Party's views

As regards ability

(395) The Notifying Party submits that the Merged Entity will lack the ability to engage into vertical customer foreclosure. First, NVIDIA is a small customer of network interconnects on the downstream market for datacentre servers.

(396) Second, NVIDIA does not currently purchase high-performance network interconnects from third parties such that there is no substantial opportunity for diversion of purchases from rivals to Mellanox.

(397) Third, according to the Notifying Party, competing network interconnect suppliers will continue to have access to a number of large, alternative datacentre server OEMs/ODMs.[383]

As regards incentives

---

381 Non-Horizontal Merger Guidelines, paragraph 50.

382 Form CO, paragraph 741.

383 Form CO, paragraphs 460-462.

(398) According to the Notifying Party, given NVIDIA's negligible presence in the datacentre server market, the Merged Entity would have an extremely limited sales base on which to enjoy any price increase downstream.

As regards effects

(399) The Notifying Party submits that any potential customer foreclosure strategy would not have any negative effects on competition in the plausible upstream network interconnect markets notably because NVIDIA has a limited market share in the downstream datacentre server market and NVIDIA currently purchases very few network interconnect products from third parties for its DGX servers.[384]

5.3.4.3. Commission's assessment

(400) For the reasons set out below, the Commission considers that the Merged Entity would not have the ability and/or incentive to engage in a customer foreclosure strategy post-Transaction.

As regards ability

(401) The Commission considers that NVIDIA's DGX servers do not represent a significant channel to market for suppliers of network interconnect solutions.

(402) First, for customer foreclosure to be a concern, the Transaction must involve a company that is an important customer with a significant degree of market power in the downstream market.[385] In 2018, NVIDIA's market share in the overall market for datacentre servers was [0-5]% worldwide.[386] Even considering a plausible narrower market for mid-range servers, NVIDIA's market share remains [10-20]% worldwide.[387] The Commission therefore considers that NVIDIA does not – and post-Transaction, the Merged Entity will not – constitute an important customer with a significant degree of market power in the downstream market for datacentre server.

(403) Second, according to paragraph 61 of the Non-Horizontal Merger Guidelines, if there is a sufficiently large customer base, at present or in the future, that is likely to turn to independent suppliers, the Commission is unlikely to raise competition concerns. The evidence shows that there are indeed sufficient economic alternatives in the downstream market for upstream rivals to sell their network interconnects.[388] Tables 4 and 5 show that these alternatives are large datacentre server suppliers with market shares that are much higher than NVIDIA's, including Dell EMC ([20-30]%), HPE (([10-20]%), Inspur ([5-10]%), IBM ([5-10]%), Lenovo ([5-10]%),

384 Form CO, paragraph 468.

385 Non-Horizontal Merger Guidelines, paragraph 61.

386 NVIDIA's market share on the datacentre server market was [0-5]% by volume and [0-5]% by value, in 2018. Source: Gartner and NVIDIA's sales data. See Form CO, Table 6 and paragraph 460. The Notifying Party notes that these market share estimates do not include the shares of companies that develop and monetise servers but do not sell them on the merchant market, such that NVIDIA's actual market shares.

387 NVIDIA's market share on a plausible mid-range datacentre server market was [10-20]% by value in 2018. Source: Gartner and NVIDIA's sales data. Source: Gartner and NVIDIA's sales data. See Form CO, Table 21.

388 Non-Horizontal Merger Guidelines, paragraph 61.

Huawei ([5-10]%), *etc*.[389] Even on the plausible mid-range datacentre server market, NVIDIA faces strong competitors such as IBM ([30-40]%), HPE ([20-30]%), Oracle ([10-20]%) and Fujitsu ([5-10]%).[390] Therefore, should the Merged Entity decide to foreclose Mellanox's competitors from NVIDIA, rival network interconnect suppliers will retain access to the vast majority of the downstream market for datacentre servers.

(404) Third, [BUSINESS SECRETS – Information redacted regarding business strategy]. As a result, there are no existing sales of high-performance fabric to NVIDIA for rivals to lose. As regards Ethernet network interconnects, the Notifying Party explained that NVIDIA [BUSINESS SECRETS – Information redacted regarding business strategy].[391] The Commission therefore considers that there is no substantial opportunity for diversion of network interconnect purchases from rivals to Mellanox. Given that pre-merger, NVIDIA's purchases from rival network interconnect suppliers represent a very small share of the available sales base for those firms, the potential loss of the Merged Entity as a customer would not represent a significant loss for upstream rivals.[392]

(405) Finally, Mellanox does not have any exclusive contracts with independent downstream customers, which further limits the ability of the Merged Entity to engage in any customer foreclosure strategy.[393]

(406) In light of these considerations, the Commission concludes that the Merged Entity will likely not have the ability to foreclose upstream competitors by sourcing internally all its requirements in high performance fabric and Ethernet NICs of at least 25 Gb/s for its DGX servers.

As regards incentives

(407) The Commission considers that the Merged Entity would not have the incentive to engage in customer foreclosure because it would not be profitable to do so.

(408) In fact, the Merged Entity would face a trade-off between the potential gains (in particular in terms of additional profits) on the upstream markets for high performance fabric and Ethernet NICs of at least 25 Gb/s from foreclosing upstream rivals (allowing the Merged Entity to raise prices in the upstream market) and the possible costs associated with reduced purchases from rival upstream suppliers.[394]

(409) First, the Commission considers that any potential gains made by the Merged Entity in the upstream markets for network interconnects would be limited. This is notably because Mellanox's market power on the high performance fabric market already enables it to extract significant profits from its InfiniBand sales.

---

389 These are market shares on the datacentre server market for 2018. Source: Gartner. See Form CO, Table 6 and paragraph 463.

390 These are market shares on the datacentre server market for 2018. Source: Gartner. See Form CO, Table 21.

391 Form CO, paragraph 468.

392 Non-Horizontal Merger Guidelines, paragraph 61, footnote 1.

393 Non-Horizontal Merger Guidelines, paragraph 61, footnote 1.

394 Non-Horizontal Merger Guidelines, paragraph 68.

(410) Second, any potential losses in the downstream market for datacentre servers (and its possible sub-segments) would likely be limited. [BUSINESS SECRETS – Information redacted regarding business strategy]. Therefore, the costs associated with reducing or stopping purchases from rival network interconnect suppliers would be minimal, if any.

(411) In light of these considerations, the Commission takes the view that it remains unclear whether the Merged Entity would have an incentive to engage in input foreclosure.

As regards effects

(412) The Commission considers that a potential customer foreclosure strategy for network interconnects would likely have no material effect on competition. First, the Commission considers that only a very limited fraction of upstream network interconnects would be affected by a potential revenue decrease resulting from a potential customer foreclosure strategy. This is so because of (i) NVIDIA's very limited position in the downstream market for datacentre servers (and its possible sub-segments) and (ii) the fact that it currently [BUSINESS SECRETS – Information redacted regarding business strategy].[395]

(413) Second, as explained above, rival suppliers are largely protected from any foreclosure strategy because of the existence of multiple suppliers of datacentre servers to which network interconnects suppliers can offer their products.

(414) Therefore, a decision post-Transaction to only purchase network interconnects from Mellanox would have no material impact on rival network interconnect suppliers.

Conclusion

(415) In light of the above considerations, the Commission considers that the Transaction does not raise serious doubts as to its compatibility with the internal market or the functioning of the EEA Agreement with respect to potential customer foreclosure.

6. CONCLUSION

(416) For the above reasons, the European Commission has decided not to oppose the notified operation and to declare it compatible with the internal market and with the EEA Agreement. This decision is adopted in application of Article 6(1)(b) of the Merger Regulation and Article 57 of the EEA Agreement.

For the Commission

(Signed)
Margrethe VESTAGER
Executive Vice-President

395 Non-Horizontal Merger Guidelines, paragraph 74.