

Unclassified

English - Or. English

30 November 2025

**DIRECTORATE FOR FINANCIAL AND ENTERPRISE AFFAIRS  
COMPETITION COMMITTEE**

**Competition in Artificial Intelligence Infrastructure – Note by Portugal**

5 December 2025

This document reproduces a written contribution from Portugal submitted for Item 7 of the 147<sup>th</sup> OECD Competition Committee meeting on 4-5 December 2025.

Antonio CAPOBIANCO  
Antonio.Capobianco@oecd.org, +(33-1) 45 24 98 08

**JT03577845**

## Portugal

### 1. Introduction

1. Artificial Intelligence (AI) is rapidly emerging as a transformative technology, with applications across multiple sectors of the economy. The Portuguese Competition Authority (AdC) has been following the generative AI sector since late-2022. The AdC has published an issues paper on AI in November 2023<sup>1</sup> and initiated a short paper series on AI in 2024. So far, three short papers have been published.

2. In its 2023 Issues Paper, the AdC maps the key determinants that affect the competitive process and anticipates the risks to competition in the Generative AI sector. The AdC concluded that these effects may result in accumulated competitive advantages to digital incumbents, as they already have access to large volumes of data and computing power.

3. The first short paper<sup>2</sup>, from September 2024, focused on the access and use of data in generative AI and its implications for competition. The AdC highlighted the shift from public to proprietary data in AI, and that data licensing agreements had become more prevalent. This results in risks to competition if these agreements include exclusivities or preferential access. The AdC highlights the importance of streamlining access to data for developers to ensure a level playing field (e.g., by serving data through open APIs, pay-as-you-go pricing structures or making public datasets easily available).

4. The second paper<sup>3</sup>, from December 2024, covered issues related to access to AI models by downstream third-party AI developers and the role of openness of AI models in promoting competition and innovation. The AdC highlights that ensuring a level playing field in access to foundation models is key to promote competition and innovation in downstream AI markets.

5. The third paper<sup>4</sup>, from July 2025, addresses competition issues related to access to talent in the generative AI sector. It notes that there is a talent shortage in the AI sector and analyses the competitive effects of strategies such as acquihires, restrictive labour clauses, and agreements between firms not to hire each other's employees. The AdC highlights that such strategies, when undertaken by a dominant firm in the relevant labour market and when likely to exclude rivals, fall under the scrutiny of competition policy.

6. In this note, the AdC aims to contribute with its views on access to computing resources and infrastructure, from a competition perspective and, whenever relevant, with a particular focus on particularities of the Portuguese landscape.

---

<sup>1</sup> Available [here](#).

<sup>2</sup> Available [here](#).

<sup>3</sup> Available [here](#).

<sup>4</sup> Available [here](#).

## 2. Access to computing resources and infrastructure

7. Beyond data, talented know-how and foundation models, the development and deployment of AI depend critically on computing resources and infrastructure, including advanced semiconductors (chips), specialized hardware, high-capacity network connectivity, as well as significant energy resources.

8. Hereafter, this note will explore three layers - chips, supercomputers, and data centres - of the AI infrastructure (physical) inputs, from a competition perspective, while highlighting, where relevant, recent public and private initiatives that have taken place in Portugal in these layers.

### 2.1. Computer chips

9. Upstream to all computing endeavours in the AI space lies access to advanced semiconductors, including both memory and logic chips. These constitute critical inputs for generative AI, both during training and inference. Focusing on logic chips, the most widely used AI chips are Graphics Processing Units (GPUs), which are better suited than traditional Central Processing Units for parallel processing, the cornerstone of the leading AI models.

10. The global market for advanced AI chips, particularly GPUs, is highly concentrated, with a small number of suppliers holding significant market power. The high degree of concentration extends across upstream stages of the semiconductor supply chain, including chip design<sup>5</sup>, software design<sup>6</sup>, memory components<sup>7</sup>, manufacturing<sup>8</sup> and advanced packaging<sup>9</sup>, which are primarily supplied by a limited number of specialised global firms.

11. Barriers to entry stem from large capital requirements, economies of scale, heavy reliance on intellectual property, and dependence on specialised manufacturing facilities (semiconductor foundries), fostering a tendency towards high concentration and benefit early movers and incumbents.

12. Meanwhile the lifecycle of AI chips is also getting shorter, with AI developers demanding increasing computational requirements and the rapid pace of innovation in chip technology<sup>10</sup>. The combination of these demand-side pressures with a high degree of supply

---

<sup>5</sup> Particularly NVIDIA. For a more detailed overview see, for example, section 4.3.2.2.1.2 of the European Commission Decision on Case M.11766 - NVIDIA / RUN:AI, published on April 2025 and available [here](#).

<sup>6</sup> Chip design depends on specialized electronic design automation software supplied by a small number of global providers, including Synopsys, Cadence, Arm and Siemens EDA.

<sup>7</sup> The supply of high-bandwidth memory is supplied by three main suppliers, namely, SK Hynix, Samsung and Micron.

<sup>8</sup> High-end AI chip manufacturing is concentrated in leading foundries, with TSMC being the primary producer of advanced-node chips.

<sup>9</sup> Assembly, Testing and Packaging services are also concentrated, with ASE and Amkor leading globally and TSMC also providing advanced packaging in-house.

<sup>10</sup> A [report](#) finds that the median time between the release of leading AI chips and the publication of the final frontier model trained on them is approximately four years. This trend is also recognized by cloud providers. CoreWeave [disclosed in its IPO filing](#) that the launch of Blackwell immediately

rigidity that makes it difficult to scale production quickly<sup>11</sup> has resulted in an appreciable scarcity of chips in the market which has led to sustained increases in prices<sup>12</sup>.

13. These are more traditional structural-type barriers to entry, and as such differ from the more typical bottlenecks in digital markets that arise at platform-level, driven by network effects<sup>13</sup>.

14. However, network effects may also be present in the chips segment<sup>14</sup>. In fact, GPU suppliers often offer software components to enable complex parallel computing. The interconnectedness of AI hardware and software attracts AI developers and creates a lock-in effect around a given GPU. As more developers use the software component of certain GPUs, more optimizations are built for those GPUs, making those more attractive and resulting in network effects. This type of barrier to entry reinforces market power of already established players in the market.

15. In addition, exclusive deals and partnerships are also becoming more prevalent. For instance, cloud providers and hyperscalers often secure long-term exclusive supply agreements with chip producers<sup>15</sup>. This can further concentrate computational capacity.

16. Overall, this environment raises barriers to entry and expansion for AI developers, due to the high cost of accessing the compute needed to train and deploy advanced models. Most AI developers, particularly startups, are unable to build their own infrastructure because of the substantial upfront investment needed.

17. While cloud computing can mitigate some of these challenges, market concentration in cloud services is high, with limited choice to customers<sup>16</sup>, which may impact AI developers' ability to innovate, and compete effectively.

18. Furthermore, to manage rising costs and reduce dependency on chip suppliers, some cloud providers have begun designing their own chips (*e.g.* Amazon's Trainium and Inferentia<sup>17</sup> and Google's TPUs<sup>18</sup>). Simultaneously chip designers are expanding into cloud

---

devalued its Hopper-based systems (previous NVIDIA's chips), requiring accelerated depreciation. AWS and others have similarly shortened server depreciation timelines ([here](#)).

<sup>11</sup> Namely high fixed costs, long lead times and time-to-market, and limited supplier capacity.

<sup>12</sup> See, for example, the Commission decision pursuant to Article 6(1)(b) of Council Regulation No 139/2004 and Article 57 of the Agreement on the European Economic Area (Case M.11766 – NVIDIA / RUN:AI) ([here](#)), paragraph 148: “NVIDIA's margins are significant: ‘For the full year [fiscal year of 2025], gross margins are expected to be in the mid-70% range’ so they have doubled since the 34% gross margins in 2000.”.

<sup>13</sup> E.g., market segments such as operating systems, social networks, and online marketplaces benefit from network effects.

<sup>14</sup> See, e.g., Hagi, A., & Wright, J. (2025). Artificial intelligence and competition policy. *International Journal of Industrial Organization*, 103134.

<sup>15</sup> See, for example, the partnerships between Amazon's AWS and Marvell ([here](#)), OpenAI and NVIDIA ([here](#)), OpenAI and AMD ([here](#)), or Oracle and AMD ([here](#)), all announced within the last 12 months.

<sup>16</sup> See, e.g., CMA, Cloud services market investigation, July 2025. See information [here](#).

<sup>17</sup> AWS Trainium (see [here](#)); AWS Inferentia (see [here](#)).

<sup>18</sup> Google's TPUs are specifically optimized for Google's own machine learning workloads, demonstrating efficiency in both inference and training processes (see more [here](#) and [here](#)). Over half of the AI workloads on Google Cloud now rely on TPU chips (see [here](#)).

services (e.g. Intel's Tiber AI Cloud<sup>19</sup>, AMD's Developer Cloud<sup>20</sup>, and NVIDIA's DGX Cloud<sup>21</sup>). This will likely reinforce ecosystem incentives that lock-in developers and raise concerns about interoperability, self-preferencing, and foreclosure risk.

19. It is key that competition authorities remain vigilant to the competition risk that incumbents, which may enjoy preferential access to AI infrastructure key inputs, such as chips, grant it to parties on an exclusive basis, or prevent competitors from accessing it, given that would undermine access to those key inputs, reinforcing barriers to entry to new players.

## 2.2. Supercomputing (HPC) capabilities

20. One step downstream in the AI supply chain one encounters supercomputers. These enable at times thousands of interconnected advanced chips in a single location, playing a strategic role in providing high-performance computing (HPC) capacity to AI model training.

21. Using a publicly owned supercomputer can be one way for AI developers to access the compute needed to develop their AI models and services.

22. Some supercomputers are publicly funded and operated under allocation criteria defined by research institutions, universities, ministries or other public entities. Such is the case of the *Deucalion* supercomputer<sup>22</sup>, managed by the RNCA, the Portuguese National Advanced Computing Network, ultimately under the Ministry of Education and Science.

23. Regarding access to public providers of HPC, competition issues relate less to market dominance and more to access conditions and governance structures in place. While public supercomputers have been mainly focused on scientific research, there have been recent developments to adapt access to public HPC capabilities for AI research projects not only directed at research, but also to industry and for AI projects (including generative AI that requires access to GPUs).

24. For example, some recent programmes in Portugal such as *InovIA*<sup>23</sup> have aimed at providing SMEs and start-ups with HPC capabilities, including those of *Deucalion*, for R&D and experimentation purposes.

---

<sup>19</sup> Intel's Tiber AI Cloud was launched in October 2024 as an evolution of the earlier Intel Tiber Developer Cloud, aiming to support production-scale AI workloads rather than just development trials ([here](#)).

<sup>20</sup> AMD's Developer Cloud launched in June 2025 as a developer platform aimed at increasing access to AMD's Instinct GPUs, with preconfigured ROCm software stacks ([here](#)).

<sup>21</sup> In June 2025, NVIDIA extended its DGX Cloud with DGX Cloud Lepton, a global GPU aggregation layer that pools spare capacity across multiple cloud partners. This enables developers to dynamically access unused GPU resources without being locked into a single provider, thereby easing transitions between different cloud platforms. See more [here](#) and [here](#).

<sup>22</sup> Deucalion is a supercomputer located in the Portuguese city of Guimarães featuring 10 PFlop of maximum processing capacity. It was jointly financed by the EuroHPC Initiative, through the European Union's Connecting Europe Facility and the Horizon 2020 research and innovation programme, as well as by the participating states Portugal and Spain. See more information [here](#).

<sup>23</sup> InovIA is a voucher programme for innovation that provides easy access to supercomputers such as Deucalion. It is aimed at start-ups and small and medium-sized enterprises (SMEs) that wish to use advanced computing for innovation. Each initial voucher includes access to 50,000 CPU hours,

25. The capabilities made available to industry players will be further augmented in the future as European Union so-called AI Factories come online<sup>24</sup>. These will eventually be made available for free to interested parties who will access them by participating in calls for interest that will be periodically made, both at national and European level<sup>25</sup>.

26. Based on publicly available information, the European Commission plans an additional 20bn € investment to build so-called AI Gigafactories with the goal that every firm in Europe can access large-scale computing, not just for R&D, but also for the development and deployment of commercial applications<sup>26</sup>.

27. Ensuring fair and open access to supercomputing resources is key to foster competition and innovation in AI. Smaller firms, in particular, may face barriers in accessing the HPC capacity needed to develop AI models and services. Absent such public supercomputing resources, the market may be more at risk of becoming dominated by few global players with privileged access to compute resources.

28. However, HPC resources are scarce, therefore it is key that allocation procedures remain transparent, objective and non-discriminatory. If selection criteria favour certain firms or consortia, this may indirectly distort competition in downstream markets for AI applications. Likewise, public access to HPC should remain, whenever possible, technology neutral.

29. A role for competition advocacy exists in this area: ensuring that the governance of public HPC facilities aligns with technological and competition neutrality, both in terms of access conditions to AI developers, as well as public procurement design for advanced chips and hardware.

### 2.3. Data centre infrastructure

30. Hosts of supercomputers and data centres do not necessarily provide computing services, but rather the necessary resources to host large scale computing operations (including but not exhausted by HPC), namely space, energy, cooling and network connectivity. They provide for the deployment of cost-effective servers that enable mass data storage and allow for the continuous delivery of AI services.

31. Investment in new data centres has increased significantly since 2022<sup>27</sup>, driven by expanding digitalization and the growing use of AI. Although most investment has been,

---

1,000 GPU hours and 1 TB of storage for a period of 6 months. Applications are currently running until 31 December 2026.

<sup>24</sup> Under the EuroHPC Joint Undertaking, the EU intends to establish at least 19 AI Factories: ecosystems built around AI-optimized supercomputers, offering computing resources and support services to the European industry and scientific users for the development of large AI models. See more [here](#) and [here](#).

<sup>25</sup> Ongoing initiatives at the European level can be found [here](#). Submitted proposals for large scale access to AI factories must successfully pass technical and peer-review evaluations organized by the EuroHPC JU Peer-review office to be granted an allocation of compute time (more details [here](#)).

<sup>26</sup> The Commission intends AI gigafactories to be capable of training the most complex AI models and provide at least four times more capacity than the AI factories being currently being set up. Investments will be made through a public-private partnership with the Commission's initial funding for InvestAI (20 bn €) coming from existing EU programmes. More details can be found [here](#) in the Commission's press release.

<sup>27</sup> See, e.g., <https://www.iea.org/commentaries/what-the-data-centre-and-ai-boom-could-mean-for-the-energy-sector>.

so far, in the United States, other locations have also seen increasing investment in this area. One recent example in Portugal<sup>28</sup> is the Data Centre SINES 4.0 project, developed by the private consortium Start Campus, which foresees the construction of six buildings forming a large-scale data centre campus with expansion potential of up to 1.2 gigawatts.<sup>29,30</sup>

32. Conversely, few AI developers have also been investing in their own data centres<sup>31</sup>. Large-scale data centre infrastructure appears to be becoming increasingly concentrated, dominated by vertically integrated global operators (hyperscalers)<sup>32</sup> that are also expanding to other layers of the digital stack, from chip design and cloud platforms to AI model deployment, as previously discussed.

33. Competition concerns in this segment stem from the risk of potential foreclosure which can be materialised, for example, through exclusivity clauses in long-term contracts or, where vertical integration is more prevalent, through self-preferencing and limited interoperability.

34. Whilst, in principle, smaller or independent operators may still emerge to meet the residual demand for data centre infrastructure, high concentration in the provision of land, energy, or network connectivity may hinder market entry.

35. It is noteworthy to mention that consolidation and joint ventures with impact in the access to data centre infrastructure can happen between hyperscalers themselves but also involve other players, such as energy or telecommunications firms. It is thus important to monitor the market developments in terms of identifying any competition risks.

### 3. Concluding remarks

36. AI infrastructure inputs, such as advanced chips, are high capital intensive and characterised by significant economies of scale, which entail strong barriers to entry, fostering a tendency towards high concentration and benefiting early movers and digital incumbents. These are more traditional structural barriers to entry; however, network

---

<sup>28</sup> See, for example, *Portugal DC (2025), Market Outlook Data Centers Portugal 2025* ([here](#)) where “it was calculated that about € 164 million was being invested in the construction of new data centres in Portugal in 2024. In 2025, this grows by more than 900% to € 1.5 billion. Between 2026 and 2031, the multi-year volume of investments is likely to exceed € 11 billion, bringing the total of construction spend to just under € 13 billion across the total time frame.”. The report was commissioned by *PortugalDC*, the national trade organisation representing the interests of the data centre community in Portugal to *Pb7 Research*, an independent data centre research bureau, and published by the former on the 22<sup>nd</sup> of October,

<sup>29</sup> See <https://www.startcampus.pt/>.

<sup>30</sup> In October 2025, Microsoft announced an agreement with Start Campus, in partnership with Nscale, a British data centre company, to rent capacity at the facility, which will use 12,600 of Nvidia’s latest chips, the Blackwell Ultra GB300 GPUs, starting in the first quarter of 2026. See more [here](#) and [here](#).

<sup>31</sup> See, e.g., Google’s recent investment on AI data centre [here](#) that could amount to 15 billion dollars or Amazon’s investment on AI data centres which could amount to 13 billion dollars ([here](#)). See also the recent partnership between OpenAI and Oracle on AI data centres (see [here](#)).

<sup>32</sup> As already acknowledged by the OECD in *OECD (2025), competition in the provision of cloud computing services* (available [here](#)). Cf., table 2 in page 22. See also footnote 19.

effects – commonly present in digital markets – may also be present in AI infrastructure inputs, reinforcing market power.

37. Partnerships and exclusive deals are also becoming more prevalent, including amongst players present in the AI infrastructure layer, which can also pose risks to competition.

38. Ensuring open and competitive access to AI infrastructure (physical) inputs, in addition to remaining key inputs such as data, talent and foundation models, is key to fostering innovation, mitigating market foreclosure and sustaining a diverse AI ecosystem.

39. Therefore, competition agencies should remain vigilant regarding market developments, monitor exclusive or vertically integrated arrangements, and seize opportunities to advocate for open and competitive access to key AI inputs, such as publicly funded compute capacity and interoperability between hardware and cloud services providers to avoid lock-in effects.