

## DIGITAL INFRASTRUCTURE DIVISION · STATEMENT

# Initial Statement

Context, bottlenecks and areas of focus of the Digital Infrastructure Division.

REFERENCE  
BSEIF-DI-ST-01

VERSION  
1.1 · Final

DATE  
10 October 2026

CLASSIFICATION  
Public

SUBJECT  
Founding statement of the Digital Infrastructure Division

LANGUAGE  
English

DISTRIBUTION  
Approved for public release

**SIGN-OFF**  
The document is issued only once the approver has signed. Each signature covers the version and classification printed above.

**PREPARED BY**  
**Lorenzo Barbagelata**  
Head of Digital Infrastructure



Signature 10 October 2026

**REVIEWED BY**  
**Lorenzo Barbagelata**  
Head of Digital Infrastructure



Signature 10 October 2026

**APPROVED BY**  
**Lorenzo Barbagelata**  
Head of Digital Infrastructure



Signature 10 October 2026

**ABSTRACT**

AI adoption now depends less on research and more on three bottlenecks: compute, energy, and capital, with regulation shaping all three. This statement sets out that context, with figures as of 9 October 2026, the five areas the Digital Infrastructure division will cover, and what the division commits to.

BSEIF-DI-ST-01

# Contents

---

|     |                                                         |    |
|-----|---------------------------------------------------------|----|
| 01  | Context                                                 | 3  |
| 02  | Artificial Intelligence as a phenomenon                 | 3  |
| 03  | The first bottleneck: compute                           | 4  |
| 04  | The second bottleneck: energy                           | 5  |
| 05  | The third bottleneck: capital                           | 7  |
| 06  | The issue of regulation                                 | 8  |
| 07  | Other uses of Digital Infrastructure                    | 9  |
| 08  | Areas of focus                                          | 9  |
| 8.1 | AI technology                                           | 9  |
| 8.2 | AI hardware                                             | 9  |
| 8.3 | Data centre energy requirements                         | 10 |
| 8.4 | Data centre financing and CapEx, including hyperscaling | 10 |
| 8.5 | Traditional digital infrastructure                      | 10 |
| 09  | This division's commitment                              | 10 |
| 10  | References                                              | 11 |

---

## 01 Context

The industrial revolution marked the start of the race for capital. Nearly every revolutionary technological advancement since has depended on external investment that allowed visionaries to transform their ideas into commercially viable products. For that is the true goal of progress: not to discover knowledge per se, but to improve the human condition overall. To convince societies to adopt an advantageous technology, the option that has historically proven most reliable is economic convenience.

We currently sit at the start of the AI race, which may well prove to be the next step for humanity after the World Wide Web, with a clear question in mind: is AI worth adopting and developing? Is it economically convenient to do so? The evidence so far is mixed. In customer support, access to a generative AI assistant raised the number of issues resolved per hour by 14% on average, and by 34% for novice workers (Brynjolfsson et al., 2025). In a randomised trial on experienced open-source developers, however, early-2025 AI tools made tasks take 19% longer (Becker et al., 2025). Meanwhile, tools such as Claude Cowork and Claude Code, ChatGPT and Codex, Grok and Grok Bots, Gemini, as well as open-weight alternatives from Meta, Alibaba, Mistral, Z.ai and others, are being adopted faster than their returns can be measured.

The main debate currently is not about the effect AI has on productivity, but rather whether the price of said productivity increase is actually more convenient than hiring more lower-productivity workers. The issue is that the development and usage of artificial intelligence tools is now heavily gated on compute power, capital and energy supply rather than on research alone. Capital is the most visible of the three: Alphabet, Amazon, Meta and Microsoft are expected to spend about \$760 billion on capital expenditure in 2026, against \$413 billion in 2025 (Buchholz, 2026), and Alphabet has already recorded its first negative free cash flow quarter. Consequently, SpaceX and Tesla are planning Terafab, a chip factory in Texas of about 100 million square feet, more than five times the floor area of the largest building in the world today (Hong, 2026), and NVIDIA, a chip designer, has become the most valuable company on Earth, worth about \$5.8 trillion in October 2026 (Crypto Briefing, 2026). Europe and China are also investing heavily in compute, with ASML supplying the lithography machines on which advanced chips are made and Huawei designing its own AI accelerators.

It is in this context that this division's activity is conceived and born.

## 02 Artificial Intelligence as a phenomenon

Contrary to popular belief, the term machine learning was coined in 1959 by IBM pioneer Arthur Samuel (Samuel, 1959), whose checkers-playing program, begun in 1952, became one of the first programs to improve by learning from its own games. However, this and many machines coming after it were far different from modern popularised software like ChatGPT.

The architecture of modern Large Language Models, the main category of frontier AI models and productivity tools, is based on the Transformer neural network, introduced by researchers at Google Brain and Google Research in 2017 in a paper named *Attention Is All You Need* (Vaswani et al., 2017). The first Generative Pre-trained Transformer, GPT-1, was released by OpenAI in June 2018 (Radford et al., 2018), more than four years before ChatGPT. It was highly relevant because it proved

that a machine learning program could be pre-trained on unlabelled text and then fine-tuned to follow the end user's needs. This first model was, however, accessible mainly via Python libraries and, as such, required a computer science background.

The real breakthrough, and the answer to the question 'Why now?', came with the public release of ChatGPT via a chatbot interface on 30 November 2022, which soon attracted curious users. This resulted in Big Tech companies realising that, if properly trained, LLMs could actually be used as productivity tools for tasks ranging from coding to research and, as such, they started diverting resources into the development of said architecture.

At present, frontier and close-to-frontier models can most easily be catalogued into two categories, closed weights and open weights, each supporting a different philosophy. In the case of closed models, as with all Claude models and the flagship ChatGPT models, the end user sends requests and receives outputs via an API key or software connecting the user's hardware to a remote data centre: the request is processed in the data centre and the output is sent back. With open-weight models, instead, the end user can either do the same, or download the binary model files from platforms like Hugging Face or Ollama and then run the model locally, carrying out the calculations directly on the user's hardware.

The main advantage of frontier closed-weight models is that they are, at present, generally considered "more capable" than frontier open-weight ones, though the debate on privacy and data storage is an argument in favour of open-weight models.

### 03 The first bottleneck: compute

The main issue with the LLM architecture is the amount of compute required, both in the training process and to actually run the model. Billions, if not trillions in modern frontier architectures, of parameters need to be adjusted through a repeating process of error-function corrections. For a frontier model this amounts to more than  $10^{25}$  calculations (GPT-4 alone is estimated at about  $2.1 \times 10^{25}$ , and over 30 models had been trained at that scale by mid-2025; Rahman et al., 2025) and, as such, requires GPU or TPU chips. Then, once the training process is complete, the model still requires around two calculations per parameter for every token it generates, and more as the KV cache fills with context. As such, to grant the user a usable speed, the same infrastructure used in training is required to use the model.

This is where modern hyperscalers come into play. Companies like Microsoft, Amazon, Google, Meta, Oracle and Alibaba, together with SpaceX, which absorbed xAI in February 2026 (Owens, 2026), have been heavily investing in compute as they aim to create larger and faster data centres (Exhibit 1). While this has increased the cost of GPUs and RAM (TrendForce forecast server DRAM prices to rise by more than 60% in the first quarter of 2026 alone; Robinson, 2026), leading to an increase in profits for hardware companies, these hyperscalers have so far been able to use their assets to their advantage, either by renting their compute capacity to AI-native companies such as Anthropic and OpenAI, or by using it to develop their own models.

## EXHIBIT 1

**Four hyperscalers will spend about \$760 billion on capital expenditure in 2026, almost double 2025**

| COMPANY                        | 2026 CAPEX GUIDANCE               | LAST UPDATE |
|--------------------------------|-----------------------------------|-------------|
| Amazon                         | c. \$220bn                        | 30 Jul 2026 |
| Alphabet                       | \$195–205bn                       | 22 Jul 2026 |
| Microsoft                      | c. \$190bn                        | 29 Apr 2026 |
| Meta                           | \$130–145bn                       | 29 Jul 2026 |
| <b>Four combined</b>           | <b>c. \$760bn</b> (2025: \$413bn) |             |
| Oracle (fiscal 2027, June–May) | up to \$95bn                      | 10 Jun 2026 |

Note: Guidance as last stated by each company; ranges where the company gave a range.


Source: Buchholz (2026); Snider (2026); TIKR (2026); Outlook Business (2026).

## 04 The second bottleneck: energy

As China's and the USA's compute capacities have increased, and with the energy supply decrease dictated by the conflict in the Middle East, the two effects led to an increase in energy prices that in and of itself threatens the sustainability of compute hyperscaling. The near-total disruption of the Strait of Hormuz cut global oil supply by 10.1 million barrels per day in March 2026 and lifted Brent by about 65% in a single month, its largest monthly rise on record (Agnolucci et al., 2026), with electricity following through the price of gas. Data centres used about 415 TWh of electricity in 2024, around 1.5% of global consumption, and the International Energy Agency (2025) expects this to more than double to about 945 TWh by 2030 (Exhibit 2). In the United States, which accounts for 45% of that demand, the share was already 4.4% in 2023 and could reach 6.7–12% by 2028 (Shehabi et al., 2024). An increase in price therefore threatens the margins and financials of hyperscalers. While this is less of an issue for some of the big caps due to their large profitability outside of the AI and compute businesses, it has created doubts on the sustainability of debt-financed data centres built by some of the more specialised companies, chief of which is Oracle, which plans to raise \$45–50 billion in debt and equity in 2026 alone (Bomey, 2026).

## EXHIBIT 2

### Data centre electricity use more than doubles by 2030 in the IEA base case



Note: 2024 consumption equals about 1.5% of global electricity use.

Source: International Energy Agency (2025).

It has thus become evident that a stable and low-cost source of energy is one of the main drivers in AI development. Energy is an operating cost: a higher price lowers operating margins, thus reducing the cash available to finance compute expansion internally and pushing hyperscalers towards debt, at precisely the moment borrowing has become more expensive (see the third bottleneck). A fluctuating energy price, instead, makes the return on each new data centre harder to forecast and might turn a profitable expansion into an unprofitable one.

That is why I was not at all surprised by the news that, on 9 September 2026, Google signed a 22-year power purchase agreement with Fortum for up to 50% of the output of the Loviisa nuclear power plant, near its Hamina data centre in Finland, reaching that share between 2030 and 2049 (Fortum, 2026). On the same day, Google announced plans to invest at least €13 billion in Finnish digital infrastructure over 2027 and 2028 (Google, 2026). If anything, it made explicit an issue that had only been noticed on the balance sheets: energy prices are as relevant as hardware prices to the evolution of AI and its medium-to-long-term sustainability.

## 05 The third bottleneck: capital

All of this would be impossible without one of the largest capital commitments ever seen. Funding is required for multiple activities, including:

- **Chips.** Many of the hyperscalers are burning through cash and borrowing at the same time. Alphabet recorded negative free cash flow of about \$5.9 billion in the second quarter of 2026, reportedly its first negative quarter in about two decades, on capital expenditure of \$44.9 billion, while its debt grew from about \$16 billion to roughly \$100 billion in twelve months (TIKR, 2026). Others, namely SpaceX, are actively searching for institutional financing and/or issuing bonds to pay for new data centres: SpaceX is reportedly in talks to raise about \$40 billion of debt secured on the NVIDIA chips it buys (Owens, 2026).
- **CapEx.** This is one of the largest issues as of the date of publishing, since CapEx is growing faster than cash generation (Exhibit 1). Due to the complex nature of the topic, it will likely be covered in detail in a future division research paper.
- **Renting.** While hyperscalers are looking for funding to buy chips, AI companies, with Anthropic and OpenAI as prime examples, are looking for funds to secure long-term chip usage, both for training and for running their models, without actually owning the infrastructure on which they are built. This is less true for open-weight models, as it is the end user who mainly needs to invest in hardware, although some data centre renting is still necessary to train them.
- **R&D.** While most headlines today discuss the cost of keeping an AI model running within a data centre, AI is still an extremely young technology that, as such, requires an immense amount of capital, in the form of investment, to improve the current models and tools, both in terms of intelligence and of efficiency.

The capital commitment is large-scale and comes from multiple sources, including (but not limited to) governments, corporate bonds and loans, bank loans, private equity takeovers, and asset management funds' equity. Exhibit 3 maps the largest of these commitments, one for each of the main actors.

### EXHIBIT 3

#### Governments, chip makers, banks and asset managers are all financing the build-out

| INVESTOR                          | COMMITMENT                                                               | SIZE                 | DATE                   |
|-----------------------------------|--------------------------------------------------------------------------|----------------------|------------------------|
| Chinese government                | Big Fund III, state-backed semiconductor fund                            | ¥344bn (c. \$47.5bn) | May 2024               |
| Blackstone                        | Acquisition of AirTrunk data centres, Asia-Pacific                       | A\$24bn              | Sep 2024               |
| NVIDIA                            | AI infrastructure financing initiative, with a backstop of up to \$125bn | over \$500bn target  | Aug 2026               |
| Goldman Sachs                     | Lender on NVIDIA's initiative, alongside Blackstone and Apollo           | n/a                  | Aug 2026               |
| BlackRock (GIP), with AIP and MGX | Acquisition of Aligned Data Centers, 6.4 GW                              | \$40bn               | Jul 2026               |
| Oracle                            | Debt and equity raise for cloud capacity                                 | \$45–50bn            | Feb 2026               |
| SpaceX                            | Debt talks led by Apollo, secured on NVIDIA chips                        | c. \$40bn            | Oct 2026 (re-reported) |

Source: Caixin Global (2024); Blackstone (2024); Azhar and Vinn (2026); Johnson (2026); Bomey (2026); Owens (2026).

Borrowing is, however, currently made more expensive by the rise in interest rates caused by inflation fears stirred by the conflict in the Middle East. On 16 September 2026 the Federal Reserve raised its target range by 25 basis points to 3.75–4.00%, its first hike since 2023, and the US 10-year Treasury yield touched 5.04%, its highest level since 2007 (FXStreet, 2026). US Treasuries carry much of the blame because most of these companies are US-based and price their debt off them.

## 06 The issue of regulation

As is often the case with newer technologies, regulatory entities are unsure how to approach the topic of AI. While some industry leaders call for lighter regulation, namely NVIDIA's Jensen Huang, who argued in September 2026 that the industry needs no new laws (Bort, 2026), and Meta's Mark Zuckerberg, who has warned that pre-emptive regulation of open-source AI will stifle innovation (Zuckerberg & Ek, 2024), others call for stricter development laws and use cases, chief amongst whom is Anthropic's Dario Amodei. On 12 September 2026 he called for a slowdown in the pace at which AI capabilities improve, a call Sam Altman and Elon Musk publicly endorsed (Berkowitz, 2026). The irony of the lead AI company's CEO asking for a slowdown ten days before the release of Claude Opus 5.5 (Let's Data Science, 2026), Anthropic's most capable generally available model, is not lost on me.

World governments are also unable to find common ground. While the People's Republic of China pushes for the idea of sovereign AI via open-weight models and state-led financing, such as the third National Integrated Circuit Industry Investment Fund (Big Fund III), with ¥344 billion (about \$47.5 billion) of registered capital (Caixin Global, 2024), the EU has adopted a much more conservative approach. Its AI Act makes the marking of AI-generated content mandatory from 2 August 2026, or from 2 December 2026 for systems already on the market (Latham & Watkins,

2026), while capital has moved more slowly: the InvestAI initiative aims to mobilise €200 billion for AI, including €20 billion for AI gigafactories, to build a European competitor to the USA's and China's mega-caps (European Commission, 2025). The United States government is assuming a deregulatory but interventionist approach: it does not limit development, but it seeks control over how frontier models are used. The dispute between Anthropic and the Pentagon became the most famous case. After Anthropic refused to allow its models to be used for fully autonomous weapons or mass domestic surveillance, the Department of Defense designated it a supply-chain risk, a designation the D.C. Circuit upheld on 25 September 2026 (Urwin, 2026), while the White House separately negotiated access to Mythos Preview for civilian agencies (Curi et al., 2026).

Legislative and regulatory bodies thus hold massive power, but face a large dilemma as well: how much regulation is too much?

## 07 Other uses of Digital Infrastructure

While the focus of research is going to be on AI, AI infrastructure and AI financing, digital infrastructure, albeit still a relatively new branch of infrastructure, does not only revolve around Artificial Intelligence. Its main other use is the maintenance of the Internet, which about 6 billion people, roughly three quarters of the world's population, used in 2025 (International Telecommunication Union, 2025). Data storage, web hosting, fibre networks and telecommunication towers, content delivery networks, hardware firewalls, DNS servers, etc., are all pieces of infrastructure dedicated to the upkeep of the digital world. Much of it behaves like traditional infrastructure, with long-lived assets and revenues contracted over many years, which is why infrastructure funds such as BlackRock's Global Infrastructure Partners now buy data centre platforms (Exhibit 3). As such, these assets may be touched upon in future papers.

As is easy to deduce, these data centres also suffer from the increased price of RAM, energy and capital.

## 08 Areas of focus

The division's research will be organised around five areas, which follow the bottlenecks and dynamics described above. Regulation is not an area of its own: it runs across all five and will be addressed wherever it bears on the analysis.

### 8.1 AI TECHNOLOGY

How models and their uses evolve, and what that means for computing demand.

### 8.2 AI HARDWARE

Chips, accelerators and servers, and what each generation asks of a data centre, as well as evolutions of said pieces of hardware.

### 8.3 DATA CENTRE ENERGY REQUIREMENTS

Power and cooling, and the grid connections and contracts that secure them.

### 8.4 DATA CENTRE FINANCING AND CAPEX, INCLUDING HYPERSCALING

How data centres are funded and what they cost, from single projects to hyperscaler capital programmes.

### 8.5 TRADITIONAL DIGITAL INFRASTRUCTURE

How the specialisation of digital infrastructure towards AI is competing for resources with the internet-purposed server farms.

## 09 This division's commitment

Bocconi Students Energy & Infrastructure Finance (BSEIF for short) is a student-led non-profit association whose goal is the dissemination of knowledge on Energy, Infrastructure, Digital Infrastructure, Markets and Venture, both through the publishing of research papers and the organisation of events. Within it, the Digital Infrastructure division exists to examine the question this note opened with in the terms in which it is actually being decided: compute, energy, capital and regulation. Whether AI is economically convenient will not be settled by a single paper, and the division does not claim to settle it. It therefore commits to the following:

- **Output.** The division aims to publish its first two research papers before BSEIF's mid-November 2026 event, each within one of the areas of focus above.
- **Method.** The division works as a flat team of analysts from several countries, who assign, review and challenge each piece of work before publication. Members' different regional perspectives are deliberate, as the build-out does not look the same from every market.
- **Standards.** Every figure carries its source and the date of that source, estimates and forecasts are labelled as such, assumptions are disclosed, and citations follow APA. Each paper states its data cut-off, and material errors found after publication are corrected in a new, dated version.
- **Independence.** Papers present the analysis of their authors and of the division, not the view of any company, sponsor or speaker, and do not constitute investment advice.
- **Skills and industry.** Members build the technical skills the subject requires, from reading hyperscaler disclosures to modelling a data centre's CapEx, and engage with practitioners at BSEIF's events.

All papers and events will be published or announced on the association's official website, [bseif.org](https://bseif.org), where departmental research appears under Insights.

## 10 References

- Agnolucci, P., Makarenko, N., & Temaj, K. (2026, May 7). *Strait of Hormuz disruption sends oil prices surging*. World Bank Blogs. <https://blogs.worldbank.org/en/opendata/strait-of-hormuz-disruption-sends-oil-prices-surging>
- Azhar, S., & Vinn, M. (2026, August 14). *Goldman in talks with investors on Nvidia financing deal after landing prized role, sources say*. Reuters, via The Star. <https://www.thestar.com.my/tech/tech-news/2026/08/14/goldman-in-talks-with-investors-on-nvidia-financing-deal-after-landing-prized-role-sources-say>
- Becker, J., Rush, N., Barnes, E., & Rein, D. (2025, July 10). *Measuring the impact of early-2025 AI on experienced open-source developer productivity*. METR. <https://metr.org/blog/2025-07-10-early-2025-ai-experienced-os-dev-study/>
- Berkowitz, B. (2026, September 12). *Anthropic, OpenAI CEOs call for slowdown in AI development*. Axios. <https://www.axios.com/2026/09/12/anthropic-ai-amodei-pacing>
- Blackstone. (2024, September 5). *Blackstone announces agreement to acquire AirTrunk in a A\$24B transaction*. <https://www.blackstone.com/news/press/blackstone-announces-agreement-to-acquire-airtrunk-in-a-a24b-transaction/>
- Bomey, N. (2026, February 2). *Oracle to raise \$50 billion as AI debt piles up*. Axios. <https://www.axios.com/2026/02/02/oracle-50-billion-ai-chatgpt-meta>
- Bort, J. (2026, September 15). *We don't need AI regulation, leave safety to us, Nvidia's Jensen Huang says*. TechCrunch. <https://techcrunch.com/2026/09/15/we-dont-need-ai-regulation-leave-safety-to-us-nvidias-jensen-huang-says/>
- Brynjolfsson, E., Li, D., & Raymond, L. (2025). Generative AI at work. *The Quarterly Journal of Economics*, 140(2), 889–942. <https://doi.org/10.1093/qje/qjae044>
- Buchholz, K. (2026, July 31). *Big Tech's AI spending to reach \$760 billion in 2026*. Statista. <https://www.statista.com/chart/35046/capital-expenditure-of-meta-alphabet-amazon-and-microsoft/>
- Caixin Global. (2024, May 28). *China piles \$47.5 billion into Big Fund III to boost chip development*. <https://www.caixinglobal.com/2024-05-28/china-piles-475-billion-into-big-fund-iii-to-boost-chip-development-102200633.html>
- Crypto Briefing. (2026, October 6). *Nvidia's valuation lead over Apple stretches toward \$1 trillion*. <https://cryptobriefing.com/nvidia-lead-over-apple-market-cap/>
- Curi, M., Caputo, M., & Lawler, D. (2026, April 16). *Trump officials weigh Anthropic Mythos access amid Pentagon blacklist*. Axios. <https://www.axios.com/2026/04/16/white-house-anthropic-ai-mythos-government-national-security>
- European Commission. (2025, February 11). *EU launches InvestAI initiative to mobilise €200 billion of investment in artificial intelligence*. <https://digital-strategy.ec.europa.eu/en/news/eu-launches-investai-initiative-mobilise-eu200-billion-investment-artificial-intelligence>
- Fortum. (2026, September 9). *Fortum and Google partner to drive sustainable growth for Finland: Sign nuclear power purchase agreement*. <https://www.fortum.com/en/media/2026/09/inside-information-fortum-and-google-partner-drive-sustainable-growth-finland-sign-nuclear-power-purchase-agreement>
- FXStreet. (2026, September 17). *US 10-year Treasury yield nears 5% after Fed rate hike, Middle East tensions raise inflation fears*. <https://www.fxstreet.com/news/us-10-year-treasury-yield-nears-5-after-fed-rate-hike-middle-east-tensions-raise-inflation-fears-202609170501>
- Google. (2026, September 9). *Google deepens commitment to Finland with two-year €13 billion investment in AI infrastructure*. <https://www.googlecloudpresscorner.com/2026-09-09-Google-Deepens-Commitment-to-Finland-with-Two-Year-EUR13-Billion-investment-in-AI-Infrastructure>
- Hong, J. (2026, August 10). *Musk's planned \$16.8 billion factory is five times bigger than the world's current largest building*. Fortune. <https://fortune.com/2026/08/10/elon-musk-chip-factory-spacex-tesla-texas/>

- International Energy Agency. (2025). *Energy and AI*. International Energy Agency. <https://www.iea.org/reports/energy-and-ai>
- International Telecommunication Union. (2025). *Facts and figures 2025*. International Telecommunication Union. <https://www.itu.int/ITU-D/ict/statistics/>
- Johnson, L. (2026, July 22). *BlackRock's GIP, Microsoft-backed AI group buy Aligned Data Centers for \$40B*. ESG Dive. <https://www.esgdive.com/news/blackrocks-gip-microsoft-backed-ai-group-buy-aligned-data-centers-for-40/825920/>
- Latham & Watkins. (2026, May 13). *AI Act update: EU resolves to change rules and extend deadlines*. <https://www.lw.com/en/insights/ai-act-update-eu-resolves-to-change-rules-and-extend-deadlines>
- Let's Data Science. (2026, September 22). *Anthropic launches Claude Opus 5.5 with safeguards*. <https://letsdatascience.com/news/anthropic-launches-claude-opus-55-with-safeguards-5d52f322>
- Outlook Business. (2026, June 11). *Oracle forecasts \$95 bn in capex for FY27, plans to raise \$40 bn in debt and equity*. <https://www.outlookbusiness.com/corporate/oracle-forecasts-95-bn-in-capex-for-fy27-plans-to-raise-40-bn-in-debt-and-equity>
- Owens, J. (2026, October 7). *SpaceX bond deal: 30-year maturities, Nvidia chip collateral worth fraction by mid-decade*. Tech Times. <https://www.techtimes.com/articles/328723/20261007/spacex-bond-deal-30-year-maturities-nvidia-chip-collateral-worth-fraction-mid-decade.htm>
- Radford, A., Narasimhan, K., Salimans, T., & Sutskever, I. (2018). *Improving language understanding by generative pre-training*. OpenAI. [https://cdn.openai.com/research-covers/language-unsupervised/language\\_understanding\\_paper.pdf](https://cdn.openai.com/research-covers/language-unsupervised/language_understanding_paper.pdf)
- Rahman, R., Heindrich, L., Owen, D., & Emberson, L. (2025). *Over 30 AI models have been trained at the scale of GPT-4*. Epoch AI. <https://epoch.ai/data-insights/models-over-1e25-flop>
- Robinson, D. (2026, January 6). *Mem-ageddon: AI chip frenzy to wallop DRAM prices with 70% hike*. The Register. [https://www.theregister.com/2026/01/06/memory\\_firm\\_profits\\_up\\_as/](https://www.theregister.com/2026/01/06/memory_firm_profits_up_as/)
- Samuel, A. L. (1959). Some studies in machine learning using the game of checkers. *IBM Journal of Research and Development*, 3(3), 210–229. <https://doi.org/10.1147/rd.33.0210>
- Shehabi, A., Newkirk, A., Smith, S. J., Hubbard, A., Lei, N., Siddik, M. A. B., Holecek, B., Koomey, J., Masanet, E., & Sartor, D. (2024). *2024 United States data center energy usage report* (LBNL-2001637). Lawrence Berkeley National Laboratory.
- Snider, S. (2026, July 31). *Amazon lifts 2026 AI capex to \$220B, still capacity-constrained*. Data Center Knowledge. <https://datacenterknowledge.com/infrastructure/amazon-lifts-ai-infrastructure-spending-to-220b-as-demand-outpaces-capacity>
- TIKR. (2026, September 18). *Alphabet's free cash flow just went negative. Here's what the debt behind it says*. <https://www.tikr.com/blog/alphabets-free-cash-flow-just-went-negative-heres-what-the-debt-behind-it-says>
- Urwin, M. (2026, September 29). *Anthropic vs. Pentagon: Fight over Claude access*. Built In. <https://builtin.com/articles/anthropic-pentagon-claude-dispute>
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*, 30. <https://papers.nips.cc/paper/7181-attention-is-all-you-need>
- Zuckerberg, M., & Ek, D. (2024, August 23). *Mark Zuckerberg and Daniel Ek on why Europe should embrace open-source AI*. Meta. <https://about.fb.com/news/2024/08/why-europe-should-embrace-open-source-ai-zuckerberg-ek/>