



The inflow of investment and research into artificial intelligence (AI) shows no sign of abating, with interest increasingly shifting to more customised and targeted uses of the technology. However, for all but the largest companies seeking to harness AI, doing so requires renting computing capacity from specialised hardware providers. Rental arrangements for GPUs are now a common feature of the AI environment. The market for GPU rental is maturing rapidly, with greater choice, flexibility, product offerings and regulation all coming into play. As we explore in this article, these dynamics bring legal issues that, in contrast to the trailblazing rate of technological progress in this area, are surprisingly familiar.

As large-scale computing power becomes more attainable, competition between rental providers is growing. Hyperscalers are increasingly competing with neoclouds: smaller, more flexible operators who often cater to more specialised functionality. This has led to more of a buyer's market for cloud-based computing arrangements, with reduced prices and more flexibility around consumption. This trend looks set to continue, with new entrants, growing investment, additional regulation and the development of more advanced chips all widening the choice available to prospective customers.

These commercial dynamics are impacting contracting arrangements. Buyers have greater scope and incentive to negotiate shorter-term contracts, and/or more flexible offtake arrangements. Sellers, who traditionally need buyers to support the capital expenditure associated with these facilities, would prefer to lock in buyers to longer-term contracts with firm payment obligations, but may find it increasingly difficult to do so. The greater availability of cloud computing providers has already purportedly made GPU rental agreements of between one to three years in length much rarer, with customers choosing to contract across shorter terms.¹

We see similarities between these trends and those witnessed in more established commodities markets. For example, sale and purchase agreements for liquefied natural gas (LNG) were historically drafted in long, inflexible terms that locked in buyers to onerous "take or pay" arrangements. Sellers required these commitments to secure financing for the large-scale capital expenditure needed to make the extraction and processing of LNG viable. However, as more sources of supply came online, buyers were gradually able to demand terms more beneficial to them. This impacted contracting practices, with mechanisms like periodic price reviews, break clauses, upward and downward flexibility provisions, as well as change of circumstances provisions becoming more and more common. In parallel, short-term spot markets developed, with prices being dictated by regional hub indexes, rather than individually negotiated between buyers and sellers. In short, contracting practices adapted to the proliferation of choice in a capital-intensive, commoditised market.

¹ D Patel, T H Chen, D Nishball, I Chiam, R Knuhtsen, The GPU Cloud ClusterMAX Rating System / How to Rent GPUs, Semianalysis, 26 March 2025 (<https://semianalysis.com/2025/03/26/the-gpu-cloud-clustermax-rating-system-how-to-rent-gpus/>).

Of course, AI companies as customers have other commercial priorities too. For these actors, protection of their intellectual property is crucial, so they may demand robust security commitments from the GPU providers on which they rely. Harnessing GPU clusters is also complex and technically difficult. This creates value for providers who offer comprehensive technical support and know-how. Because providing these services involves more costs, it can lead to trade-offs. The smaller neoclouds providing the greatest flexibility or specialisation of GPU capacity may not be able to offer the same security and technical Key Performance Indicators (KPIs) as a more established hyperscaler.

All of these dynamics have the potential to breed disputes. In the fast-paced world of AI, even buyers committed to shorter-term GPU rental arrangements may find themselves commercially disadvantaged if cheaper and more flexible options become available. Contracts signed at haste or on ill-fitting, standard-form terms may not adequately address parties' needs, particularly in periods of disruption. Intricate or bespoke flexibility mechanisms that are vague, or even contradictory, can give rise to confusion or end up being applied oppressively. Blanket limitation of liability clauses might leave an innocent party significantly out of pocket when there is a breach of contract by the other. Finally, the regulatory landscape for cloud-based computing is undergoing significant change across many jurisdictions, with the impact on private commercial contracting practices yet unknown. These are all sources of legal risk that could conceivably affect GPU rental arrangements, but these can also be mitigated if given appropriate thought at the outset.

Seeking advice from external counsel will always be helpful in this regard. But there are still practical points that any party negotiating a contract will benefit from bearing in mind – and that are particularly relevant to GPU rental arrangements, given the issues described above:

- **Make sure what you have agreed is reflected in your contract.** Merely discussing something during negotiations, or even relying on customary market practice, is often not enough to create a binding contractual obligation, however uncontroversial you may consider the point to be at the time.
- **If your agreement is spread across more than one document, think about how these interact with each other.** Commercial contracts often provide for how discrepancies between, for example, framework agreements and standard terms are to be resolved, but this may not be the case or these priority mechanisms might operate in a way that is prejudicial to your interests.
- **Give thought to provisions that leave something to future agreement, or at the discretion of one party.** When parties' commercial interests diverge, these provisions can be exercised in ways that are unexpected, or which harm one side.

Assessing in advance how this might transpire, and ensuring clarity in the drafting of any necessary clauses, is often more prudent than assuming (or hoping) that your counterparty will behave in a certain way.

- **Assess whether any exclusion or limitation of liability provisions are appropriate.** This involves thinking about what you might stand to lose if your contract isn't performed: for AI companies working to tight product development timetables, the financial cost of even a small period of interruption may be particularly significant. Now consider whether any liability caps specified in the contract would adequately compensate you. If not, it may be sensible to negotiate further. Disputes often centre upon the application and interpretation of these provisions, rather than the existence of the underlying breach itself.
- **Be aware of forthcoming changes to the regulatory landscape in the jurisdictions that are relevant to your arrangement and how these may impact your contract.** If the contractual terms are not aligned with any such regulation, this may make enforcing those terms more difficult. As far as preexisting contracts are concerned, parties may well benefit from giving thought to how performance might be affected by future regulatory and legislative changes.
- **Think about how you would enforce the contract, if needed.** This involves making sure you understand which court or arbitration process you would need to follow to do this. This is not only significant in terms of obtaining a financial remedy for any breach. Different rules around the confidentiality of any dispute may apply, which can be very significant for parties operating in technology sectors. In addition, the nature of activities taking place under the agreement may mean you need to force the other party to do or not do something (for example in the event a data centre is compromised by a cyber-attack). The availability of wider remedies is not uniform across courts or jurisdictions.

Computing and AI are complex and exciting fields that look set to profoundly reshape life as we experience it. GPU rental agreements are at the vanguard of this technological revolution. In contrast, the legal principles governing these transactions are well-established and time-honoured, with parallels in the development of other industries in the past. With so much interest and investment in this sector, parties will be well-served by not overlooking the same legal pitfalls that have been afflicting commercial relationships for centuries. These issues, however mundane, are far more enduring than the average AI chatbot.

Contact



Tim Flamank
Director, London
T +44 207 655 1062
E tim.flamank@squirepb.com