
Can the Iran War Derail the AI Trade?

The question is not whether the Iran war is negative for the AI trade. Clearly it is. The more relevant question is whether it changes the direction of travel, or simply raises the cost, complexity and selectivity of the build-out. AI infrastructure is a long and highly interdependent supply chain. It starts with power, fuel and financing, runs through chips, memory, cables, racks, transformers, switchgear and cooling systems, and ends with the data centres themselves. A disruption in the Gulf does not need to halt GPU shipments outright to matter, it only has to tighten a few already fragile links.

Power prices and the cost of running AI

One obvious concern is energy. Data centres are power-intensive assets, and many electricity markets still clear off gas at the margin. As explained above, even if the Strait of Hormuz could fully reopen within months, global energy markets will embed a higher risk premium in oil and gas prices going forward. That is clearly negative for operators and developers of data centres, particularly in Europe and Asia, which are more exposed to global gas prices.

The counterarguments are that most existing and planned capacity sits in the US, where natural gas prices are down since the outbreak of the war, and that hyperscalers are so profitable that even a substantial increase in power prices is unlikely to alter their build-out plans. A more realistic risk is that if elevated power prices persist, especially outside the US, regulators and utilities impose stricter additionality and “bring-your-own-power” requirements, complicating grid access and adding to project complexity.

Chip supply and the helium constraint

The more serious risk is physical rather than financial. Semiconductors remain the most valuable and constrained component of the AI stack, not only GPUs from **Nvidia (NVDA)** but also high-bandwidth memory from suppliers such as **Micron (MU)** and **SK Hynix**. These markets were already supply-constrained before the war, with backlogs stretching well into next year.

Higher energy costs are clearly unwelcome for energy-intensive industries, but with customers desperate for compute, much of the incremental cost can likely be passed through. That alone is unlikely to stop the AI build-out. The bigger concern lies in inputs that could physically limit production.

Helium is the most important of these. It is essential for cooling, plasma etching and EUV lithography, and there is no viable substitute in advanced semiconductor manufacturing. Helium is a by-product of gas processing, and Qatar accounts for more than a third of global supply. Its production has been offline since early March following Iranian drone strikes, and even with a partial reopening of Hormuz under the ceasefire negotiations, it remains unclear when LNG and helium production can restart.

South Korea and Taiwan, which dominate advanced semiconductor manufacturing, source a large share of their helium from Qatar and have already triggered mitigation measures, including helium recycling on selected lines that could reduce demand by 15-20% if fully implemented. In the meantime, buffer stocks estimated at three to six months are being drawn down. While helium represents less than 0.5% of chip production costs, the spot market is extremely thin as most volumes are sold under long-term contracts. If supply is not restored by late-summer, declining inventories risk becoming a binding constraint on AI chip output.

That said, it is worth maintaining perspective. Semiconductor manufacturing makes up about 25% of global helium demand and the industry’s ability to absorb cost increases in the current AI driven boom

will price out any other users. Industrial welding, for example, accounts for about 8% of US helium demand, but since helium can be substituted with argon, we would expect long term contracts in such segments to be resold at a premium to the semiconductor customers. If not, government are likely to intervene and reallocate volumes to the economically and strategically critical semiconductor industry.

Financing risk, private credit and Gulf capital

Financing is another transmission channel. Historically, sustained energy shocks, even if supply driven, have tightened financial conditions as central banks raise rates to contain inflation. Given that data centre development is very capital intensive, higher rates increase total construction costs and can eventually weigh on demand. For now, the picture remains mixed as the ECB is expected to tighten further, while the Fed has signalled that near term hikes are unlikely.

Even without additional rate increases, financial conditions have tightened since the war began. While hyperscalers are largely internally financed and can absorb higher capital costs, smaller data-centre developers, neo clouds, power project SPVs and many AI related start ups all rely on external capital and may struggle to raise funds on attractive terms. Key equipment suppliers, themselves a bottleneck in the AI build out, may also hesitate to pursue aggressive capacity expansion if their own funding costs rise. Still, it is important to remember that the rate-insensitive hyperscalers account for the majority of AI capex.

Concerns around private credit have added another layer of uncertainty. Several funds gated redemptions earlier this year, raising fears about the availability of non bank capital. That risk is real, but it matters how it is framed. Blue Owl, which sat at the centre of the debate, reportedly sold assets at close to 99.7 percent of par within a week of the gating, while restrictions imposed by banks such as JP Morgan, which made headline news, relate to lending against private credit funds seeking leverage, rather than to end projects themselves. This is more analogous to adjusting repo haircuts than to a withdrawal of capital. Most capital in private credit vehicles is unlevered, 'real money' funding; while defaults in software or AI companies will reduce returns for savers, the risk looks manageable rather than systematic in the near to medium term.

Gulf capital is another area of focus. Sovereign wealth funds from the region have been important partners in data centre and AI projects. Their ability to deploy abroad could be constrained if capital is needed domestically to rebuild infrastructure or offset lost energy revenues. That risk exists, but for now capital availability is seemingly ample. For example, OpenAI's latest funding round reportedly closed at around USD 122bn raised at a post-money valuation in the USD 850bn range, which hardly points to a market starved of capital, at least not yet.

Policy risk

Policy risk was rising even before the war. Political resistance to data centre power consumption has been building as grid constraints become more visible. Although US power prices have not been impacted by the Iran war, the gasoline price has risen nearly 40% since before the conflict. This could make consumers even more hostile towards future data centres as their costs are squeezed from several directions. Stricter rules around grid access, demand response, emissions, local content and "bring your own power" requirements appear increasingly likely. This will lengthen timelines, raise capex and shift risk away from smaller, less well capitalised developers.

Conclusion

We do not currently see the Iran war as a direct showstopper for the “Powering AI” theme, provided two conditions are met: 1) helium production in Qatar restarts by late-summer, and 2) oil and gas flows through Hormuz normalise sufficiently to avoid a sustained energy price shock that forces broad demand destruction.

Even under those assumptions, AI infrastructure is likely to become more expensive, more regulated and more geographically selective. Higher energy costs, tighter credit and greater physical and political risk will raise the cost of capital and tilt the playing field in favour of large, well capitalised hyperscalers over smaller neo clouds. On-site and behind the meter power and grid upgrades are likely to gain share versus marginal developments, while companies that can pass through costs and manage complexity will be preferred over those reliant on cheap funding and benign policy.

Crucially, demand for compute still appears to be growing faster than supply, and hyperscaler capex intentions remain intact. In our view, this points to an AI investment cycle that is more likely to be inflationary than derailed. Cost inflation is nothing new for the hyperscalers, which have been dealing with rising power and infrastructure costs for years. If anything, the additional inflationary pressure from the Iran war may push total capex higher, as more is spent to secure power, chips and resilient infrastructure. The direction of travel of more power, more capacity and more grid is likely to continue, and rather accelerate in regions like Europe and Asia because of this conflict. For now, we view the AI trade as dented but far from derailed.

Date Created

2026/04/10