

Silicon, Not Electrons: Ulanqab and China's Data Centre Bottleneck

Introduction: Capacity Without Compute

Until recently, Ulanqab was known mostly for potatoes. The prefecture sits 1,500 metres up on the Inner Mongolian plateau, about 350km from Beijing, and is cold for most of the year. In six years it has also become China's largest intelligent computing base. By the end of June 2026 it had signed 89 data centre projects, 86 of them for AI, backed by more than RMB500bn of announced investment. Alibaba Cloud [NYSE: BABA] runs close to half the facilities. Huawei and DeepSeek have their own campuses, VNET [NASDAQ: VNET] hosts ByteDance, and ByteDance is negotiating another 5-6GW for early 2028. You might assume the halls are empty, and across much of western China you would be right, with utilisation as low as 20-30% in places. Ulanqab is different. Electricity use at its data centres was up 85% year on year in the first half of 2026, and its tenants are named and under contract. The empty capacity is mostly elsewhere, in the smart computing centres provincial governments bought in 2024 without anyone lined up to use them.

China has 24GW of data centre capacity running and another 50GW in the pipeline, yet its real gap with the US is in compute, not capacity. What holds China back is accelerators, and that limit is now partly self-imposed. Washington keeps Nvidia's [NASDAQ: NVDA] best chips out of reach, and Beijing is rationing the ones it is still allowed to buy so that inference runs on Huawei silicon instead. Capacity figures therefore overstate how much compute China really has, which raises the question of who absorbs the difference. Local government parks have already taken the hit. Alibaba, Tencent [HKEX: 700] and ByteDance are building up exposure as capex outpaces free cash flow. The least exposed are landlords with signed leases, who own the shell but not the chips. The rest of this article looks at the global picture, explains why Ulanqab works, traces the bottleneck to the chip stack and follows the money through the layers that fund it.

Global Overview

On September 22nd, speaking at the UN General Assembly, Donald Trump told world leaders that whoever wins the "super intelligence" race is the one that "wins", and that the US would not surrender its lead to China. The line captures how far Washington has committed itself to the AI race. As in the Cold War, the US is competing with a communist superpower; this time, it is China.

Western readers seldom get a clear view of what is happening inside China, beyond the occasional breakthrough that unsettles Western governments. One of the race's key fronts is the AI infrastructure build-out, and here Europe is already trailing both the US and China. Between the two leaders, the US is ahead on capacity, but its real lead is in compute. US operational capacity is estimated at 56GW against China's 24GW, although China has a further 50GW under construction or announced. Europe's total capacity is expected to reach 13GW by the end of 2026, a 20% increase on 2025.

What sets the players apart is the constraint each one faces. For the US and Europe, the scarce input is electricity: in the UK, grid connection timelines for data centres now extend past eight years, and in the US the average interconnection wait is around five years, against a build time of 18-24 months for an AI data centre. China has abundant, cheap power in its western hubs, though not unlimited: new data centres in Inner Mongolia must source over 80% of their annual electricity from renewables, and energy-consumption quotas may cap capacity over time. Its real bottleneck is silicon. That is why, although China's operational capacity is around 43% of America's, its AI compute is only around a fifth: the US held about three-quarters of global GPU cluster performance as of May 2025, against 15% for China.

All the views expressed are opinions of Bocconi Students Investment Club members and can in no way be associated with Bocconi University. All the financial recommendations offered are for educational purposes only. Bocconi Students Investment Club declines any responsibility for eventual losses you may incur implementing all or part of the ideas contained in this website. The Bocconi Students Investment Club is not authorised to give investment advice. Information, opinions, and estimates contained in this report reflect a judgment at its original date of publication by Bocconi Students Investment Club and are subject to change without notice. The price, value of and income from any of the securities or financial instruments mentioned in this report can fall as well as rise. Bocconi Students Investment Club does not receive compensation and has no business relationship with any mentioned company.

The gap is set to persist, because American capex remains unmatched. Microsoft [NASDAQ: MSFT] alone has lifted its 2026 capex outlook to \$190bn, and Amazon [NASDAQ: AMZN] is expected to exceed \$230bn. By comparison, Alibaba, China's largest cloud provider, has committed RMB380bn (approximately \$53bn) over three years to AI and cloud infrastructure.

China's difficulty is getting enough chips into its data centres, and part of that difficulty is now self-imposed. Beijing has slowed imports of the American chips Washington has allowed it to buy, to protect its technological sovereignty and push demand towards domestic suppliers such as Huawei, a bet that may cost it in the near term but could pay off over time. China has played this game before: over three decades it built near-total control of rare earths, and according to the IEA it now accounts for 59% of global rare earth mining and 84% of the refining of the rare earths used in permanent magnets. In the upper ranks of the Communist Party, the logic of self-reliance carries significant weight.

Why Ulanqab Is China's Northern Virginia

Every AI superpower has its data centre capital. For the US, it is Northern Virginia, whose fittingly named Data Center Alley holds more than 200 data centres in a 30-square-mile stretch. Built up over 30 years as demand grew, it sits atop a global crossroads of fibre-optic networks. For China, it is Ulanqab, a city of under 2m people where nearly all the data centres have gone up within the past decade, driven by policy rather than demand.

In December 2021, under Beijing's "East Data, West Computing" programme, the National Development and Reform Commission (NDRC) made Inner Mongolia one of eight national computing hubs and named Ulanqab's Jining Big Data Industrial Park one of the two starting zones of its flagship cluster. It set targets for occupancy (at least 65%) and energy efficiency (a PUE of 1.2 or below) rather than a quota of buildings. Since then, operators have committed 12.5-15GW to Ulanqab, according to Goldman Sachs and the FT, making it the fastest-growing data centre hub in Asia. Its location helps too: about 350km from Beijing, two hours by high-speed rail and 4ms of latency away.

What makes the city work is power. Inner Mongolia has far more electricity than it can use: with 117GW of wind capacity (nearly four times the UK's) and 130GW of fossil-fuel capacity, largely coal, it probably has the largest electricity oversupply of any region in the world, according to the Lantau Group. As a result, power is cheap. The standard tariff in western Inner Mongolia is about RMB0.52 per kWh, against more than RMB0.60 in many Chinese cities and RMB0.84 in Beijing, and some data centres have negotiated rates of around RMB0.32-0.36.

On a like-for-like basis in 2025, Ulanqab's standard tariff was about 60% below the EU average and less than a third of the UK's, but only about 20% below the US's, and since then the gap has, if anything, widened, with Ulanqab prices falling while US and UK prices have risen. A 20% saving is significant, but America does not lack cheap power; what it lacks is access to the grid.

While most of China is served by the State Grid, western Inner Mongolia, including Ulanqab, is served by the independently operated Mengxi Grid. According to Jefferies, this lets some operators build power plants connected directly to their data centres, avoiding State Grid transmission charges. In Northern Virginia, by contrast, Dominion Energy [NYSE: D] has 51GW of data centre demand contracted and more than 20GW queued. Construction in Ulanqab takes roughly 12-18 months, against 18-24 months in the US, and costs about 20% less than in China's big cities, according to Goldman Sachs. The cold climate, with an average temperature of about 4.3°C, also cuts the heavy cooling load of AI clusters.

All the views expressed are opinions of Bocconi Students Investment Club members and can in no way be associated with Bocconi University. All the financial recommendations offered are for educational purposes only. Bocconi Students Investment Club declines any responsibility for eventual losses you may incur implementing all or part of the ideas contained in this website. The Bocconi Students Investment Club is not authorised to give investment advice. Information, opinions, and estimates contained in this report reflect a judgment at its original date of publication by Bocconi Students Investment Club and are subject to change without notice. The price, value of and income from any of the securities or financial instruments mentioned in this report can fall as well as rise. Bocconi Students Investment Club does not receive compensation and has no business relationship with any mentioned company.

About 68 firms are building or operating in Ulanqab, and they fall into four groups. Self-builders, such as Huawei and Alibaba, build data centres for their own use. Operators build data centres and rent out the capacity: VNET's Bayin campus runs about 221MW, with 700MW planned. Anchor customers rent capacity from operators; ByteDance alone accounts for about a fifth of China's data centre capacity. Power developers build the electricity infrastructure behind the halls, such as Envision, which is building a 2GW power campus beside several data centres.

Unlike much of China's build-out, Ulanqab's capacity is being used. Goldman Sachs estimates that about 1.2GW was in operation in 2025, of which about two-thirds was in use, narrowly above the 65% occupancy target set in the hub's approval. Since then, computing power in operation has more than doubled. VNET reports that 94% of the 585MW it has under construction, mostly in Ulanqab, is already pre-committed.

Ulanqab vs Northern Virginia: both in demand, opposite bottlenecks

	Ulanqab	Northern Virginia
Power price (pre-tax, € cents/kWh, Q2 2025)	5.9 (standard tariff; data centres pay less)	7.6 (US industrial average)
Typical build time	12–18 months	18–24 months (US)
Data centres	89 projects signed (June 2026)	200+ in Data Center Alley
Capacity in operation	~1.2 GW (2025), ~66% utilised	4.5 GW in service; 0.2% vacancy
Pipeline	12.5–15 GW committed	2.4 GW under construction; 51 GW contracted at Dominion (Virginia-wide)
Demand signal	Compute in operation +112% YoY; VNET pipeline 94% pre-committed	467.6 MW net absorption in H1 2026
Main tenants	Alibaba, Huawei, Tencent, ByteDance, DeepSeek, GDS (68 firms in total)	Amazon, Microsoft, Google, Meta, ByteDance
Binding constraint	Chips: Ascend output, HBM supply, Beijing's H200 rationing	Grid power: 20+ GW queued at Dominion; new loads phased in over ~4 years

Source: FT, 36Kr, Goldman Sachs, Science and Technology Daily, VNET Q2 2026, Reuters, CBRE (H1 2026), Synergy Research, Virginia Business, Inner Mongolia Power Group, EIA, ECB; BSIC calc.

Note: Power prices pre-tax, converted at ECB rates of 9 Oct 2026; Ulanqab = Western Inner Mongolia standard tariff; US = national industrial average (Virginia: 8.99 US cents, 2024). Ulanqab operating capacity and utilisation are Goldman Sachs 2025 estimates; Northern Virginia figures are CBRE market data (H1 2026). VNET's 94% pre-committed rate covers its whole 585 MW pipeline, most of it in Ulanqab.

Both Ulanqab and Northern Virginia are in demand, but they are short of different things. Northern Virginia has the customers and the chips but cannot connect power fast enough; Ulanqab has the power and the halls but depends on a chip supply far tighter than America's. And Ulanqab is the exception. Across large parts of China, the data centres were built and the chips never came.

The Real Bottleneck

China has the halls to support large amounts of compute; what it lacks is chips. Since 2022, Washington has steadily tightened China's access to Nvidia hardware. In October 2022, the US restricted exports to China of top AI chips and of the equipment needed to make advanced chips. In October 2023, it blocked Nvidia's A800 and H800, which Nvidia had designed to fall just under the 2022 limits. In December 2024, it restricted sales of high-bandwidth memory (HBM), including memory made abroad with US technology, limiting Huawei's ability to produce its own chips. In April 2025, it imposed a licence requirement on the H20, a chip built to fit under the

All the views expressed are opinions of Bocconi Students Investment Club members and can in no way be associated with Bocconi University. All the financial recommendations offered are for educational purposes only. Bocconi Students Investment Club declines any responsibility for eventual losses you may incur implementing all or part of the ideas contained in this website. The Bocconi Students Investment Club is not authorised to give investment advice. Information, opinions, and estimates contained in this report reflect a judgment at its original date of publication by Bocconi Students Investment Club and are subject to change without notice. The price, value of and income from any of the securities or financial instruments mentioned in this report can fall as well as rise. Bocconi Students Investment Club does not receive compensation and has no business relationship with any mentioned company.

new limits, which made exports to China practically impossible and forced Nvidia to take a \$4.5bn charge on unsold inventory and purchase commitments.

The US is not the only side blocking the flow. In November 2025, China banned foreign chips from new state-funded data centres. In December 2025, Washington authorised H200 sales to China, licensing up to 75,000 chips per approved buyer, but in January 2026 Chinese customs told agents the chips were not permitted to enter, and suppliers paused output. By July 2026, Beijing was reportedly set to approve fewer than 200,000 H200s, less than half of what companies had requested, confining them to training and directing companies to use domestic processors such as Huawei's for inference. Squeezed from both sides, domestic suppliers are expected to capture nearly 80% of China's AI server market in 2026, betting that many weaker domestic chips can stand in for a smaller number of more powerful foreign ones.

That bet is still a long way from paying off. Huawei, China's largest AI chipmaker, produced about 800,000 chips in 2025 and is expected to make 1.5m in 2026, against Nvidia's 4.5m in 2025 and a forecast 5.9m in 2026. Huawei's most powerful chip, the Ascend 950, delivers about half the performance of Nvidia's four-year-old H100. Measured in H100-equivalents, Huawei is set to produce 0.9m this year, while Nvidia, which already ships more advanced chips such as the B300, is set to produce 23m, according to Epoch AI.

However many chips Huawei plans to make, it needs enough HBM to build them. Just before the US restrictions took effect, Huawei stockpiled more than 10m HBM stacks, a reserve expected to run out in 2026. China's main HBM producer, CXMT, is set to make enough for only 250,000-300,000 chips this year, forcing Huawei to use slower conventional memory on half of all chips. The second constraint is manufacturing: SMIC [HKEX: 981], China's leading chipmaker, has advanced capacity of about 60,000 wafers a month, and each chip takes about eight months to make, twice as long as at TSMC [NYSE: TSM]. Memory, not chip design, now sets the ceiling on China's output.

The result is a data centre build-out that has outrun its chips. Across China, many data centres reportedly sit idle or at 20-30% utilisation, and one 2025 estimate put up to 80% of capacity unused. Four factors explain most of it: the lack of chips; remote sites with high latency; facilities built for model training rather than inference; and mixes of chips that do not work together. Beijing now requires data centres to run at a utilisation of at least 60%, yet SMIC's co-CEO has admitted that "what exactly these data centers will do... has not been fully thought through". Ulanqab, at about two-thirds utilisation and 4ms from Beijing, is what the others were supposed to be.

Halls without chips are not just wasted capacity: they are costly to maintain and lose value while producing nothing. Who financed them, and who is left carrying the cost?

The Financing Angle

There is not a single funder behind China's data centres but four different layers, and each owns a different part of the asset. Local governments and their financing vehicles paid for the speculative parks, the state telecoms build where policy points them, the hyperscalers buy the chips, and listed operators such as VNET and GDS Holdings [NASDAQ: GDS] own the shells and lease them out.

The first layer has already taken its loss. In 2024, provincial and city governments, often through local-government financing vehicles (LGFVs), procured "smart computing centres" on subsidy and debt, frequently with no anchor tenant, betting that East Data, West Computing would bring the demand with it. Most of the idle capacity described above sits there, at 20-30% utilisation; Shaanxi, which was not even selected as a hub, built 22 large data centres on subsidies, while Guizhou, the programme's showcase, has yet to turn a profit on its data centres and ranks only


All the views expressed are opinions of Bocconi Students Investment Club members and can in no way be associated with Bocconi University. All the financial recommendations offered are for educational purposes only. Bocconi Students Investment Club declines any responsibility for eventual losses you may incur implementing all or part of the ideas contained in this website. The Bocconi Students Investment Club is not authorised to give investment advice. Information, opinions, and estimates contained in this report reflect a judgment at its original date of publication by Bocconi Students Investment Club and are subject to change without notice. The price, value of and income from any of the securities or financial instruments mentioned in this report can fall as well as rise. Bocconi Students Investment Club does not receive compensation and has no business relationship with any mentioned company.

22nd by GDP. Beijing's answer has been to support demand rather than write off supply: computing power vouchers that subsidise smaller firms buying compute, a 60% utilisation floor, a ban on new large facilities in cities running below 50%, and more than 100 projects cancelled in the 18 months to mid-2025. In effect, the losses are socialised. The rules also protect incumbents: by blocking new supply wherever utilisation is low, they push demand towards hubs that are already full, and Ulanqab is the main beneficiary.

The second layer is the state. China Mobile [HKEX: 941] and China Telecom [HKEX: 728] are cutting total capex while raising spending on computing sharply, and both report high utilisation of what they have built. Their returns are set by policy rather than by the market, and state ownership cushions them in a way no private landlord enjoys, which makes them the natural buyer of last resort if the local parks have to be consolidated. Their utilisation figures, however, are self-reported and not comparable with an operator's utilised megawatts, so they say more about policy than about demand.

The third layer is where the risk is building. Alibaba spent RMB67.7bn of capex in the June quarter against free cash flow of negative RMB44.7bn, and Tencent RMB52.8bn against negative RMB13.8bn, on top of the RMB380bn Alibaba has already committed. ByteDance, unlisted but holding about a fifth of China's data centre capacity, is the largest single tenant and is negotiating a further 5-6GW in Ulanqab. Unlike the local parks, the hyperscalers' halls are not empty, so their risk is not vacancy but obsolescence and returns. They own the chips, and with Beijing steering inference onto weaker domestic silicon, the return on each accelerator is uncertain. With free cash flow negative, the marginal funder of the build-out is no longer operating cash but the debt and equity markets.

Spending ahead of cash: capex vs free cash flow, June quarter 2026 (RMBbn)



Source: Alibaba Group and Tencent Holdings quarterly results, SCMP, BSIC calc.

Note: June quarter 2026 (Alibaba fiscal Q1 2027). Negative free cash flow shown in brackets.

What separates the layers is what each one owns. A data centre shell and its grid connection last 20 years or more; the accelerators inside it are obsolete in about five. "Stranded" therefore means different things at each layer. An

All the views expressed are opinions of Bocconi Students Investment Club members and can in no way be associated with Bocconi University. All the financial recommendations offered are for educational purposes only. Bocconi Students Investment Club declines any responsibility for eventual losses you may incur implementing all or part of the ideas contained in this website. The Bocconi Students Investment Club is not authorised to give investment advice. Information, opinions, and estimates contained in this report reflect a judgment at its original date of publication by Bocconi Students Investment Club and are subject to change without notice. The price, value of and income from any of the securities or financial instruments mentioned in this report can fall as well as rise. Bocconi Students Investment Club does not receive compensation and has no business relationship with any mentioned company.

unleased shell is a vacancy problem for its owner, while a leased hall full of idle chips is a write-down for the tenant, and the landlord keeps collecting rent. Lease terms, meaning tenor, take-or-pay clauses and early-termination rights, decide who actually bears the loss, and chips that are weaker to begin with lose their value well before the lease runs out.

The fourth layer is where the halls meet the capital markets. VNET is the clearest case: it has moved from retail colocation to wholesale campuses built for a few large tenants, with Ulanqab and ByteDance at the centre. At the end of June it had 744MW utilised out of 1,007MW in service, and 94% of the 585MW under construction was already pre-committed. That concentration is both its strength and its main risk. The build-out is financed largely with debt, while new state-linked shareholders such as CATL [SZSE: 300750] and Shandong Hi-Speed have brought in fresh equity. GDS has shown the exit route: its C-REIT priced at a 5.2% yield, equivalent to roughly 19 times distributable cash flow, with 166 times institutional cover. The market is sorting the assets. Stabilised, leased halls can be recycled at tight yields and the cash redeployed; unleased or under-chipped ones cannot.

Who carries the empty halls: China's data centre capital stack

	What they own	Utilisation	Main risk	Who backstops
Local governments and LGFVs	Subsidised smart computing parks, often with no anchor tenant	20-30% in many parks	Stranded shells; loss already crystallised	Beijing: vouchers, utilisation floor, consolidation
State telecoms	Policy-directed computing capacity	High (self-reported)	Low, policy-set returns	State ownership
Hyperscalers	The chips, plus self-built halls	Halls in use; return on chips uncertain	Obsolescence; capex ahead of free cash flow	Debt and equity markets
Colocation operators	Long-lived shells and grid connections, leased out	VNET: 744MW of 1,007MW in service	Tenant concentration; refinancing	Signed leases; C-REIT investors

Source: ASPI, Reuters, Notebookcheck, Caixin, TelecomLead, VNET Q2 2026, GDS, BSIC analysis.

Note: layers ordered by who bears the loss first. Telecom utilisation is self-reported and not comparable with VNET's utilised megawatts.

Put together, the risk is bifurcated. It has already crystallised in the local parks, where the state is quietly absorbing it, and it is growing on hyperscaler balance sheets and with whoever refinances them. Landlords with signed leases sit at the back of the queue, and how quickly the queue moves depends on whether the silicon arrives.

Outlook: Who Ends Up Paying

China now has roughly 74GW of shells either built or planned, and most of them are still waiting on chips. How that gap closes matters less for China's AI ambitions than for the question of who ends up paying for it, and we see four ways it could play out. In each of them we use Ulanqab as the yardstick, since it is the hub with the least to lose: if the plateau struggles, it is hard to see any other hub doing better.

All the views expressed are opinions of Bocconi Students Investment Club members and can in no way be associated with Bocconi University. All the financial recommendations offered are for educational purposes only. Bocconi Students Investment Club declines any responsibility for eventual losses you may incur implementing all or part of the ideas contained in this website. The Bocconi Students Investment Club is not authorised to give investment advice. Information, opinions, and estimates contained in this report reflect a judgment at its original date of publication by Bocconi Students Investment Club and are subject to change without notice. The price, value of and income from any of the securities or financial instruments mentioned in this report can fall as well as rise. Bocconi Students Investment Club does not receive compensation and has no business relationship with any mentioned company.

The optimistic case is simply the plan Beijing drew up in 2021 actually working. Huawei is on track to ship about 1.5m Ascend chips this year, close to twice what it managed in 2025. Should those volumes materialise, and should SMIC keep getting acceptable yields out of its enhanced 7nm node, domestic accelerators would fill the western halls faster than any import quota could. Utilisation in the local government parks would drift up towards Beijing's 60% floor, and the C-REIT market, which so far has only taken stabilised colocation, could start absorbing AI campuses too. GDS's first vehicle, priced at a 5.2% yield with 166 times institutional cover, suggests the appetite is there. For this to be the world we are in, three things have to hold. Ascend shipments need to come in on plan, Huawei needs HBM from somewhere other than its shrinking stockpile, and VNET's utilised capacity (744MW out of 1,007MW in service at the end of June) needs to keep rising as new halls come online rather than being diluted by them.

Our base case is less flattering. H200 approvals stay small, Huawei remains three to four years behind Nvidia, and training capacity stays scarce. Beijing is in fact steering things this way on purpose, with imported chips reserved for training and domestic ones for inference. Under this scenario the western campuses end up as inference and storage parks, handling workloads that can live with 4ms of latency to Beijing and a good deal more to Shenzhen. The economics are thinner. Inference on older silicon is sold by the token, and the provincial compute vouchers subsidise the buyer, not the landlord. Ulanqab still comes out ahead. Its negotiated power costs RMB0.32-0.36 per kWh and its modular halls go up fast, which makes it the cheapest place in China to run this load. It should stay full while parks without a cost or connectivity advantage empty out. We would know we are on this path if H200 quotas keep getting renewed each quarter in small batches, if hyperscalers start saying openly that their western capacity is for inference, and if token prices carry on falling.

The bear case comes down to enforcement. On paper Beijing has a 60% utilisation floor and a ban on new large facilities in cities running below 50%. Whether Shaanxi or Guizhou will actually apply those rules against their own financing vehicles is another matter. If they don't, the local government parks get folded into the state telecoms or written down, and the pressure travels up the stack. Alibaba and Tencent already spend more on capex than they bring in as free cash flow, so the first guidance cut made to protect cash would be the moment things turn. It is also the only route by which the problem reaches VNET and GDS, whose pre-commitments (94% of VNET's pipeline) rest on a handful of wholesale tenants. The warning signs would be utilisation rules nobody enforces, capex guidance being trimmed and ByteDance's gigawatt talks going quiet.

There is also a more speculative outcome. The Ministry of Industry and Information Technology wants surplus compute pooled into a nationally scheduled cloud by 2028, mainly to push domestic utilisation up. Compute that cannot find a buyer at home, though, will sooner or later find one abroad if the price is low enough. Chinese labs already sell tokens far below Western rates. Put idle western capacity behind open-weight models and offshore APIs, and the marginal price of inference worldwide could end up being set in Guizhou instead of Northern Virginia. We think this is the least likely of the four, for the simple reason that it needs chips the halls don't have yet.

The indicators should move roughly in sequence. SMIC's output and Ascend shipments against the 1.5m estimate come first, then domestic HBM, then the size and conditions of each H200 approval round along with any change in US licensing. Together these tell us which scenario China is actually in. Who pays is a separate set of numbers: hyperscaler capex against free cash flow, computing capex and utilisation at China Mobile and China Telecom, utilised against in-service megawatts at VNET and GDS, and C-REIT issuance and yields.

All the views expressed are opinions of Bocconi Students Investment Club members and can in no way be associated with Bocconi University. All the financial recommendations offered are for educational purposes only. Bocconi Students Investment Club declines any responsibility for eventual losses you may incur implementing all or part of the ideas contained in this website. The Bocconi Students Investment Club is not authorised to give investment advice. Information, opinions, and estimates contained in this report reflect a judgment at its original date of publication by Bocconi Students Investment Club and are subject to change without notice. The price, value of and income from any of the securities or financial instruments mentioned in this report can fall as well as rise. Bocconi Students Investment Club does not receive compensation and has no business relationship with any mentioned company.

Conclusion: The Landlord Pays Last

Ulanqab shows that China can put up a hyperscale campus in just over a year, on power far cheaper than Beijing's, and fill it, provided the tenant turns up with chips. Nationally, though, the build-out has outrun the silicon, and Beijing is happy to leave it that way if that is what it takes to force a domestic chip industry into being. Put simply, it is accepting slower AI now so as not to depend on Washington later. Local governments are already paying for that decision, stuck with parks running at a fraction of their design utilisation. The hyperscalers, whose free cash flow went negative to pay for all this capex, are next in line, followed by whoever refinances them. Landlords holding signed long-term leases pay last. So the potato town on the plateau ends up as the safest asset here, which is good news for Ulanqab and not much comfort for anywhere else. If investors take one thing from this, it should be to read the lease before worrying about where the building is, because what can go wrong is the chips inside it.

References

- [1] Caixin Global, "In Depth: How a Windswept City in North China Became an AI Powerhouse", 2026. [Link](#)
- [2] ASPI The Strategist, "Abundant Electricity Isn't Enough: China's Overbuilt AI Computing Power Is Underused", 2026. [Link](#)
- [3] Financial Times, "China Races to Build Data Centres in Bid for AI Supremacy" (via Dealroom), 2026. [Link](#)
- [4] TrendForce, "China Reportedly to Allow NVIDIA H200 Imports, but Approvals May Be Capped Below 200K Chips", 2026. [Link](#)
- [5] Bloomberg, "China Is Building an AI Future Far from Its Biggest Cities" (via The Star), 2026. [Link](#)
- [6] Financial Times, Ulanqab data centre hub, 2026. [Link](#)
- [7] VNET Group, "Unaudited Second Quarter 2026 Financial Results", 2026. [Link](#)
- [8] Epoch AI, "Will Huawei Catch Up to Nvidia by 2030?", 2026. [Link](#)

TAGS: China, Data Centres, Artificial Intelligence, Ulanqab, Export Controls, Semiconductors, Real Estate, Infrastructure, VNET, GDS, C-REIT

All the views expressed are opinions of Bocconi Students Investment Club members and can in no way be associated with Bocconi University. All the financial recommendations offered are for educational purposes only. Bocconi Students Investment Club declines any responsibility for eventual losses you may incur implementing all or part of the ideas contained in this website. The Bocconi Students Investment Club is not authorised to give investment advice. Information, opinions, and estimates contained in this report reflect a judgment at its original date of publication by Bocconi Students Investment Club and are subject to change without notice. The price, value of and income from any of the securities or financial instruments mentioned in this report can fall as well as rise. Bocconi Students Investment Club does not receive compensation and has no business relationship with any mentioned company.