



Capvion Insights

Strategic perspectives on capital, growth and transformation

The AI Gold Rush Has the Wrong Bottleneck ***Capacity, Clarity, Consumption: The Three C's of AI Adoption***

The industry solved the capacity problem. The harder problem, and the durable value, sit further up the stack

Sander Barens | August 26, 2026
with associates

Capvion Insights provides strategic perspectives on value creation, transformation, and leadership in technology and telecommunications businesses.

Each publication examines how strategy, capital, and execution interact to shape enterprise value in complex markets.

Drawing on extensive board-level, advisory, and operational experience, Capvion and its associates combine strategic thinking with real-world execution across numerous M&A transactions, platform integrations, and growth transformations in the technology and telecom sectors.

Our insights are designed to support founders, executives, and investors in navigating complex growth, restructuring, and transaction environments.

Each publication reflects Capvion's commitment to clarity, independence, and practical relevance.

Executive Insight

Everyone is talking about GPUs. They are looking in the wrong place.

For two years the industry has answered a single question: how do we build more capacity? It was the right question, and the answer was extraordinary. But in recent conversations with operators, infrastructure providers, investors and enterprise leaders across Europe and the Middle East, almost nobody asks any more whether AI infrastructure will exist. They ask something harder. Which of these hundreds of options should we choose, and once we have chosen, why is it still so difficult to actually use?

That is not an infrastructure problem. It is a consumption problem, and it is what now stands between AI and large-scale enterprise adoption.

We see enterprise AI moving through three phases, and it is worth naming them clearly. Capacity was the challenge of the last two years. Clarity is the challenge of today. Consumption is the challenge that will decide where the value ends up.

Every infrastructure market in living memory has run that same sequence, and in every one of them the value eventually moved away from the companies that built the capacity. There is no obvious reason why AI should be the exception.

The only real question is who moves first.

Key Takeaways

- **The industry solved the right problem, two years ago.**
Capacity was the binding constraint of 2024 and 2025. It is no longer the constraint that determines outcomes.
- **Clarity is today's bottleneck.**
Models, architectures, clouds, sovereign initiatives and regulations are multiplying faster than any enterprise can evaluate them. The shortage is not compute. It is judgement.
- **Consumption is the bottleneck that follows.**
Choosing a provider is where the work begins, not where it ends. Contracts, APIs, governance and compliance all have to be solved again for every one of them.
- **The Three C's are sequential, not parallel.**
Capacity creates capability. Clarity creates confidence. Consumption creates adoption. Skipping a phase does not accelerate the market, it stalls it.
- **Telecom and cloud already ran this experiment.**
Both built capacity first, standardised access second, and then watched consumption platforms capture a disproportionate share of the value.
- **The winners will not necessarily own the infrastructure.**
They will own the consumption layer, and it is far less capital intensive to build.

Context

Everyone is talking about GPUs. They are looking in the wrong place.

For the past two years the industry has answered demand the only way infrastructure industries know how. More data centres. More GPUs. More sovereign AI programmes. Given the question being asked, that was exactly the right response, and it was executed at remarkable speed.

Over the past few months I have been having a different kind of conversation. Operators, infrastructure providers, investors and enterprise leaders, across Europe and the Middle East. Interestingly, almost nobody is asking whether AI infrastructure will exist. Instead I keep hearing two very different frustrations, and neither of them is about capacity.

The first is too much choice. The second is too little connectivity. One is a problem of decision-making, the other a problem of friction, and together they explain why so much installed capacity is still waiting for serious enterprise workloads to arrive.

Analysis

The market is still measuring progress in gigawatts and GPU counts, while the constraint that actually governs adoption has quietly moved somewhere else. It helps to hold the whole picture in one frame. We call it the Three C's of AI Adoption.




Figure 1. The Three C's of AI Adoption

Capacity: The Challenge We Have Largely Solved

Capacity was the honest constraint of the last two years. Demand exploded, supply did not exist, and the industry responded the way every infrastructure industry always has. It built. Enormous amounts of capital went into data centres, accelerators, power and sovereign programmes, and the response worked.

That phase is not over, but it is no longer where the interesting problems are. Capacity is becoming a question of financing, energy and location rather than a question of strategy. Necessary, capital intensive, and increasingly difficult to differentiate.

Clarity: Today's Real Shortage

Every few months brings another foundation model, another GPU architecture, another cloud provider, another sovereign AI initiative and another wave of regulation. The pace of innovation is extraordinary, but it comes at a cost that rarely appears in any business case. As the number of options grows, making the right decision becomes harder, slower and more expensive.

This is what I hear most often from boards. Not "can we get compute", but "how do we know this is the right choice, and how much does it cost us if it is not". Ironically, the industry is no longer facing a shortage of compute. It is facing a shortage of clarity.

Consumption: The Challenge That Decides the Winners

Even after an organisation has decided which infrastructure and models to use, the journey has only just begun. Every provider has its own commercial model, APIs, contracts, governance framework and ordering process. Comparing providers is still largely a manual exercise. Moving workloads between them is complicated. Proving compliance often means starting all over again.

The capacity exists. Consuming it efficiently is the real challenge, and consumption is a much broader problem than procurement. It covers the ability to discover, compare, procure, connect, govern, switch and operate AI services across providers, without rebuilding the plumbing every time.

These are no longer infrastructure bottlenecks. They are consumption bottlenecks, and they are the reason so many AI programmes look impressive in a pilot and disappointing at scale.

I Have Seen This Movie Before

Twenty years ago, global telecom looked remarkably similar. Every operator had different products, different contracts, different ordering processes and different interfaces. Buying connectivity across multiple providers was slow, expensive and highly manual. It took years of industry collaboration and standardisation to turn connectivity into something that could be quoted, ordered and provisioned programmatically rather than through spreadsheets, emails and endless negotiations.

Cloud computing followed exactly the same path. AWS did not change the industry because it built the largest data centres. It changed the industry because it removed friction. Developers could suddenly provision infrastructure in minutes instead of spending months negotiating contracts and waiting for hardware to arrive. The capacity mattered. The consumption model was what made it a revolution.

The Pattern Behind Every Infrastructure Revolution

Looking back, the sequence is almost mechanical. First we build capacity. Then we standardise access. Then platforms emerge that make the underlying infrastructure easy to discover, compare, consume and manage. And then, reliably, the economics tilt towards those platforms.




Figure 2. Where AI sits in the evolution every infrastructure market has followed

“Capacity creates capability. Clarity creates confidence. Consumption creates adoption.”

AI is now somewhere between the first and second stage. The capacity is being built, standardisation has barely started, and the consumption layer is largely absent. That gap is the opportunity.

Where the Durable Value Goes

If history repeats itself, much of today's GPU capacity becomes the commodity layer. Essential, yes. Capital intensive, certainly. But hard to defend and harder to differentiate once supply catches up with demand.

The durable value moves higher up the stack, to whoever makes a fragmented ecosystem simple to navigate, simple to connect to, simple to govern and simple to leave. Telecom learned this the hard way. Operators built the networks and platforms captured a disproportionate share of the value. Which raises an uncomfortable question for an industry that has already lived through this once: the telcos know this pattern better than anyone, so why are they not the ones building the AI consumption layer?

Capvion Perspective

In our experience advising companies, investors and operators across technology and telecommunications markets, the organisations that get value out of AI are the ones that treat the Three C's as three separate decisions rather than one procurement exercise. Three habits separate them from the rest.

Separate the Capacity Decision from the Consumption Decision

Most AI business cases still collapse these into a single question. They are not the same. Where the compute sits is a capacity decision. How easily you can compare, connect, govern and switch is a consumption decision, and it is the one that determines whether the capacity ever gets used well. Approve them separately, and hold different people accountable for each.

Buy for Switching, Not Only for Scale

The best commercial terms available today are worth less than the ability to move tomorrow. Negotiate for portability, exit and interoperability with the same seriousness applied to price per GPU hour. Lock-in is rarely priced at signature. It is always paid later, and usually at the worst possible moment.

For Investors, Look One Layer Above the Estate

Capital is still flowing overwhelmingly into the layer that is commoditising. The scarcer asset is the layer that removes complexity, and it requires a fraction of the capital to build. When infrastructure matures, the multiple migrates upward. It has done so in every previous cycle.

Conclusion

Every infrastructure revolution ends the same way. Capacity gets built. Access gets standardised. And then a small number of companies emerge whose entire purpose is to make the thing simple to consume. They rarely own the underlying assets. They almost always own the customer relationship, the workflow, and the multiple.

I do not believe AI will be the exception. My prediction is that within three years the breakout companies in AI will not be the ones with the largest GPU estates. They will be the ones that made AI as easy to consume as cloud became fifteen years ago: a universal way to discover, compare, procure, connect, govern, switch and operate AI services, regardless of who owns the infrastructure underneath.

Capacity is being solved with capital. Clarity is being solved slowly, and painfully, inside every enterprise on its own. Consumption has barely been started.

The winners will not necessarily own the infrastructure. They will own the consumption layer.

Which leaves one question worth sitting with.

Who is going to build the AI consumption layer?

About Capvion

Capvion is a strategy and capital advisory firm that works with leadership teams, investors, and entrepreneurs on complex growth, transformation, and value creation challenges.

The firm combines strategic insight with hands-on operational and transaction experience across technology and telecommunications markets, advising companies and investors on corporate strategy, mergers and acquisitions, transaction readiness, and enterprise value optimisation.

Capvion operates as a senior-led advisory boutique, working closely with a network of experienced associates and industry specialists. This model allows the firm to bring in the right expertise at the right moment, combining strategic clarity with deep domain knowledge across transactions, integrations, and operational transformation.

Through its advisory work and publications such as Capvion Insights, the firm shares practical perspectives on how strategy, execution, and capital interact to shape enterprise outcomes.

For more information visit www.capvion.com or contact insights@capvion.com