
PERSPECTIVE

Battle for the Edge

Whitepaper

Bart de Graaff · BdG Advisory · October 2026

The edge is not a place to put servers. It is the point where data is admitted, identified, governed, computed on and billed. Hyperscalers are extending their control plane into it with far deeper pockets, telcos own it physically but keep drifting towards being landlords, and neither side has won. While the contest lasts, buyers hold bargaining power that disappears once it is settled.

Telcos own the last mile. Most are still selling the cable.

The edge is not a place to put servers. It is a position of control, and two industries are fighting for it.

Telcos and cloud providers want the same thing: control of data access at the edge. That is the point where data from devices, sites and people enters a network, is identified, meets policy and inference, and is billed. Hyperscalers want that point to be one more extension of their cloud control plane. Telcos own the physical edge, but a business model built on turning their capital spending into customers' monthly subscriptions keeps pulling them towards selling capacity instead of services. The outcome decides who governs enterprise intelligence at the edge, and what it costs an enterprise to leave.

This paper maps both sides' positions, explains why each acts rationally, sets out what has happened so far, and shows how enterprises should choose infrastructure partners while the battle is still open.

The prize is the control point, not the compute

In 1984 John Gage of Sun Microsystems coined the phrase "The Network is the Computer". Sun let the trademark lapse after its acquisition, and Cloudflare later registered it to describe its own strategy. For most of the forty years since, the phrase was a slogan. Now it is an architecture.

Inference is moving into the network itself. Cloudflare said in 2024 that its serverless inference platform had GPUs in more than 180 cities. In October 2025 NVIDIA agreed to invest \$1 billion in Nokia, as part of a partnership to put accelerated computing into the radio access network, so that the equipment at the base of a mobile mast becomes a computing platform as well as a radio. The network is becoming the place where decisions are computed, not just the road the data travels on.

That changes what is being fought over. Compute is increasingly fungible. A GPU in one building does much the same work as a GPU in another. What is not fungible is the control point that sits in front of it. Five functions make up that control point:

- **Access:** who and what is allowed to connect at the location.
- **Identity:** how a device, user or transaction proves what it is.
- **Policy:** what that identity is allowed to do, and with which data.
- **Inference:** where the decision runs and on whose stack.
- **Billing:** who sends the invoice and therefore owns the customer relationship.

Whoever holds those functions sees the data flows, shapes the enterprise's architecture around their own tools, and sets the cost of switching. Margins in infrastructure come from control, not from the infrastructure.

Consider how this plays out at a typical multi-site business. A national retailer puts loss-prevention inference into its stores. The carrier supplies the store network and a 5G backup link. A cloud provider supplies an on-site appliance that runs the models and reports to its console. The camera maker ships its own analytics on the camera. The integrator wires it together.

Now ask who holds each control point. The carrier controls access and the SIM identities. The cloud provider controls the inference stack, the model registry and the update channel. The camera maker controls what its own analytics decide. The integrator holds the configuration nobody else has documented. Billing arrives from four directions.

Two years later the retailer wants to change model provider. It discovers that the appliance runs only one cloud's stack, the device identities live in that cloud's console, and the camera analytics overlap with the new models in ways nobody can untangle quickly. Nothing in the original purchase was unreasonable. The control points were simply allocated by whoever arrived first, and the cost of reallocating them is now measured in store visits.

That is what the battle for the edge looks like from the buyer's side. Not a dramatic contest between giants, but a slow allocation of control points, site by site, contract by contract.

The prize is large. IDC expects service providers' infrastructure investment for edge services, including multi-access edge computing, content delivery networks and virtual network functions, to reach almost \$100 billion by 2028. That is the capital cost of the battlefield. The control points are what the capital is meant to buy.

For the enterprise, the stakes are higher than the price of a service. The companion paper, *The Intelligent Edge Changes Who Decides*, argues that when inference runs on site, decisions move to the point of work. This paper looks one layer down. Whoever controls access at the edge holds a claim on those decisions: on the data behind them, the channel that updates them, and the terms on which the enterprise can change its mind.

Hyperscalers extend their control plane outwards

The hyperscaler strategy is coherent and well funded. Their product is not data centres. It is a control plane: one set of interfaces, one identity system, one bill and one developer toolchain, from the cloud region to the device. Every location brought inside that control plane raises the switching cost for everything else.

They do not need to own the last mile to win. They need to run inside it, and they need the developer.

The developer is the kingmaker at the edge. Whoever's interfaces the application is written against owns the location in practice, whatever the network diagram says. Telcos expose their networks through their own portals and standards bodies. Developers write against cloud software development kits they already use every day. Each edge application built that way pulls one more site into a control plane the telco does not hold.

AWS Wavelength is the clearest example. It embeds AWS compute and storage inside operators' 5G networks, so that application traffic reaches AWS servers without leaving the mobile network. AWS currently lists Wavelength in 31 cities across five operators, twenty of them on Verizon's network. The operator provides the site, the power, the radio and the local licence. The developer sees AWS.

Google has gone a step further with sovereignty. Under its Managed GDC Provider initiative, announced in 2024, selected partners deploy and operate Google Distributed Cloud as a fully managed service for customers with strict sovereignty needs. The first partners included a telco group's IT services arm, T-Systems, and Clarence, a joint venture between a Luxembourg state-owned company and a telco's Luxembourg subsidiary. The telco side runs the facility and carries the local trust. Google supplies the stack.

Microsoft went into the network core. In 2021 AT&T agreed to move its 5G mobile network to Microsoft's Azure for Operators cloud, and Microsoft acquired AT&T's Network Cloud technology and engineering team. One of the largest operators in the world handed the software platform under its mobile core to a cloud provider.

The hyperscalers are also building network themselves. Meta announced Project Waterworth in 2025, a subsea cable system of more than 50,000 kilometres reaching five continents, built to carry AI-era traffic. The companies that once bought capacity from carriers now lay their own.

And they can outspend everyone. In July 2026 Alphabet raised its 2026 capital expenditure guidance to between \$195 billion and \$205 billion. Verizon, one of the largest operators in the world, expects 2026 capital expenditure of \$16.0 billion to \$16.5 billion. One hyperscaler's annual capital budget is more than ten times that of a major national carrier.

The mechanism is rational on both sides. For the hyperscaler, every edge location run on its stack extends a control plane it already monetises. For the operator, hosting the hyperscaler brings utilisation, revenue and a developer audience it struggled to reach alone. The result is a division of roles: the telco as landlord, providing space, power, fibre and licence; the

hyperscaler as ruler, holding the interface, the identity, the bill and the customer.

The telco trap: every rational move goes down the stack

Telcos built a business on converting capital into subscriptions. They raised money, built networks, and sold access back to customers as a monthly operating expense. For decades that was a fine business, because connectivity was scarce and switching was hard.

That model is now under strain. Boston Consulting Group's 2025 analysis found that the median telco's return on invested capital fell to 6.7% in 2024, below a weighted average cost of capital of 7.1%, and that five-year annualised shareholder returns for the telcos it studied were about 4%, against about 12% for the S&P Global 1200. On average, the industry is not earning its cost of capital.

When a business does not earn its cost of capital, management does rational things. Each of them moves the telco further down the stack.

It separates and monetises infrastructure. BCG notes that Telecom Italia completed the first spin-off of a former monopoly fixed-line network in Europe in 2024, cutting its debt ratio sharply. Towers, fibre and data centres become assets to sell or share, which turns the operator's own capital spending into someone else's balance sheet and its own rent.

It disciplines capital spending. Verizon's guidance above is not a sign of weakness. It is what investors reward when returns are below the cost of capital.

It partners with the hyperscalers for what it cannot build. In 2024 Vodafone and Microsoft signed a ten-year partnership in which Vodafone committed \$1.5 billion to cloud and AI services developed with Microsoft, including migrating its virtual data centres to Azure. That is a sensible procurement decision. It also places the operator's own IT and its enterprise services on the other side's platform.

And it uses new technology to sell more of the old product. When Verizon launched AI Connect in January 2025, the first users it named were Google Cloud and Meta, using capacity on its network for AI workloads. Verizon's business chief said the company had well over \$1 billion in its sales funnel using services it already had. That is good business. It is also the definition of a pipe: the AI boom arrives, and the telco's first win is selling fibre to the hyperscalers.

The attempts to move up the stack have been expensive. Network application programming interfaces (APIs), which let developers call network functions such as location checks or quality-of-service requests directly, were meant to turn the network into a platform. Ericsson paid \$6.2 billion for Vonage in 2022 to build that platform, then wrote down about \$3 billion in 2023 and a further \$1.1 billion in 2024, citing lower anticipated growth. In 2025 Ericsson and twelve large operators finalised Aduna, a 50:50 joint venture to aggregate and sell network APIs globally. It is the industry's collective bet that operators can sell a platform together that none of them could sell alone.

The deeper problem is incentive, not technology. Telcos have the assets: last-mile ownership, physical sites, spectrum, local licences and operations people in every region. What they lack is a reason, inside the current business, to cannibalise the margin that keeps them alive. A business-unit head paid on connectivity revenue and margin will rationally sell more connectivity. A partnership with a hyperscaler moves slowly because neither side wants to be the junior partner, and the telco usually is.

Consider a pattern familiar to anyone who has bought edge services from an operator. A manufacturer asks for a site-level service: private wireless, local compute for vision models, security and someone accountable when the line stops. The account team is measured on contract value for connectivity and managed network services. The compute and model layers are easier to source from a hyperscaler partner than to build. The proposal that goes out is a connectivity contract with a cloud reference architecture attached.

Everyone in that chain has acted sensibly. The account team sold what it is paid to sell. The product team avoided building what it could not build at scale. Finance protected margin. The customer got a workable answer. And the control plane for the manufacturer's edge went to the hyperscaler, while the operator kept the part of the deal with the lowest growth.

So the trap closes one sensible decision at a time. No board ever votes to become a landlord. It votes for capital discipline, an infrastructure sale, a cloud partnership and a capacity deal. Add them up and the operator holds the asset while someone else holds the control point.

Neither side has won the edge

It would be easy to conclude that the hyperscalers have already won. The evidence says otherwise. Both sides have retreated from positions they once claimed, and both are making new moves.

The hyperscalers pulled back from private networks. AWS ended support for AWS Private 5G in May 2025, and Microsoft disclosed the retirement of Azure Private 5G Core, due in September 2025. Analysys Mason's assessment was blunt: "The private networks market has not developed quickly enough for the hyperscalers' public-cloud-centric vision to materialise." In private networks the hyperscalers now play secondary roles to operators and their partners.

The telco-hosted cloud edge has also contracted in places. In 2021 AWS announced Wavelength Zones on Vodafone's network in London, Berlin, Munich and Dortmund. None of those cities appears on AWS's current Wavelength location list. Embedding a hyperscaler inside a telco network turned out to be harder to sell than to announce.

Meanwhile, some operators are moving up the stack rather than down. In February 2026 Deutsche Telekom opened an Industrial AI Cloud in Munich with nearly 10,000 NVIDIA Blackwell GPUs, run under German data protection and security requirements, with SAP providing the business platform layer. The same month, its IT arm said its sovereign public cloud already offered 80% of the core functionality of leading cloud providers, with 100%

targeted by the end of 2026. Whether that claim holds is for customers to test. The strategic intent is clear: a telco offering itself as the sovereign alternative, not as the landlord.

The radio network is becoming contested ground too. The NVIDIA–Nokia deal is aimed at AI-native mobile networks, with T-Mobile US planning field trials from 2026. If the base station becomes a computer, the operator owns computing sites in every neighbourhood. Whether it also owns the control plane on those computers is the open question.

The mechanism behind this standoff is that the edge rewards two different kinds of asset. It rewards physical presence, local regulation and operational accountability at the location, which are telco assets. It also rewards a developer-facing control plane with global reach, which is a hyperscaler asset. Neither side has both. Whoever acquires the other's asset on credible terms wins the location.

Sovereignty is where the two strategies collide most directly. Governments and regulated industries increasingly want data, compute and operations under local jurisdiction and local control. That plays to operators, which are already licensed, regulated and physically present in-country, and it explains why a large European operator now markets itself as the sovereign alternative. The hyperscalers' answer is to run their stack through local partners, so that the facility and the operators are local while the control plane stays theirs.

Both models will sell. They allocate control differently. In one, the local operator holds the control plane and buys technology. In the other, the hyperscaler holds the control plane and rents local trust. An enterprise buying "sovereign edge" should know which of the two it is buying, because the label on the box is the same.

For telcos, that leaves two exits. They can become business service providers: accountable for an outcome at a location, combining secure connectivity, edge compute, sovereignty and operations into one service the enterprise buys as such. Or they can become landlords: valuable, regulated, low-growth hosts for someone else's control plane. Some will make the first move. The incentives described above suggest most will drift into the second.

The business service provider is a different company, not a new product. It treats connectivity as a consumable inside a larger service rather than as the thing being sold. It embeds security at the device and transaction level instead of selling it as a separate line. It lets workloads scale across more than one cloud from the operator's own edge, so the customer is not forced into a single control plane. It is paid on the outcome at the location, with people on the ground who can act. And it gives the enterprise, not the supplier, the keys to identity, data and updates.

Buyers can tell which exit an operator is taking without inside knowledge. Look at what its enterprise sales teams are paid on. Look at whether its edge offers run on its own control plane or simply host someone else's. Look at whether it will put portability and data access in a contract. An operator that cannot answer those questions is a landlord, whatever its marketing says.

The strongest objection: it is only infrastructure

The best counter-argument is the one a good procurement director makes. Why should a buyer care who wins?

Infrastructure is becoming a commodity. Connectivity is sold on price and service levels. Compute is compute. The rational buyer multi-sources, buys the lowest unit cost, writes strong service-level agreements and lets suppliers fight over margin. The battle is their problem. Competition between telcos and hyperscalers only pushes prices down.

Much of that is right. For connectivity and raw compute, competition helps buyers, and enterprises should take the benefit. A buyer who pays a premium for a brand or a partnership logo is wasting money.

But the objection confuses the commodity with the control point.

First, control points carry switching costs, and regulators have found them material. In 2025 the UK Competition and Markets Authority concluded its cloud market investigation and found that Amazon and Microsoft hold positions of significant market power, with data egress fees and interoperability barriers restricting switching and multi-cloud. The EU Data Act goes further: from 12 January 2027, cloud providers may no longer charge customers for switching at all. Lawmakers do not legislate against lock-in that does not exist.

Second, at the edge the switching costs are physical. Leaving a cloud region means moving data. Leaving an edge control plane means replacing hardware on hundreds of sites, re-integrating network functions and renegotiating site agreements. No regulator has yet written an egress rule for a factory floor.

Third, the buyer who ignores the battle inherits the winner's defaults: its identity model, its update channel, its data terms and its price list once the contest is over.

And fourth, the battle is the buyer's best bargaining position. Suppliers concede terms on portability, data access, co-investment and sovereignty while position is contested. They stop conceding once it is settled. The procurement director is right to want competition. The way to keep it is to buy so that the edge stays contestable.

How buyers should choose

Enterprises that understand the battle choose infrastructure partners for control, not just for price. They make six moves, in this order.

First, map the control points at each class of site. For factories, stores, depots, branches or clinics, record who holds access, identity, policy, inference, updates and billing today. Most enterprises will find that these functions are spread across a carrier, a cloud provider, equipment makers and integrators, with no design behind the split.

Second, separate what suppliers want to bundle. Connectivity, edge compute, models and the control plane are four different purchases. A bundle is fine when it is chosen. It is dangerous when it arrives by default inside a network renewal or an enterprise cloud commitment.

Third, choose for accountability at the location. Ask who answers when a site's edge stops at three in the morning: who has people, spares, local licences and authority to act on site. That is the asset telcos can offer and hyperscalers usually cannot. Value it explicitly, and make suppliers who claim it prove it with named people, response times and a record of site incidents they have actually handled.

Fourth, write exit into the architecture. Specify portable workloads, open interfaces and tested data export before the first site goes live. In Europe, use the Data Act's switching rights. Prefer edge hardware that can run more than one supplier's stack. Every edge contract states how the enterprise leaves and what it costs.

Fifth, use the contest while it lasts. Operators want to move up the stack and will offer sovereign options, local operations and co-investment to do it. Hyperscalers want locations inside their control plane and will offer to run through local operators, as Google's managed partner model shows. Take those offers on the enterprise's terms: control, portability and data access in writing, now.

Sixth, give edge partner strategy to whoever owns decision architecture. Choosing edge infrastructure is choosing who holds a claim on edge decisions. It belongs with the leader accountable for where intelligence lives in the enterprise, as argued in *The Intelligent Edge Changes Who Decides*, not with network procurement alone.

And stop doing four things.

Stop buying edge capability as an add-on to a connectivity contract. Stop signing enterprise-wide cloud commitments that pull every site into one control plane by default. Stop assuming the telco is only a pipe or the hyperscaler is the only platform; the evidence shows neither assumption holds. Stop evaluating edge partners on unit price, when the cost that matters is the cost of leaving.

Telcos have the last mile. Hyperscalers have the control plane. Make sure the enterprise keeps the decision.

What to do about it

- **Map the control points at each class of site.** Record who holds access, identity, policy, inference, updates and billing today, and where the split was never designed.
- **Separate what suppliers want to bundle.** Buy connectivity, edge compute, models and the control plane as distinct decisions, so any bundle is chosen rather than inherited.

- **Choose for accountability at the location.** Value the supplier that has people, spares, licences and authority on site, and make claimants prove it.
- **Write exit into the architecture.** Specify portable workloads, open interfaces and tested data export before go-live, and use the EU Data Act's switching rights where they apply.
- **Use the contest while it lasts.** Take operators' and hyperscalers' competing offers on the enterprise's terms, with control, portability and data access in writing now.
- **Give edge partner strategy to the owner of decision architecture.** Choosing edge infrastructure is choosing who holds a claim on edge decisions, so it cannot sit with network procurement alone.

Sources

- Cloudflare, The Network is the Computer, 2019: <https://blog.cloudflare.com/the-network-is-the-computer/>
- Cloudflare, Workers AI GPU upgrade and GPUs in more than 180 cities, 2024: <https://www.cloudflare.com/press/press-releases/2024/cloudflare-enhances-ai-inference-platform-with-powerful-gpu-upgrade/>
- NVIDIA, NVIDIA and Nokia to pioneer the AI platform for 6G, 2025: <https://investor.nvidia.com/news/press-release-details/2025/NVIDIA-and-Nokia-to-Pioneer-the-AI-Platform-for-6G--Powering-Americas-Return-to-Telecommunications-Leadership/default.aspx>
- IDC, Global spending on edge computing to reach nearly \$380 billion by 2028, 2025: <https://www.idc.com/resource-center/press-releases/idc-estimates-global-spending-on-edge-computing-to-grow-at-13-8-reaching-nearly-380-billion-by-2028/>
- AWS, AWS Wavelength locations: <https://aws.amazon.com/wavelength/locations/>
- AWS, General availability of AWS Wavelength in Germany, 2021: <https://aws.amazon.com/about-aws/whats-new/2021/12/aws-general-availability-aws-wavelength-germany>
- Google Cloud, Managed GDC Provider initiative, 2024: <https://cloud.google.com/blog/topics/hybrid-cloud/new-managed-gdc-provider-initiative/>
- Microsoft, AT&T to run its mobility network on Microsoft's Azure for Operators cloud, 2021: <https://news.microsoft.com/2021/06/30/att-to-run-its-mobility-network-on-microsofts-azure-for-operators-cloud-delivering-cost-efficient-5g-services-at-scale/>
- Meta Engineering, Project Waterworth, 2025: <https://engineering.fb.com/2025/02/14/connectivity/project-waterworth-ai-subsea-infrastructure/>
- W.Media, AI infrastructure demand pushes Alphabet's 2026 capex guidance to US\$205 billion, 2026: <https://w.media/ai-infrastructure-demand-pushes-alphabets-2026-capex-guidance-to-us-205-billion/>

- Verizon, Second-quarter 2026 results, Form 8-K Exhibit 99, 2026: <https://www.sec.gov/Archives/edgar/data/0000732712/000073271226000040/a2026q2exhibit99.htm>
- Boston Consulting Group, Boosting value creation in telcos, 2025: <https://www.bcg.com/publications/2025/boosting-value-creation-in-telcos>
- Microsoft, Vodafone and Microsoft sign 10-year strategic partnership, 2024: <https://news.microsoft.com/2024/01/15/vodafone-and-microsoft-sign-10-year-strategic-partnership-to-bring-generative-ai-digital-services-and-the-cloud-to-more-than-300-million-businesses-and-consumers/>
- Verizon, Verizon unveils AI strategy to power next-gen AI demands, 2025: <https://www.verizon.com/about/news/verizon-unveils-ai-strategy-power-next-gen-ai-demands>
- TelecomTV, Why Verizon is excited about AI Connect, 2025: <https://www.telecomtv.com/content/telcos-and-ai-channel/why-verizon-is-excited-about-ai-connect-52203/>
- UC Today, Ericsson announces \$1.1bn writedown after weak Vonage forecast, 2024: <https://www.uctoday.com/unified-communications/ericsson-announces-1-1bn-writedown-after-weak-vonage-forecast/>
- RCR Wireless, Ericsson finalises Aduna network API joint venture, 2025: <https://rcrwireless.com/20250730/5g/ericsson-aduna-jv>
- Analysys Mason, Hyperscalers' retreat from the private networks market, 2025: <https://www.analysysmason.com/research/content/articles/hyperscalers-private-networks-rma17/>
- Deutsche Telekom, Germany's first AI factory for industry goes into operation in Munich, 2026: <https://www.telekom.com/en/newsroom/latest-updates/media-information/2026/2/germany-s-first-ai-factory-for-industry>
- Deutsche Telekom, T Cloud Public: hyperscaler power with European sovereignty, 2026: <https://www.telekom.com/en/newsroom/latest-updates/media-information/2026/2/t-cloud-public-sovereign-power>
- UK Competition and Markets Authority, CMA announces package of actions on business software and cloud services, 2026: <https://www.gov.uk/government/news/cma-announces-package-of-actions-on-business-software-and-cloud-services>
- Regulation (EU) 2023/2854 (Data Act), Article 29: switching charges: <https://www.springlex.eu/en/packages/data-act/data-act-regulation/article-29/>

Read online: bdgadvisory.com/perspectives/infrastructure-power-battle