BERNSTEIN
SOCIETE GENERALE GROUP

20 October 2026

US Telecom

US Telecom: Telcos getting an AI narrative? Assessing the Edge AI opportunity


Laurent Yoon
+1 917 344 8502
laurent.yoon@bernsteinsg.com


Madison Rezaei
+1 917 344 8622
madison.rezaei@bernsteinsg.com


Daniel Zhu
+1 917 344 8309
daniel.zhu@bernsteinsg.com


Martin Boruchowicz
+1 917 344 8564
martin.boruchowicz@bernsteinsg.com


Andrew Chung
+1 917 344 8302
andrew.chung@bernsteinsg.com

Telcos and CableCos have increasingly discussed Edge AI as a potential growth opportunity, leveraging existing network infrastructure and real estate. Needless to say, AI is the buzzword du jour, so we took a deeper look into what the opportunity actually is. Four key takeaways:

We've been here before. Telecom has a long history of trying to move beyond connectivity and capture value further up the stack. Cloud computing was one such opportunity in the 2000s, which did not work out. Then came IoT, where much of the value accrued elsewhere else in the ecosystem. More recently, MEC and 5G edge compute promised a wave of low-latency applications that never arrived. Is Edge AI is the latest iteration of a familiar theme?

Perhaps it's different this time. Unlike prior efforts, operators are not proposing a large capex project to compete with hyperscalers on datacenter capacity. Rather, the opportunity is to monetize assets they already own—fiber-connected facilities with power, cooling, permits, and space. In other words, this is less about building the next AI factory and more about finding a productive use for infrastructure already sitting on their balance sheets.

The question is demand. An orchestration layer such as NVIDIA AI Grid can pool hundreds of distributed locations into a single computing platform, making fragmented assets commercially usable. If that's the case, however, proximity becomes less differentiated. The more compelling edge thesis is ultra-low latency. Processing inference closer to end users can reduce latency to under 10ms, potentially enhancing latency-sensitive applications such as robotics, industrial automation, real-time video analytics and other physical AI applications. The challenge is that this sounds awfully familiar. Low-latency applications were also the core argument for MEC, yet meaningful demand never emerged. It remains unclear which applications truly require shaving tens of milliseconds and whether those use cases will develop at enough scale to matter.

Let's assume demand arrives. Using Charter as a case study, we estimate its edge footprint could generate $0.8 to $1.5B of annual revenue and potentially >$2B under favorable monetization scenarios, though substantially less under a colo scenario where Charter does not own compute. That's not trivial for a sector starved for growth, but it is unlikely to change the investment narrative anytime soon, given the time it takes to reach that level of utilization. More importantly, capturing the higher-value economics requires owning the compute. By our estimates, fully utilizing Charter's available capacity could require 150K - 200K NVIDIA GPUs, implying $2-$3B of GPU spend alone. The opportunity is meaningful, but so is the investment required to capture it. If operators want AI economics rather than data-center landlord economics, they will likely need to own the compute, implying billions of dollars of incremental capital. For a sector increasingly focused on FCF and capital discipline, that may prove a harder sell than the Edge AI narrative itself. For now, Edge AI looks less like a near-term growth engine and more like a potential option that has yet to arrive.

See the Disclosure Appendix of this report for required disclosures, analyst certifications and other important information. Alternatively, visit our Global Research Disclosure Website.

First Published: 20 Oct 2026 17:10 UTC Completion Date: 20 Oct 2026 17:10 UTC

BERNSTEIN TICKER TABLE

			19 Oct 2026		TTM	Adjusted EPS				Adjusted P/E (x)	
Ticker	Rating	Cur	Closing Price	Price Target	Rel. Perf.	Cur	2025A	2026E	2027E	2025A	2026E	2027E
CHTR (Charter)	M	USD	108.51	120.00	(71.4)%	USD	36.20	38.21	40.14	3.0	2.8	2.7
CMCSA (Comcast)	M	USD	21.90	23.00	(35.7)%	USD	4.30	3.43	3.71	5.1	6.4	5.9
T (AT&T)	M	USD	23.20	24.00	(25.4)%	USD	2.12	2.35	2.68	10.9	9.9	8.7
TMUS (T-Mobile)	M	USD	156.09	180.00	(46.5)%	USD	9.72	10.68	13.24	16.1	14.6	11.8
VZ (Verizon)	M	USD	44.15	45.00	(6.2)%	USD	4.68	5.00	5.20	9.4	8.8	8.5
SPX			7,690.73									

O - Outperform, M - Market-Perform, U - Underperform, NR - Not Rated, CS - Coverage Suspended

Source: Bloomberg, Bernstein estimates and analysis.

INVESTMENT IMPLICATIONS

We maintain our Market-Perform rating for T (PT: $24), VZ (PT: $45), TMUS (PT: $180), CMCSA (PT: $23), and CHTR (PT: $120) —all unchanged.

DETAILS

CAN EDGE AI BE A REAL OPPORTUNITY FOR CABLE COMPANIES AND TELCOS?

As AI shifts from training massive models toward running them at scale, a growing share of inference may need to happen close to users and devices, where latency, bandwidth costs, and data locality matter. Cable operators and Telcos sit on dense national footprints of powered, fiber-connected network facilities, many of them a few milliseconds from the homes and businesses they serve.

Nearly every company we cover has raised it this year. Comcast launched an edge AI field trial with NVIDIA in March, placing GPUs in regional facilities that reach 65 million homes and businesses, and said in September that they now run more than 200 AI-powered edge compute centers, with Fastly and Oracle already using the platform. Verizon told investors on its Q2 call that it is in the early stages of retrofitting central offices into data centers for edge inference, pitching their ready power and permits to multiple prospective partners as part of its broader AI Connect initiative.

We use Charter as our case study as they have given specific details of their opportunity. In September, Charter announced partnerships with Cast AI, HP, Hydra Host, and World Wide Technology to commercialize its Edge Compute Infrastructure platform, which places NVIDIA GPUs inside its hubs and headends. CEO Chris Winfrey recently quantified more than 1,200 local edge data centers, with about 250 MW of fallow capacity across roughly 600 of them, freed up as network equipment was virtualized and available without additional investment.

In this note, we explain what edge computing is, walk through the use cases where proximity to users matters, and lay out what Charter has and what it is offering. We then size the potential revenue opportunity, translating Charter's fallow megawatts into revenue, and consider what it implies for the rest of the Cable and Telco operators pursuing the same strategy.

EXHIBIT 1: **Major TelCos/CableCos are talking about Edge AI as a potential growth opportunity**

Every major operator is talking about edge AI, but they are at very different stages

	Charter	Comcast	Verizon	AT&T
Stage	Early commercial	Field trial + early partners	Retrofitting, partner talks	Evaluating, selective
Edge footprint	1,200+ local edge data centers (former hubs and headends)	200+ AI-powered edge compute centers reach of 65M homes and businesses	Thousands of central offices with power, space and cooling	GPUs at 6 regional Cisco data centers
Disclosed capacity	~250 MW fallow across ~600 sites			
Partners & hardware	NVIDIA AI Grid, Cast AI, Hydra Host, HP, WWT, RTX PRO 6000 Blackwell	NVIDIA AI Grid, HPE, Fastly, Oracle, RTX PRO 6000 Blackwell	AI Connect suite, partners not named	NVIDIA and Cisco AI Grid for enterprise IoT
What management said	*"250 megawatts of fallow capacity without any additional investment*	*"We now treat the network as a platform"*	*"Retrofitting many of our central offices into data centers for inference*	*Does it "make sense to insert some compute capability" closer to the edge?*

Source: Company reports and presentations; Bernstein analysis

WHAT IS EDGE COMPUTE AND WHY MIGHT IT BE IMPORTANT?

Edge computing moves data processing and storage out of centralized data centers and onto servers physically close to where users and devices generate and consume data. The benefits of edge computing are lower latency, less congestion, better resiliency, and easier compliance when data must stay local.

Edge sites are a very different asset from the large data centers driving today's AI build-out. Central AI campuses are purpose-built for scale, with most individual deployments of 50 MW or more, liquid-cooled racks that each draw well over 100 kW, and thousands of tightly networked GPUs training frontier models or serving them in bulk. Edge sites instead are hundreds or

thousands of small, distributed facilities, sitting closer to the end user, and each with modest power, best suited for smaller compute workloads or inference tasks, where a model answers a request, rather than training.

How scattered sites work together

In a large AI datacenter campus, thousands of GPUs sit side by side in the same building, already wired together on one high-speed network. When a customer needs computing power, the operator hands a slice of one big, connected pool. Pooling the capacity is built into the design.

In edge computing, Charter's GPUs would be spread across hundreds of separate buildings in different towns, each with its own amount of power, space, and cooling. To turn those scattered sites into one product a customer can buy, an operator needs a software layer that acts as the central dispatcher. It keeps track of what every site has available in real time, and when a request comes in, it sends that request to the best location.

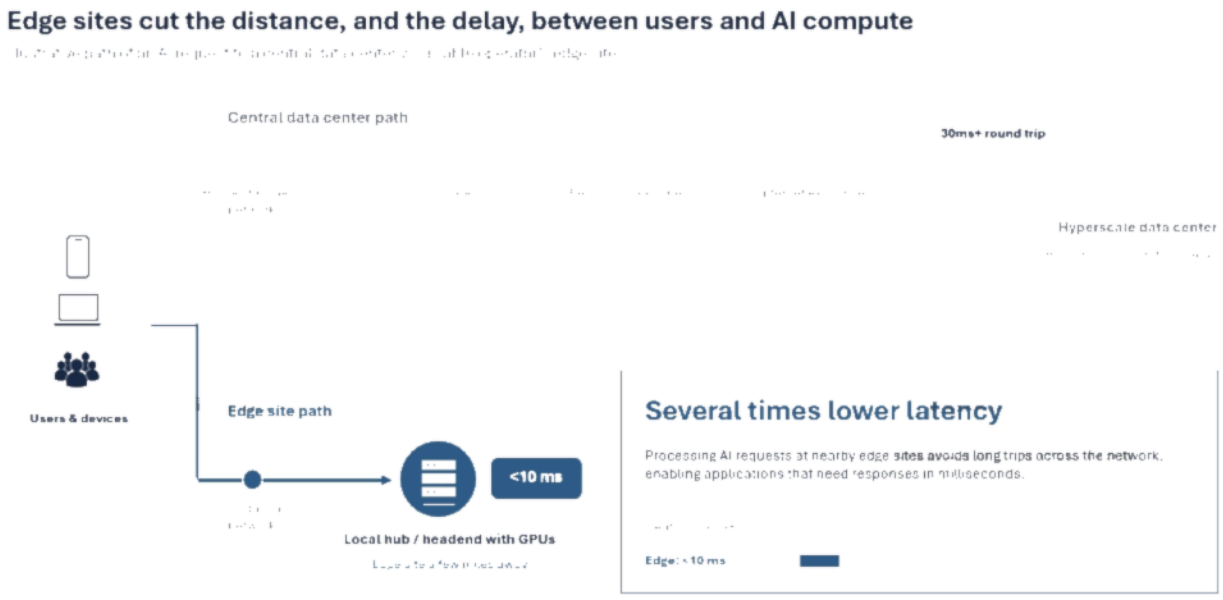
NVIDIA's AI Grid is the reference design Charter is using for this. It's an intelligent control plane that knits fragmented infrastructure into a single, programmable platform, along with NVIDIA's inference software for workloads like low-latency video and speech

What Charter's sites actually are

Charter's edge locations are its headends and hubs, the facilities that make up the backbone of a cable network. Headends are where traffic is aggregated, historically where video signals were received and processed, and where the cable modem termination systems (CMTS) that manage broadband sit. Hubs are smaller distribution points that connect headends to the fiber nodes serving individual neighborhoods. Both are already fitted with fiber, utility power, backup power, and cooling, because they have always run racks of network equipment.

However, cable operators are virtualizing their access networks, replacing large racks of dedicated CMTS hardware with software running on standard servers, and pushing other functions out into field nodes, and video's decline has reduced the equipment needed to process TV signals. This results in spare capacity, and for Charter it's about 250 MW of fallow power across roughly 600 of its more than 1,200 sites.

EXHIBIT 2: **Processing AI requests at nearby edge sites cuts the trip across the network, reducing latency to under 10ms**



Source: Charter at SCTE TechExpo; Bernstein analysis

WHY IT MATTERS AND EDGE COMPUTING USE CASES

As AI shifts from training to inference, more workloads involve machines reacting to the physical world in real time. Spectrum's Justin Colwell argued at SCTE that physical AI cannot work with a roughly 100-millisecond round trip to a distant cloud, and that Spectrum's 1,000-plus sites have become a strategic advantage. Charter says its footprint puts compute within about 10 milliseconds of 500 million devices.

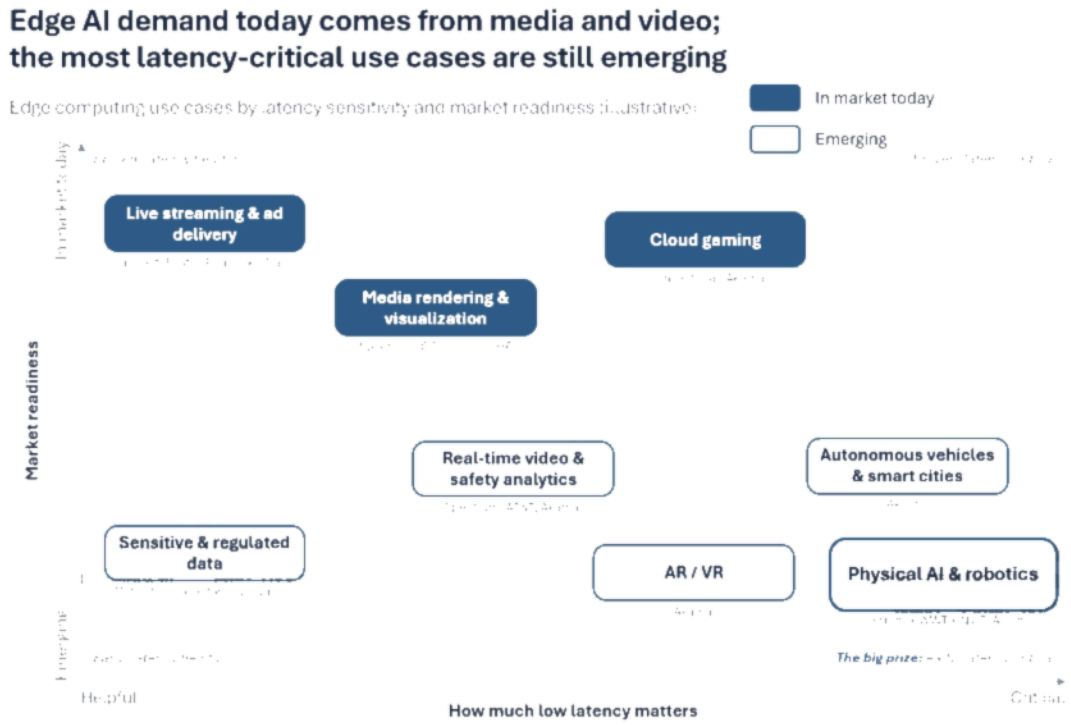
Use cases

The use cases where proximity matters fall into a handful of categories, several of which Charter and its partners have already demonstrated. The most prominent is physical AI: robotics, industrial automation, and machines that must react to the real world in split seconds, a response time a distant cloud cannot guarantee. Spectrum showcased this at SCTE with WWT and NVIDIA through a humanoid robot demo.

Media rendering and visualization is the use case Charter has taken furthest commercially. Its March deployment gives animation studios near-local GPU power over Spectrum fiber for CGI rendering, and HP's Z Boost extends that model to graphics studios that want more compute without building their own. Closely related is real-time video, vision, and safety analytics, where cameras and sensors produce data too heavy and time-sensitive to haul back to a central cloud.

Beyond those, AR, and VR is another use case, where lag directly degrades the experience, and to IoT, smart cities, and connected vehicles, where massive volumes of device data need local, real-time processing. Comcast is using its edge for live streaming, with Fastly running on its platform to handle peak sports traffic for Peacock and Oracle using it for low-latency media delivery. Finally, workloads involving sensitive or regulated data benefit from staying local, easing privacy and data sovereignty requirements.

EXHIBIT 3: **Edge AI demand today comes from media rendering, video, and live streaming; the most latency critical use cases like robotics and autonomous veichles are still emerging**



Source: Charter/Spectrum press releases (Mar-26, Sep-26); Comcast/NVIDIA (Mar-26); Akamai; AT&T; Bernstein analysis

WHAT'S THE REVENUE OPPORTUNITY?

Notably, Charter is not trying to build this business on its own. It brings the sites, power, cooling, fiber, and the compute hardware inside them, and is leaning on partners for everything else. As described above, NVIDIA's AI Grid provides the orchestration layer that ties the edge sites together into one network. On top of that, Charter announced collaborations with Cast AI, HP, and Hydra Host, each filling a different gap.

Cast AI runs a marketplace that makes Spectrum's GPUs show up inside customers' existing Amazon, Google, or Microsoft cloud setups as if they were simply more machines in their own accounts, giving Charter a ready-made storefront without its own sales force. HP's Z Boost works like a car-sharing service for GPUs, letting design, animation, and engineering teams tap Spectrum-hosted GPU power on demand instead of each workstation needing its own expensive card. Hydra Host acts as a broker that finds AI companies needing GPU capacity and rents them Spectrum-owned capacity.

Charter owns the infrastructure; NVIDIA supplies the software that connects it; Cast AI and Hydra Host bring in demand; HP packages it for specific uses.

HOW MUCH IS CHARTER'S ~250 MW OF FALLOW POWER ACTUALLY WORTH?

The starting point is the roughly 250 MW of fallow power across about 600 of Charter's hubs and headends. Management did not specify whether that figure is total facility power or critical IT load, the power that actually reaches the servers. We treat it as facility power, the more conservative approach, which means some of it must be set aside for the building itself before any of it can run GPUs.

The standard way to measure that set-aside is power usage effectiveness (PUE), the ratio of a facility's total power to the power that actually reaches its servers. Our colleagues covering hardware, semiconductors, and data centers have done extensive work on AI rack economics, showing that purpose-built AI facilities run at a PUE (power usage effectiveness) of roughly ~1.25, meaning about 80% of power reaches the compute racks while the rest goes to overhead: lighting, HVAC, security, power conversion losses, and cooling.

Charter's hubs and headends were built to run network equipment, not dense compute, so they are likely to be less efficient than central data centers designed for AI. Some overhead, such as lighting, security, and installed cooling equipment, is already in place, but the extra cooling and power conversion for new GPUs is still required. We therefore estimate a PUE of 1.3 to 1.5 for Charter's sites giving roughly **165 to 190 IT MW** available for GPUs and racks.

WHAT IS EACH MEGAWATT WORTH?

How much revenue a megawatt generates depends heavily on the business model. Our Bernstein data center team estimates a wide range of annual revenue per megawatt across the industry. When looking at longer term data center contracts, we see that at one end are neoclouds such as CoreWeave, which own the GPUs and the software that integrates them with customers, and typically earn around $9 million to $10 million per megawatt per year. At the other end are colocation operators such as Digital Realty and Equinix, which do not own the compute; they rent out the building, power, and cooling to tenants who bring their own hardware, and earn roughly $3 million to $4 million per megawatt per year. At the lowest end, emerging AI infrastructure players like RIOT and WULF see rates of $1 million to $2 million per megawatt per year.

Where does Charter fit? Its September press release states that the GPUs on its platform are owned, hosted, and operated by Spectrum, and its deployments use NVIDIA's RTX PRO 6000 Blackwell Server Edition GPUs. Owning and operating the compute places Charter closer to the neocloud model than to colocation, which implies it should earn more than a colocation landlord on each megawatt.

However, we do not expect Charter to reach full neocloud rates. Charter relies on partners such as Cast AI and Hydra Host to bring in customers and integrate its compute, so those partners will take a share of the economics. Additionally, Charter's capacity is spread across hundreds of fragmented edge sites rather than concentrated in large central data centers and also, its compute is older and less powerful than the flagship hardware in central campuses.

We therefore apply a haircut to typical neocloud rates and assume Charter earns $5 million to $8 million per megawatt per year, reflecting pricing on long-term contracts of around five years. Applied to the 165 to 190 IT MW we estimate above, that implies a revenue opportunity of roughly **$0.8 billion to $1.5 billion** a year once the capacity is fully deployed.

If Charter can sell capacity on shorter-term contracts, pricing could be meaningfully higher. Neoclouds can earn anywhere from $15 million -$25 million per year on shorter-term deals. Applying a similar discount, we estimate Charter could earn $10 million to $20 million per megawatt, implying a revenue opportunity of roughly **$1.7 billion to $3.8 billion a year**, though shorter

contracts also bring less predictable demand and more utilization risk.

Ultimately, contract terms will depend on Charter's customers. Hyperscalers and AI labs typically sign multi-year contracts for large blocks of capacity at lower prices, while enterprises buy smaller amounts on shorter terms at higher prices. Charter's early partners, such as Cast AI, Hydra Host, and HP, point toward enterprise customers and shorter-term deals.

Timing matters as well. These rates reflect today's market, where demand for AI compute outpaces supply. What Charter can charge will depend on how quickly it brings capacity online, since hyperscalers and neoclouds are adding compute rapidly and NVIDIA's next generation of chips will push Blackwell cards down the performance curve. As supply catches up, pricing could compress, especially for older hardware.

Although it has been implied that Charter will own and host the compute, acquiring the hardware to fill this capacity would be difficult and would require a sizable capital commitment. If Charter instead shifts toward a colocation model, renting out its space, power, and cooling, we would expect rates closer to $0.5 million to $3 million per megawatt per year, applying the same haircut to centralized data center rates. That would imply a revenue opportunity of roughly **$80 million to $580 million a year**, well under $1 billion. Against Charter's ~$58 billion in annual revenue today, this opportunity remains small.

EXHIBIT 4: **Charter owns the sites, power, fiber, and GPUs; NVIDIA connects them, and partners bring the customers and take a share of the revenue**

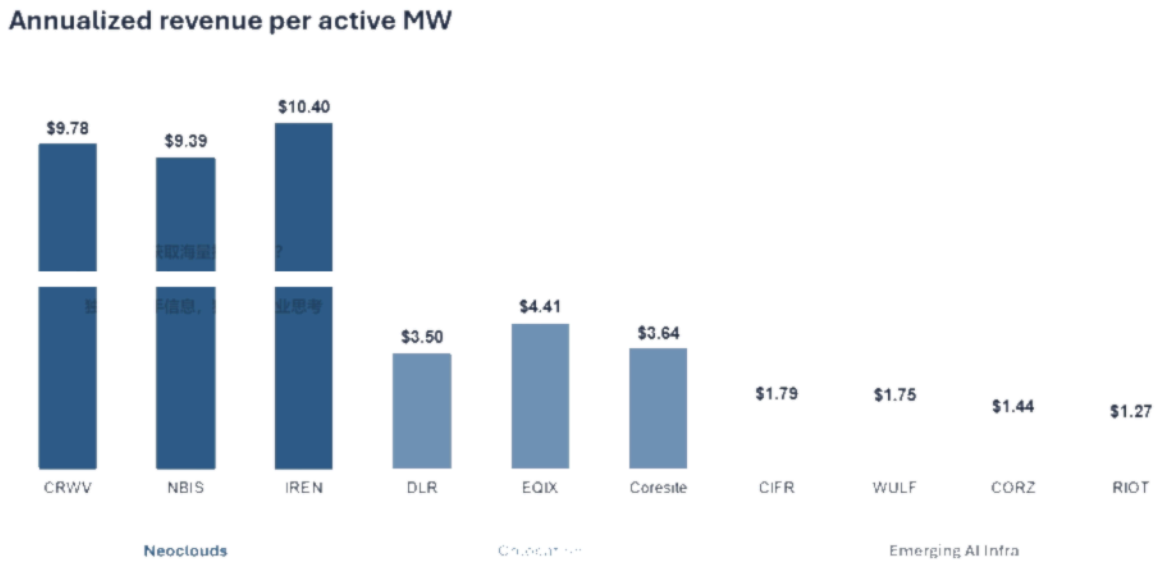


Source: Charter/spectrum press releases (Mar-26, Sep-26); Charter management at SCTE TechExpo; Bernstein analysis

人在草木间

EXHIBIT 5: **Revenue per megawatt depends on who owns the compute and the contract length**

Prices shown are for longer term contracts



Source: Company reports; Bernstein estimates & anlaysis

EXHIBIT 6: **Charter's ~250 MW of fallow power could be worth $0.1 - $3.8B a year, depending on whether it owns the GPUs and the length of its contracts**



Source: Charter (Sep-26); Charter/Spectrum press releases; Bernstein estimates & analysis

CHALLENGES IN THE EDGE COMPUTING OPPORTUNITY

While we cited a $0.8 billion to $3.8 billion revenue opportunity for Charter, getting there will likely be difficult. If Charter chooses to own the compute, it has to buy the GPUs and hardware. A single RTX PRO 6000 GPU, the card Charter used in its demonstrations, now sells for about $13,250, up roughly 55% since launch because of a memory shortage. Using the framework in our blackbook, The AI Infrastructure Value Chain, which estimates that GPUs account for roughly two-thirds of the cost of an AI rack, we estimate that filling Charter's 165 to 190 IT MW would take roughly 150,000 to 200,000 GPUs. That

implies about $2 billion to $3 billion in GPUs alone and more once servers, memory, and networking are included. This is a large amount of capex at a time when Charter is trying to bring capex down and grow free cash flow. Even with the capital, ramping up will take time. Charter relies on partners like Cast AI and Hydra Host to bring in customers, and running hundreds of separate compute sites is new for the company.

The bigger question is whether the demand is there, and telcos have tried this before. Between 2019 and 2022, the industry pushed "5G edge" computing. Verizon placed AWS servers inside its network across 19 US metros, and AT&T struck edge deals with Microsoft, Google Cloud, and IBM. Ericsson's own edge platform, Edge Gravity, shut down in 2020. None became a meaningful business. The apps they were built for, like AR/VR and self-driving cars, never scaled, and the cloud providers kept the customers and most of the economics. An AT&T executive admitted this year that edge compute was long a use case without a market. Today's version is different: it runs on fiber-connected hubs rather than 5G cell networks, uses GPUs built for AI, and the operator, not a cloud provider, owns the compute.

Still, demand remains early. AI traffic makes up only about 0.1% of downstream usage on Spectrum's network. Widespread adoption of robots and self-driving cars, the headline use cases, are still many years ahead of us. The market may be closer than it was in 2019, but still too early for large scale Capex as demand is yet to arrive.