



Telecom & Digital Infrastructure
Technology and Investment Trends
A look ahead into 2026

December 10, 2025

Digital Infrastructure: Momentum in 2025, Validation in 2026

Key Takeaways

- **AI-driven infrastructure** spend is reshaping the market with a wave of data center buildouts, placing data center interconnect at both inter and intra facility levels in sharp focus as a premium asset class, with U.S. demand forecast to double backbone capacity by 2029 through 92,000 new route miles supported by hyperscaler contracts and 30-40% annual traffic growth, while subsea cable investment surges with \$13 billion committed between 2025 - 2027 through projects such as Meta's 50,000 kilometer Project Waterworth and new builds by Google, Amazon, and Microsoft to diversify global routes.
- **Telecom consolidation is accelerating** as capex spending, which weakened in 2024, remains muted and is projected to extend further into 2026, compounded by weak 5G monetization. Landmark deals on the service provider side such as Charter and Cox's \$34.5 billion merger, AT&T's \$5.75 billion acquisition of Lumen fiber assets, and EchoStar's \$43 billion spectrum divestitures highlight survival through scale as the dominant strategy. At the same time vendors are scaling down, exiting non-core businesses, and reorganizing to align with AI infrastructure demand.
- **Capex intensity in AI data centers is unprecedented**, with gigawatt-scale campuses announced worldwide, in some cases backed by sovereign AI initiatives that treat compute as strategic infrastructure. Yet the sector faces mounting risks of overbuild, fragile neo-cloud business models, and above all energy constraints that have become the defining bottleneck, driving demand for powered land strategies, alternative energy solutions, and heightened regulatory scrutiny across major regions.
- **Valuation risks remain elevated**, with record fundraising concentrated in mega-funds and annual AI infrastructure spending exceeding \$300 billion in 2025. Investor caution around AI economics is rising, driven by short GPU depreciation cycles, uncertain workload maturation, and limited RoI from generative AI deployments. Given the multilayered nature of the AI stack, asset value disinflation is uneven across segments, reinforcing the importance of disciplined evaluation of both technical fundamentals and commercial viability.
- **LEO satellites are emerging as an early convergence point** between terrestrial and space infrastructure, while traditional mobile and fixed access providers struggle with monetization and GEO operators face shrinking revenues and consolidation pressures, exemplified by SES and Intelsat. This dynamic is heightening investor interest in Ground Station as a Service (GSaaS) models that resemble tower and data center investments. At the same time, the satellite sector is undergoing rapid transformation as Amazon's Kuiper prepares for commercial rollout to challenge Starlink's five-year lead, and sovereign programs increasingly treat satellites as strategic assets, positioning national constellations as critical to digital infrastructure resilience.

In 2025, the **digital infrastructure market continued its "K"-shaped trajectory**, with explosive growth in AI-driven segments set against persistent challenges in traditional telecom. Data centers expanded rapidly to meet unprecedented demand for compute power, while low Earth orbit (LEO) satellite networks and subsea fiber investments sustained strong momentum. In contrast, telecom operators struggled with ongoing capital constraints and revenue pressures, prompting a wave of mergers and asset sales as 5G monetization stalled and fiber rollout under programs such as BEAD in the U.S., and equivalent programs in the EU and UK, achieved only limited progress. This divergence highlighted a broader shift, as hyperscalers and technology giants increasingly shaped infrastructure priorities, leaving legacy players to consolidate in order to survive.

As 2026 unfolds, the sector confronts **growing uncertainties** that could reshape investment flows and operating models. Concerns about a potential "AI bubble" have left the market cautious as it awaits proof of sustainable revenue streams and returns on investment. Aligning revenues with expenditures, particularly given the short depreciation cycle for GPU capital spending, remains a central issue for prospective investors. Additional pressures include worsening global power shortages that hinder data center expansion, shifting AI dynamics in which efficiency gains may temper hyperscale spending, persistent geopolitical tensions that disrupt supply chains, evolving rules around data sovereignty, and rising exposure to cybersecurity threats.

AI has pushed data-center spending to unprecedented levels, with capex accelerating sharply as governments and hyperscalers treat **compute as strategic infrastructure** in the U.S.–China geopolitical race. The surge is increasingly financed through heavy debt issuance by major cloud providers, rising from \$15 billion in 2024 to \$125 billion in 2025 according to UBS, and

through emerging "circular" models where key suppliers reinvest into customer build-outs to secure demand. There is no obvious end to this business model. One consequence of this flywheel effect is that Gigawatt-scale campuses are being announced across the U.S. and Europe, though many remain early-stage and unproven. Analysts increasingly warn that global data-center **power demand could more than double** by 2030, driven by AI workloads and a long-term shift toward latency-sensitive inference infrastructure. Yet the build-out also carries real risk: soaring capex forecasts, uncertain workload maturation, grid constraints, and the growing exposure to neo-cloud intermediaries whose business models are already straining under scaling and pricing pressures. These factors raise the possibility of overbuild and uneven market consolidation as AI deployment patterns become clearer. To balance these risks, modular and flexible smaller-scale deployment models are emerging. They are less capital intensive to initiate and sustain, far less constrained by power and land. Positioned closer to end-user communities and faster to deploy, they follow a "grow and fund as needed" approach to compute. This model will develop alongside hyperscaler buildouts and is amenable to smaller investments with a correspondingly lower risk profile.

The **data center interconnect** market is becoming a premium asset class as AI workloads drive surging demand for dark fiber and high-capacity wavelengths. The U.S. alone is forecast to need 92,000 new route miles by 2029, with providers like Zayo and Lumen already expanding networks through hyperscaler partnerships. Traffic growth of 30–40% annually and multi-year contracts with cloud giants underpin durable returns, though capital intensity and power constraints remain key risks.

The **incumbent telecom operator capex malaise**, which deepened in 2024, showed signs of stabilization in 2025 but remained muted, as global spending normalized under

the weight of heavy debt and flat revenues from disappointing 5G adoption. U.S. fiber deployments under BEAD began to take shape, though only modestly, and failed to offset broader stagnation. Stagnation was evident in the UK, where numerous small fiber AltNets are under severe pressure. Apart from a few outliers such as Belgium, Europe has largely passed the peak of its FTTH activity for the near future. Achieving acceptable ROI has remained challenging in the fiber sector. M&A activity and divestments accelerated, highlighted by EchoStar's exit from the mobile telecom business through a \$23 billion spectrum sale to AT&T and a \$20 billion sale of AWS-4 and AWS-3 assets to SpaceX. The latter deal intensified interest in SpaceX's wireless ambitions and its potential to compete rather than collaborate with terrestrial operators. AT&T also expanded its fiber footprint through the \$5.75 billion acquisition of Lumen's fiber assets, reinforcing the strategic importance of scale in fixed access. The \$34.5 billion merger of Charter and Cox further underscored consolidation pressures in the U.S. market, while similar transactions across regions, including the breakup of Altice in France and consolidation in the Argentine telecom market, reinforced the prevailing mindset that **survival depends on scale**.

The **satellite sector** saw Amazon's Project Kuiper rebranded as LEO and launching a preview program for select enterprise customers to test its services ahead of a broader commercial rollout in 2026. Starlink, with a five-year lead on LEO, strengthened its dominance in 2025 by adding more than 3,000 satellites, expanding the constellation to over 9,200 in orbit. It reported service in 42 countries and boosted peak-hour network speeds by 50% through technological advances, reaching typical performance of 200/30 Mbps. Partnerships such as direct-to-cell with T-Mobile achieved commercial rollout milestones, while Starlink's fixed broadband offerings captured significant rural market share, surpassing traditional wireless ISPs and GEO

satellite operators. Competitors in the direct-to-device market continued efforts to close the gap, but funding challenges spurred early consolidation, exemplified by the merger of Omnispace and Lynk under SES, which became a major strategic shareholder. These developments underscored the mounting pressure on GEO satellites, reflected in declining revenues and intensifying competition from LEO operators, culminating in consolidation moves such as the SES-Intelsat merger. The convergence of these trends, the rise of LEO and the strain on GEO, has heightened investor interest in the **Ground Station-as-a-Service (GSaaS)** model, which mirrors the appeal of tower infrastructure and data center markets. GSaaS is expected to drive corresponding M&A activity in modern software needed to bring GEO ground services in line with current expectations. Once the field consolidates, the dominant players in software are even more likely to control the GSaaS domain than those focused on ground station hardware platforms.

The **subsea cable market showed steady progress** in 2025, with investments exceeding \$13 billion across 2025–2027, nearly double the amount committed between 2022 - 2024, according to TeleGeography. Growth was driven by insatiable hyperscalers demand for bandwidth and redundant routes for AI data flows. Meta led activity via Project Waterworth, a 50,000-kilometer cable that will connect five continents – the world's longest subsea cable project. Meta also completed the core 2Africa infrastructure, recognized as the world's longest open-access subsea cable system, and advanced the Bifrost cable linking Southeast Asia with the U.S. at 260 Tbps. Google, Amazon, and Microsoft were similarly active in funding and launching new projects. Trans-Pacific and intra-Asian routes dominated new builds, as geopolitical tensions highlighted vulnerabilities and spurred diversified routing along with sovereign-backed initiatives to mitigate risk.

In 2025, **our team contributed to cutting-edge projects** in data center, AI compute infrastructure and telecom network infrastructure. Highlights include:

- Led assessments of AI data center buildouts, including sovereign AI initiatives; developed go-to-market strategies for greenfield ventures; and advised investors on opportunities across the AI stack, spanning compute and storage chipsets, networking, data management, and cybersecurity solutions.
- Collaborated with stakeholders to evaluate the impact of AI developments on data centers, cloud infrastructure, connectivity, and cybersecurity strategies.
- Provided advisory services for the deployment of 5G private wireless networks across various geographies and sectors, including airports, utilities and manufacturing.
- Advised on multiple investments in subsea cable projects, including a significant trans-oceanic initiative.
- Offered guidance on the planning and assessment of LEO satellite constellations, covering broadband and direct-to-device services.

During 2025 Xona Partners published **Insight Notes** on some of the prominent topics, including:

Part 1: AI Infrastructure Network Reinvention

Powering the AI Revolution: The Urgent Need to Reinvent the Network. AI training now moves petabytes daily and current fiber topology cannot keep up. We map how data centers cluster near cheap power, why long haul and metro dark fiber become strategic assets, and which optical and energy efficient technologies will capture value as bandwidth demand inflates. Read it to identify where long term fiber builds, wavelength services, and low power interconnects offer durable returns.

Betting on Sovereignty: Assessing Risks in the Sovereign AI Race. Nations talk sovereignty but often underwrite projects that lack paying customers, secure power, or portable software stacks. The note breaks sovereignty into data, operational, technical, and legal pillars and exposes common execution failures that destroy value. Read it to learn how to vet sovereign AI deals for real demand, enforceable exit rights, and resilient funding rather than political rhetoric.

Navigating Risks and Opportunities in Sovereign AI Infrastructure. Countries can create strategic advantage, but only when they pair funding with usable data, talent, and realistic procurement. The analysis highlights capital intensity, supply chain dependence, and the temptation to hide vendor lock in behind sovereignty branding. Read it for a practical checklist investors can use to separate credible national programs from contingent spending that risks idle compute.

AI Infrastructure at a Crossroads: Navigating Volatility and Opportunity. The market faces both durable demand for compute and acute financial stress from rapid depreciation and customer concentration. The note explains why neoclouds absorb hyperscaler overflow, how circular financing raises contagion risk, and where disciplined due diligence can protect capital. Read it to understand which infrastructure pieces justify long horizons and which require skeptical stress testing.

Part 2: Non-Terrestrial Networks

Orbiting on Demand: How GSaaS is Rewiring the Satellite Industry's Ground Game. Ground Stations as a Service turns fixed ground infrastructure into a scalable, API driven asset with recurring revenue. The paper traces GSaaS evolution, shows how LEO traffic drives utilization, and explains why private equity treats GSaaS like tower assets. Read it to spot acquisition targets,

carve outs, and the commercial segments that will produce steady cash flow.

The Disruptive Force of Starlink: Reshaping Rural Connectivity, Challenging Incumbents.

Starlink forced incumbents to confront a new cost structure and changed expectations for remote connectivity. The note documents subscriber growth, the pressure on rural WISPs, and the strategic responses from GEO operators and governments. Read it to assess where Starlink's model creates durable winners and where public subsidies and subsea fiber still win on economics.

Friend or Foe: Rethinking Mobile Strategy in the Age of Non-Terrestrial Networks.

Satellites that secure mobile spectrum can bypass mobile operators and shift market power. The briefing examines regulatory frameworks, the operational demands of NTN, and high profile spectrum transactions that redraw competitive lines. Read it to evaluate how spectrum ownership and regulatory moves will affect mobile incumbents, satellite operators, and potential consolidation plays.

Navigating the Future of HAPS: Challenges, Investments, and Innovations. High altitude platforms offer targeted coverage and complement satellites but they still face cost, spectrum, and station keeping hurdles. The report compares fixed wing, airship, and balloon platforms and outlines commercial use cases that can reach payback. Read it to identify niche deployments and where regulatory clarity or technology advances could unlock investor scale.

Optical Space Communications: Enabling the Next Leap in Connectivity.

Laser feeder links promise orders of magnitude more throughput and stronger security than radio, yet atmospheric turbulence and automation limits block broad commercial scale. The note highlights defense as the early anchor market and portable ground stations as a nascent commercial pathway. Read it to judge where technical leadership meets

defendable contracting that can support returns.

Part 3: Telecom Networks

Turning the Page on 5G: Service Providers' Vision for 6G. Operators want a simpler, monetizable 6G that fixes 5G's cost and complexity. The paper summarizes operator priorities around efficiency, software first architectures, and realistic capex discipline. Read it to position investments around incremental, deployable technologies and services that operators can actually pay for.

Fortifying Global Connectivity: Enhancing Resiliency and Strategic Investments in Subsea Fiber Cable Networks. Subsea fiber remains the backbone of global data flows and AI driven demand elevates its strategic value. The analysis shows where new long haul corridors and redundancy investments will concentrate, and why dark fiber and wavelength economics improve with hyperscaler contracts. Read it to find resilient subsea projects and partnership structures that secure long term cash flows.

AI's Speculative Surge: Unpacking Mobile Traffic's Uncertain Path Forward. AI is transforming long-haul networks, but its mobile impact is uncertain. With saturated video and social media, mobile data growth is slowing, weakening investment cases for new networks. AI could spur demand if apps need low latency, edge execution, or heavy uplinks. While ecosystem players highlight uplink-driven scenarios to justify 5G Advanced, evidence is inconclusive. This framework helps investors assess whether AI will meaningfully change traffic, signaling, caching, or network architecture, or follow past hype cycles.

In 2025, the digital infrastructure narrative centered on momentum. By 2026, we expect it to shift toward validation and resilience, with emphasis on proving the economics of AI, addressing energy constraints, navigating geopolitical risks, and consolidating

fragmented markets. Looking ahead across segments of digital infrastructure, key themes to watch include:

- **AI platforms:** Scaling large language models is approaching its economic and technical ceiling. Investors are beginning to recognize diminishing returns, with rising capex for GPUs and energy costs outpacing revenue growth. The shift could signal a potential reallocation of capital toward more efficient architectures, hybrid approaches, and AI models that deliver measurable ROI rather than sheer scale.
- **LEO satellite services:** The anticipated launch of Amazon's LEO commercial services and expansion of Starlink's Direct-to-Cell service the first service with further development in the GSaaS segment as interests of satellite players and investors converge to develop new revenue streams.
- **Telecom networks:** Funding flow from programs such as the BEAD initiative to

breathe new life into the fixed access market with potentially significant spillover into fixed wireless access and satellite connectivity.

- **Infrastructure resiliency:** The growing strategic importance of wireline, wireless, and cloud infrastructures for national sovereignty and economic development.
- **Quantum technology:** Applications in both computing and cybersecurity.
- **Sustainability:** Challenges facing cloud and telecom infrastructures, particularly with the demands of AI and 5G.

In conclusion, we expect economic conditions and geopolitical tensions to be pivotal factors shaping investments in telecom, cloud, and digital infrastructure assets in 2026. A potential disinflation in asset valuations reminds investors of the critical importance of thorough technical and commercial due diligence to mitigate losses from hype-driven bubbles.

In This Edition

Part 1: AI Infrastructure

Powering the AI Revolution: The Urgent Need to Reinvent the Network

Betting on Sovereignty: Assessing Risks in the Sovereign AI Race

Navigating Risks and Opportunities in Sovereign AI Infrastructure

AI Infrastructure at a Crossroads: Navigating Volatility and Opportunity

Part 2: Non-Terrestrial Networks

Orbiting on Demand: How GSaaS is Rewiring the Satellite Industry's Ground Game

The Disruptive Force of Starlink: Reshaping Rural Connectivity, Challenging Incumbents

Friend or Foe: Rethinking Mobile Strategy in the Age of Non-Terrestrial Networks

Navigating the Future of HAPS: Challenges, Investments, and Innovations

Optical Space Communications: Enabling the Next Leap in Connectivity

Part 3: Telecom Networks

Turning the Page on 5G: Service Providers' Vision for 6G

Fortifying Global Connectivity: Enhancing Resiliency and Strategic Investments in Subsea Fiber Cable Networks

AI Bandwidth Bubble: Unpacking Mobile Traffic's Uncertain Path Forward

Part 1:

AI Infrastrcuture

Powering the AI Revolution: The Urgent Need to Reinvent the Network

Overview. The rise of AI is transforming the digital landscape, but its growth is constrained by a fundamental bottleneck: the physical movement of data within the data center and between data centers. AI model training and synchronization require the transfer of petabytes daily, a demand that far exceeds traditional internet traffic and is driving an unprecedented need for high-performance, low-latency, and power-efficient connectivity. This creates a significant and long-term investment opportunity in the underlying fiber and broadband infrastructure that forms the central nervous system of the AI revolution.

Evolution of Connectivity. Internet connectivity has centered on population hubs, with broadband and fiber infrastructure scaled to meet demand in urban centers. This model prioritized high-capacity fiber backbones in cities (metro networks) with long-haul interconnect, generating stable returns through consumer and enterprise traffic. The rise of cloud services introduced hyperscale data centers, optimized for cost efficiency (e.g., land, power and cooling) with some consideration for latency. As demand for cloud computing services increased bandwidth demand, it largely reinforced existing topology, by augmenting inter-data center bandwidth and incorporating edge computing data centers to address issues such as data sovereignty, localisation and lower latency. Fiber investments continued to yield predictable RoI tied to user proximity.

The rise of AI is transforming the connectivity landscape and creating significant investment opportunities in fiber and broadband infrastructure. AI training

A Snapshot of the AI Interconnect Market

A recent whitepaper by the Fiber Broadband Association underscores the massive infrastructure demand driven by AI. It estimates that AI data centers will require an additional 92,000 route miles of dedicated long-haul fiber in the U.S. by 2029, effectively doubling the existing ~95,000-mile backbone currently interconnecting data centers. Of this new build, approximately 77,000 miles will connect new, often remote, data centers to the network, while 15,000 miles will be for upgrades to existing routes.

This forecast is already manifesting in the market. For example, Zayo announced a major plan to construct 5,000 new long-haul route miles over the next five years, including five new routes. Similarly, Lumen Technologies is deploying ~1,500 new route miles in 2025, a strategic expansion supported by its recently announced partnerships with Microsoft and Meta to provide dark fiber for AI data center interconnects.

While these examples are U.S.-focused, reflecting the market's early lead in AI infrastructure deployment (closely followed by China), the demand for AI interconnect is a powerful global catalyst. It is driving fiber service providers worldwide to add and upgrade infrastructure. Key examples include Digital Realty and Colt Technology Services enabling AI hubs across Europe, and NTT Ltd.'s planned \$11 billion investment to expand its data center and optical infrastructure across Japan and the APAC region.

data centers increasingly prioritize locations with scalable and low-cost power, often near hydroelectric plants or renewable energy hubs. This shift is driving demand for new high-capacity fiber routes that connect these remote sites to each other and to global networks. Hyperscalers report that data center interconnect traffic is five to ten times higher than traditional access traffic and is growing at a much faster rate, between 30% - 40% annually, in countries with advanced AI infrastructure. This trend is fundamentally changing the dynamics of fiber connectivity, requiring new network topologies, greater scale, and substantial capital investment. For investors, this translates into rapid growth in fiber deployment as AI-driven workloads such as model training and synchronization move petabytes of data daily, far outpacing traditional internet traffic.




Figure 1 Worldwide data center computing forecast in \$Billion. [Source: Intersect360 Research]

Power constraints also shape AI data center design, creating additional opportunities for connectivity investments. To optimize power usage, providers often build data center campuses or clusters within a 50-100 km radius, sharing infrastructure like substations. This trend increases demand for inter-data center dark and lit fiber links with ultra-high throughput (e.g., 400 and 800 Gbps).

Evolving Connectivity Technologies

The growing demand for AI workloads is reshaping connectivity by introducing new requirements to support both AI and High Performance Computing (HPC). In response, industry alliances have begun to form, including the Ultra Ethernet Consortium (UEC), which aims to enhance Ethernet performance. One of UEC's primary objectives is to develop specifications for a new transport protocol known as Ultra Ethernet Transport (UET). The 1.0 specification for UET was released in June 2025.

UET features Remote Direct Memory Access (RDMA), enabling data to be delivered directly from the network into application memory. In addition to transport protocols, UEC is also focused on advancing other layers of Ethernet, including the physical, link, and software layers, along with a range of related initiatives.

Within data centers, connectivity between server racks is a critical bottleneck. High-speed links (e.g., 400 Gbps Ethernet or NVIDIA's InfiniBand) enable rack-to-rack communication but consume 10%-15% of a data center's power budget. This drives up operational costs, impacting RoI for data center operators but creating opportunities for investors in energy-efficient connectivity solutions.

To round the picture with a technology perspective, two primary interconnect technologies dominate AI data centers: Ethernet, valued for its flexibility and ubiquity, and InfiniBand, prized for its ultra-low latency and optimized fabric for GPU clustering. Both of these protocols are continually pushing bandwidth higher and latency lower.

Market Response. As connectivity evolves along certain trends to meet the demands of AI, these trends are providing new opportunities for investors in digital

infrastructure assets. These technical and economic drivers are catalyzing several key market trends that define the investment landscape for AI connectivity:

1. **Surging Demand for Dark Fiber and Wavelength Services:** AI infrastructure prioritizes dark fiber and high-capacity wavelengths (e.g., 400/800 Gbps) to connect data centers to cloud locations and training data sources, due to their scalability, flexibility and control. Dark fiber's appeal lies in its ability to future-proof capacity as it offers ultimate control over the light signals (wavelengths), allowing for the use of the latest coherent optical technology (e.g., 800ZR/+ plugs) for massive capacity without being locked into a carrier's equipment upgrade cycle. Companies that deploy and sell dark fiber and wavelengths services are benefiting from strong demand for these products. For instance, Zayo reports metro dark fiber demand grew 268% year-over-year in 2024 and wavelength demand grew by 180% between 2020 - 2024. Lumen and AT&T in the US and Interxion (Digital Realty), Colt Technology Services, and Deutsche Telekom in Europe are examples of companies that are benefiting from this trend.
2. **Hyperscaler-Driven Capacity Growth:** with AI training and inference requiring massive data transfers (e.g., petabytes daily), much of that demand for dark fiber and wavelength came from hyperscalers and carriers who between 2020 – 2024 accounted for 91.2% of metro dark fiber purchases and 66.8% of all wavelength deals exceeding 1 Tb of capacity. Software and tech companies also boosted purchases by 450% from 2020 to 2024 (9.6 Tb to 52.12 Tb). This demand translates to high-margin revenue streams for fiber providers, with contracts often spanning 5-10 years. Investors can capitalize on this by funding scalable terrestrial and submarine fiber networks, particularly those supporting intra-data center connectivity.

AI Networking Requirements

AI model training at scale relies on distributed computing with clusters of interconnected GPUs. The characteristics of training workloads place specific demands on interconnectivity across the entire stack, including GPUs, servers, racks, and data centers. These workloads require sustained high bandwidth for data transfer, low latency for synchronization, efficient handling of both large tensors and small control messages, and consistent performance without jitter or packet loss. These requirements have led to the development of a new class of networking solutions optimized for AI workloads, including NVLink, which is used for GPU-to-GPU connectivity, InfiniBand, which is used for server-to-server and rack-to-rack connectivity within clusters, and advanced Ethernet fabrics for larger scale data center networking and interconnect.

As an example, training GPT-4 class models with more than 100 billion parameters requires massive bandwidth, both for on-GPU memory access and for GPU-to-GPU tensor exchange. Each Nvidia H100 GPU provides over 3 TB/s of HBM3 memory bandwidth, and GPU clusters collectively move tens of terabytes per second during training. The fifth-generation NVLink, introduced with the Blackwell architecture in 2024, provides up to 1.8 TB/s of bidirectional bandwidth per GPU, with cluster-level fabrics scaling to well over 100 TB/s of aggregate bandwidth. For server-to-server and rack-level communication, the Nvidia Quantum-2 InfiniBand switch delivers 64 ports at 400 Gbps each, with a total switching capacity of 51.2 Tb/s.

At larger scale, Ethernet-based fabrics such as platforms from Arista, Cisco and Nvidia Spectrum-X are designed to support millions of GPUs within a data center and extend connectivity across data centers.

3. **New Corridors in Non-Traditional Markets:** AI data centers prioritize low-

cost power and land, shifting connectivity demand to low-population regions and creating new opportunities for long-haul fiber. For example, Zayo reports the highest purchased capacity being along the San Jose, CA to Pendleton, OR, a town of under 20,000 that emerged as an AI data center hub due to abundant space and power. Similarly, new routes are gaining tractions, such as Chicago to Columbus and Atlanta to Ashburn. Thus, demand for long-haul fiber is expected to grow particularly along these new routes.

4. Demand for higher bandwidth:

Connectivity customers are increasingly relying on higher bandwidth, primarily today's workhorse 400 Gbps, which is outpacing demand for lower services. AI workload requirements for high-throughput and low latency are driving the deployment of 800 Gbps, considered state of the art today, particularly in backend data center connectivity of GPU clusters for training and inference where synchronized computation is essential. Thus, both inter-data center connectivity in metro and regional networks, and intra-data center connectivity are increasingly demanding 800 Gbps services.

Challenges in AI Connectivity. Deploying connectivity infrastructure for AI data centers presents significant challenges due to the scale and intensity of AI workloads. Below, we outline two key hurdles and their implications for investors in fiber and broadband assets, alongside strategies to capitalize on these challenges.

1. High Capital Expenditures for Bandwidth Upgrades. Supporting 800 Gbps connectivity, critical for AI training and inference, demands substantial investment in optical transceivers (\$5,000-\$10,000 per unit), high-capacity switches, and testing tools like bit error rate testers. Upgrading a single data center's intra- and inter-data center networks can cost millions of dollars, favoring large, well-capitalized operators

and fiber providers. In addition, building new fiber routes demands substantial capital to construct physical infrastructure and to establish connections with cloud hubs, internet exchanges, and telecom points of presence.

2. Soaring Power Consumption. Beyond the power consumed by GPUs, the high-speed interconnect fabric itself (e.g., 800 Gbps Ethernet/InfiniBand) can consume between 10%-15% of the data center's power budget. This drives up operational costs and necessitates advanced, expensive cooling solutions. This challenge directly fuels investment in next-generation, energy-efficient technologies like Linear Pluggable Optics (LPOs) and co-packaged optics (CPOs).

Investment Implications. The rapid expansion of AI workloads represents a structural transformation, not a temporary trend, positioning high-performance connectivity as a vital and increasingly valuable asset class. Investors should prioritize companies located along emerging fiber corridors that connect power-rich, non-traditional markets, as well as owners of metro and long-haul dark fiber infrastructure. Equally important are firms developing advanced optical components and energy-efficient interconnect solutions that address rising cost and power constraints. Those building and operating the physical networks that enable the AI revolution are well positioned to capture durable, long-term value. In parallel,




Figure 2 Optical solutions to address power efficiency in AI-HPC.

evolving demands for high-speed interconnect—both within and between data centers—are driving renewed investment in optical and electrical components, subsystems, and networking platforms

designed to meet the next generation of performance requirements.

Key Takeaways

- **From Population-Centric to Compute-Centric Design:** AI is shifting fiber networks from urban-centric designs to remote, power-optimized locations, where data centers can scale efficiently.
- **Emergence of New Long-Haul Corridors:** AI is fueling demand for long-haul fiber to link remote data centers with global networks, opening new terrestrial and submarine fiber investment corridors beyond metro regions.
- **Intra- and Inter-Data Center Connectivity Upgraded:** AI's bandwidth and latency demands are driving major upgrades to campus and regional fiber, with high-count dark fiber paths and high-capacity wavelengths emerging as core infrastructure.
- **Power-Sharing Clusters Drive Local Fiber Demand:** AI data centers are clustering near shared power infrastructure, driving demand for high-throughput fiber links and reshaping local network design.
- **Bandwidth Inflation Is Structural, Not Cyclical:** AI workloads move petabytes daily, driving 800 Gbps fiber adoption and ongoing investment in advanced optics and infrastructure.
- **Energy-Efficient Connectivity Becomes Critical:** High-speed interconnects consume up to 15% of a data center's power budget. This is accelerating adoption of technologies like Linear Pluggable Optics (LPOs) and Co-Packaged Optics (CPOs), creating new investment opportunities in energy-efficient fiber infrastructure.
- **Dark Fiber as a Strategic Asset Class:** AI's bandwidth demands make dark fiber a premium asset, with long-term hyperscaler contracts fueling stable, high-margin returns, especially in emerging and metro markets.
- **Resurgence of venture capital investment in high-speed connectivity startups:** Rapidly evolving demands for data center interconnect—both within and between facilities—are driving renewed investor interest in startups developing optical and electrical components, subsystems, and advanced networking solutions tailored to next-generation performance requirements.

About Xona Partners

Xona Partners (Xona) is a boutique advisory services firm specializing in technology, media, and telecommunications (TMT). Established in 2012 by a team of seasoned technologists, startup founders, managing directors in global ventures, and investment advisors, Xona leverages its founders' cross-functional expertise to offer a unique, multidisciplinary approach to technology and investment advisory services. Our clientele includes private equity and venture funds, technology corporations, regulators, and public sector organizations. We assist our clients with pre-investment due diligence, post-investment lifecycle management, and strategic technology management, helping them identify new revenue streams and navigate the complex landscape of the TMT sector.

E-mail: advisors@xonapartners.com | Web: xonapartners.com


Betting on Sovereignty: Assessing Risks in the Sovereign AI Race

Overview. Sovereign AI refers to the development, deployment, and control of artificial intelligence models, infrastructure, and data by domestic actors in line with a nation's or region's security, economic, cultural, and legal interests. The term "sovereign AI" gained prominence after Jensen Huang's February 2024 speech at the World Governments Summit, though related concepts predate it, such as sovereign cloud and digital sovereignty.

In this Insight Note, we highlight key aspects of sovereign AI, including its tenets, challenges, and risk reduction factors, to help prospective investors in sovereign AI infrastructure develop a more methodical approach to evaluating these investments. A similar concept, sovereign cloud, was once positioned as critical but never fully materialized, offering valuable lessons that can inform how investors assess the viability and execution of sovereign AI initiatives.

The pillars of AI sovereignty. AI sovereignty consists of four interlocking dimensions:

1. *Data sovereignty:* Beyond data residency, this entails exclusive legal and technical control over digital information generated within national boundaries. It includes the right to govern how data is collected, used, shared, and protected, free from foreign surveillance laws such as the U.S. CLOUD Act. Data controlled by U.S. entities may still fall under the CLOUD Act even if stored domestically. To counterbalance this, jurisdictions such as the European Union have established frameworks like the General Data



Worldwide. 11/22/24 - 11/22/25. Web Search.

Figure 1 Google Trends interest for "Sovereign AI".

- Protection Regulation (GDPR) and the Data Act, which explicitly restrict unauthorized access to EU data and reinforce local jurisdiction.
2. *Operational sovereignty:* AI infrastructure, including data centers, training platforms, and inference systems, is nationally owned, operated, and managed. This ensures that critical decisions about uptime, access, and maintenance remain with domestic entities.
3. *Technical sovereignty:* Control extends to the hardware and software stack, including semiconductors, cloud infrastructure, foundational models, and development tools.
4. *Legal sovereignty:* AI systems must comply exclusively with national or regional laws, including data protection, AI ethics, and liability frameworks. The aim is to prevent extraterritorial legal orders, such as a U.S. court compelling a U.S. AI firm to hand over third-party national AI training data, from overriding domestic jurisdiction.

Additional dimensions of AI sovereignty include governance, ethics and culture,

economics, and innovation. This four-layer model prioritizes the core technical, operational, and legal elements most relevant to infrastructure investments. Together, these dimensions provide criteria to assess the efficacy of sovereign AI initiatives and to compare them across contexts.

In essence, sovereign AI seeks to reduce strategic dependence on a small number of foreign governments, primarily the United States and China, and corporations that dominate the global AI technology stack. Given the interconnectedness of the AI stack, the level of independence will vary with the maturity of the national or regional market. This leads to different emphases in sovereign AI depending on the country or region, which we address later in this Note.

Drivers and strategic purposes of sovereign AI. To characterize sovereign AI, it's important to understand the critical drivers, which include:

1. *National security and defense autonomy:* AI is increasingly vital to defense, intelligence, and critical infrastructure. Sovereign capabilities prevent adversaries and even allies from disrupting or manipulating strategic systems.
2. *Economic competitiveness and growth:* Sovereign AI catalyzes new industries, retains intellectual property domestically, and boosts productivity through localized innovation.
3. *Data sovereignty and privacy protection:* Countries and regions seek to shield their data from extraterritorial legal reach. A sovereign architecture can strengthen both technical and legal protection.
4. *Preservation of national culture, language, and values:* AI models trained on local languages, histories, and social norms produce more relevant, respectful, and representative outputs. This is especially crucial in regions who seek

The AI Technology Stack

The complexity of the AI stack creates national interdependency, scaled by the strength of each ecosystem. The United States and China, the two largest players, dominate much of the stack and come closest to end-to-end control, yet neither achieves full sovereignty. The table below presents the classification as outlined by the Brookings Institution. While we use it as a reference point, we hold reservations about elements of the structure and the entities identified within the stack. In our view, it remains a messy framework that is ripe for disruption.

Table 1 Representation of the AI technology stack. [Source: [Brookings Institution](#)]

Key firms	Key countries
Data and AI	
OpenAI, Microsoft, Google, Meta, Anthropic, Alibaba, Baidu, DeepSeek	U.S., China
Software	
Microsoft, Apple, Google, Baidu, Tencent, Alibaba, Huawei, SAP	U.S., China, Germany
Cloud	
AWS, Microsoft, Google, Alibaba	U.S., China
IoT & devices	
Apple, Google, Huawei, Xiaomi, Samsung, Siemens, Bosch	U.S., China, South Korea, Germany
Networks	
Huawei, ZTE, Nokia, Ericsson, NEC, Cisco	China, EU, Japan
Chips	
TSMC, Samsung, SK Hynix, Nvidia, AMD	U.S., Taiwan, South Korea, Netherlands
Machine tools	
ASML, DMG Mori, Yamazaki, Mazak, Trumpf, Okuma, Haas	U.S., Japan, Germany
Raw materials, energy and water	
Exxon, Chevron, China Rare Earth Group, Gazprom, Rosatom	U.S. China, Russia

protection from Western or Chinese-dominated AI.

5. **Regulatory independence:** Sovereign AI allows nations to enforce their own AI ethics, safety, and accountability regimes, such as the EU's AI Act, without being overridden by foreign providers' policies or lobbying. This holds to the extent that infrastructure and models are domestically controlled or procured from jurisdictions subject to local regulation.
6. **Geopolitical leverage:** In a world where the U.S. and China dominate most of the AI stack, from rare earths to cloud services, sovereign AI is increasingly positioned as a strategic alternative for middle powers. It avoids forced alignment and enhances diplomatic and technological agency.

Key Challenges in Building Sovereign AI.

Despite strong motivations, sovereign AI faces significant hurdles, such as:

1. **High upfront costs:** In some AI infrastructure deployments, GPU costs may approach 50%. Power, cooling, data centers, and connectivity add billions for national-scale deployments (e.g., >100MW facilities). Sovereign AI demands massive, sustained investment, and only a few countries possess the sophisticated financing ecosystems needed to raise funds at this scale. For example, the EU has proposed tens of billions for AI infrastructure. Keeping costs in check is a critical issue. Various proposals suggest shared regional compute, though no large-scale implementation exists, such as European AI gigafactories. Other approaches include regional alliances, phased rollouts focusing on high-return use cases such as healthcare, and open-source models. Open source has gained traction in large language models and beyond, not only reducing training costs but also expanding access to software and tools that lower barriers for domestic ecosystems.
2. **Talent shortages:** A lack of AI engineers, data scientists, and policy experts

Evolution of AI policies

While sovereign AI is a relatively new concept, countries have pursued AI policies for about the past decade. Early initiatives such as the Global Partnership on AI (GPAI) and similar frameworks launched 6–7 years ago addressed AI development, governance, and even sovereignty concerns, laying the groundwork for today's debates. Analysis reveals three dominant strategic imperatives: exploiting AI's potential, building capacity for AI, and governing AI. These imperatives are shared globally but differ in emphasis across countries and have evolved over time. Initial efforts focused on AI use cases and potential, but governance has surged since the emergence of generative AI and large language models in late 2022. Throughout, capacity building has remained a core imperative, underscoring its foundational role.




Figure 2 Imperative prevalence over time across national AI policies. [Source: [AI in a Comparative Perspective: Linking Policy Formulation and Implementation](#)]

hampers sovereign operations. Some regions face retention challenges, particularly when compared with U.S. compensation and scale. Addressing talent shortages requires investment in national AI institutes, university-industry partnerships, upskilling programs, and visa incentives for global experts.

3. **Data access and quality:** Sovereign models depend on high-quality, localized datasets. Past failures highlight risks from poor data, whether synthetic or unrepresentative, algorithmic bias, ethical misalignment with cultural contexts,

inadequate validation in areas such as clinical AI, and cybersecurity breaches that erode public trust. These risks are compounded by limited customer demand in many countries: without sufficient paying users, nations struggle to generate the real-world data and revenue streams needed to fund AI development. This explains why some countries have built AI data centers but lack clear pathways to utilization. Mitigation approaches include curating real-world data through national digitization projects, establishing strong governance and audit frameworks, and creating ethical review boards to ensure cultural alignment — all of which must be anchored in actual demand to be sustainable.

4. *Innovation and scalability gaps:* Global leaders innovate at speed and scale, while smaller sovereign ecosystems struggle to match performance or keep pace. Mitigation strategies include leveraging open-source models, deploying multi-cloud and hybrid architectures for interoperability, and focusing on niche excellence, such as the EU in AI safety or Japan in lithography.
5. *Geopolitical and supply chain dependencies:* Complete independence is difficult to achieve. Even the United States lacks full stack control, with TSMC manufacturing Nvidia's chips, ASML in the Netherlands supplying EUV machines, and rare earths concentrated in limited global sources. True sovereignty remains elusive, as China faces gaps in advanced lithography and electronic design automation. A key theme here is supply chain control, which advanced countries such as Japan have prioritized for more than a decade to secure critical inputs and reduce vulnerability. What is achievable is strategic interdependence, where nations mitigate single-point exposure and build leverage through diversified partnerships and resilient supply chains.

Risks in sovereign AI investments. Risks in sovereign AI investments tend to concentrate on operational, commercial, and financial realities rather than technological feasibility. The most common value-destroying issues include:

1. *Lack of customer demand:* The bigger challenge for sovereign AI is not simply underutilized compute, but the absence of sufficient domestic demand with the means to pay for and use it. In many cases, high-end GPU capacity ends up serving customers in third-party countries or companies, leaving national infrastructure underleveraged. Since accelerators depreciate over three to four years, extending useful life assumptions to five or six years can artificially inflate reported earnings. With GPU capital expenditure representing a large share of total infrastructure spend, investors must rigorously assess real demand, paying customers, workload forecasts, and binding commitments. At this stage of market maturity, relying on a "build it and they will come" thesis is speculative at best.
2. *Power availability and constraints:* Many data center builds are completed before power delivery, pricing, or interconnection capacity is secured, leaving hundreds of millions in stranded capital. Power procurement, scalability over time, and price certainty are as important as the compute itself. Investors should scrutinize power delivery timelines, firm allocation agreements, and penalty or curtailment provisions.
3. *Vendor lock-in disguised as sovereignty:* Sovereign branding can mask deep reliance on a hyperscaler, a foreign foundation model, or a proprietary orchestration stack, ultimately undermining national control. This mirrors the long standing problem in cloud, where sovereign offerings often concealed dependence on foreign providers, and the same risk now extends into AI. Red flags include limited data or model portability, APIs without

contractual exit rights, and lack of ownership or access to weights, logs, and pipelines. Without enforceable portability, sovereignty remains aspirational rather than operational.

4. *Data availability and readiness:* Without sufficiently digitized, representative, and legally usable datasets, sovereign AI cannot be trained, validated, or deployed effectively. Real, high-quality domestic data, not theoretical access, is the bottleneck. Weak data governance frameworks, fragmented custodianship, or unclear privacy regimes materially increase execution and regulatory risk.
5. *Political and budget cyclicalities:* Government support in financial, regulatory, and procurement terms can shift abruptly with elections, fiscal tightening, or geopolitical reorientation. Discretionary funding is fragile. Investors should prioritize initiatives backed by multi-year mandates, statutory obligations, or cross-party alignment rather than ministerial enthusiasm.

6. *Talent availability:* Sovereign AI requires skilled engineers and data scientists, data governance specialists, infrastructure operators, and safety and compliance professionals. Reliance on imported contractors, below-market pay structures, or the absence of national workforce development plans introduces execution risk. Operational sovereignty without human capability is illusory.

Additional market-dependent issues include governance clarity, operator accountability, refresh cycles and depreciation assumptions, and alignment of build timelines with actual demand. These factors materially influence investment outcomes. Taken together, these risks do not invalidate the sovereign AI thesis, but they demand disciplined underwriting, verifiable commitments, and a bias toward projects with demonstrated utilization, secured power, and credible long-term funding.

Key Takeaways

- **Pillars of sovereignty:** Data, operational, technical, and legal control form the foundation of sovereign AI. Together they define who owns the infrastructure, who manages it, and which laws govern its use.
- **Key drivers:** Governments pursue sovereignty to secure national defense, strengthen economic competitiveness, enforce local regulation, preserve cultural identity, and gain geopolitical leverage. These drivers create long-term demand for sovereign infrastructure.
- **Main challenges:** Sovereign AI requires massive capital, skilled talent, high-quality data, and innovation ecosystems. Global interdependencies in chips, tools, and supply chains mean no country achieves full independence.
- **Corporate dependency and sovereignty beyond governments:** Sovereignty claims often overlook reliance on global technology providers. Dependence on hyperscalers, foundation models, or proprietary stacks can undermine national control just as vendor lock-in did in cloud. True sovereignty requires reducing exposure to corporate dependencies as well as navigating government policy.

Investor Risks

- **Capital risk:** GPU and infrastructure costs are enormous, and idle or underutilized compute erodes returns. Accurate demand forecasts and binding customer commitments are critical. RoI is at risk within short timeframes due to depreciation and obsolescence, compounded by rapid technology disruptions in large language models and broader AI stacks.

- **Power risk:** Many projects stall when energy supply, pricing, or interconnection capacity is not secured. Investors must verify delivery timelines and firm allocation agreements.
- **Vendor risk:** Sovereignty claims may conceal reliance on foreign hyperscalers or proprietary stacks. Without enforceable portability, national control remains aspirational.
- **Data risk:** Weak governance, fragmented custodianship, or poor quality datasets undermine training, validation, and public trust. Real, usable domestic data is the bottleneck.
- **Political risk:** Government support can shift with elections or fiscal cycles. Multi-year mandates and statutory obligations provide more reliable backing than discretionary enthusiasm.
- **Execution risk:** Misalignment between build timelines and actual demand, refresh cycles, or operator accountability can materially affect returns.
- **Liquidity risk:** Long payback periods and uncertain exit options make sovereign AI investments less flexible than mainstream infrastructure.

About Xona Partners

Xona Partners (Xona) is a boutique advisory services firm specializing in technology, media, and telecommunications (TMT). Established in 2012 by a team of seasoned technologists, startup founders, managing directors in global ventures, and investment advisors, Xona leverages its founders' cross-functional expertise to offer a unique, multidisciplinary approach to technology and investment advisory services. Our clientele includes private equity and venture funds, technology corporations, regulators, and public sector organizations. We assist our clients with pre-investment due diligence, post-investment lifecycle management, and strategic technology management, helping them identify new revenue streams and navigate the complex landscape of the TMT sector.

E-mail: advisors@xonapartners.com | Web: xonapartners.com

Navigating Risks and Opportunities in Sovereign AI Infrastructure

Overview. Several ventures, referred to as neoclouds, have launched or are planning GPU-as-a-Service offerings to provide specialized AI infrastructure. International telecom service providers, in partnership with Nvidia, have launched 'AI Factories,' aiming to explore new revenue streams. Recently, the U.S. government proposed its 'Framework for Artificial Intelligence Diffusion,' which establishes export-control rules for advanced computing chips and AI models. This Insight Note examines the risks faced by such infrastructure operators in the context of U.S. export regulations that aim to define benchmarks and guardrails for this emerging market.

AI Export Control Rules. In the final days of the Biden administration, the U.S. Department of Commerce's Industry and Security Bureau (BIS) published the 'Framework for Artificial Intelligence Diffusion' on January 15, 2025. This framework updates and tightens rules governing the export and transfer of advanced computing processors and AI models (referred to hereafter as BIS Rules or AI Diffusion Rules interchangeably). The regulations aim to safeguard the primacy of the U.S. AI ecosystem while restricting access to U.S. AI technologies by countries such as China and Russia. Notably, the framework extends beyond hardware export limits by introducing controls on AI model weights—a novel step in regulating the proliferation of AI technologies. For the purposes of this Note, we focus primarily on the hardware controls, which are outlined.

Three Tiers of Countries. The BIS AI diffusion rules divide the world into three tiers of countries:

- Tier 1: Consists of 18 countries that are close allies of the United States, with nearly unrestricted access to advanced GPUs.
- Tier 2: Encompasses most countries outside Tier 1 and Tier 3, such as Mexico, Brazil, India, Malaysia, Indonesia, and Gulf States. Exporting GPUs to these nations is permitted but subject to limits on computational capabilities and quantities.
- Tier 3: Includes embargoed states like China and Russia, which are completely prohibited from accessing GPUs.

BRICS countries are classified as either Tier 2 or Tier 3 under the BIS AI Diffusion Rules. Tier 2 countries, such as India, Brazil, and South Africa, are particularly noteworthy due to their limited access to advanced AI technologies, governed by restrictions on computational capabilities and export quantities. Other Tier 2 countries of interest




Figure 1 Country tiers under BIS Rules—Green: Tier 1; Orange: Tier 2; and Red: Tier 3.

include Singapore, known as a hub for data centers, and Malaysia, which has often served as a destination for overflow from Singapore's data center demand.

Quantity Limits. Enterprises in Tier 2 countries have three primary pathways to obtain export licenses for GPUs:

- **License Exception Low Processing Performance (LPP):** Allows access to the equivalent of 1,700 H100 GPUs per year¹. This provides the simplest route to accessing GPUs.
- **National Validated End User license (NVEU):** Grants data centers access to the cumulative equivalent of 100,000 H100 GPUs in 2025, 270,000 in 2026, and 320,000 by the end of 2027. However, obtaining NVEU status is an involved process with strict requirements for applicants. [BIS recommends securing government-to-government assurances between the U.S. and the host country before applying for NVEU authorization.]
- **Export License:** Permits Tier 2 countries to import up to the equivalent of 50,000 H100 GPUs between 2025 and 2027 inclusive.

Additionally, a second class of data centers—Universal Validated End User (UVEU), located in Tier 1 countries—face restrictions on deploying GPUs in Tier 2 countries. A UVEU holder cannot deploy more than 25% of their compute power in Tier 2 countries, with a maximum of 7% of total compute power allocated to a single Tier 2 country. Moreover, U.S. companies must maintain at least 50% of their compute power within the United States.

The rise of 'AI Factories' has been driven by the advent of Generative AI and Large Language Models, the maturation of the AI technology ecosystem, and Nvidia's aggressive expansion into the data center GPU market. Governments see AI Factories as key enablers of sovereign AI. For example, the European Commission has identified their establishment as a strategic priority, aiming to bring 15 AI Factories into operation

by 2026 and allocating a €20 billion fund to create 5 AI Gigafactories.

Many telecom service providers have adopted AI Factories as a strategic response to counter stagnating revenues. While this Note examines Telco AI Factories as examples, the insights apply broadly to other organizations. Leading adopters of this approach include Japanese and South Korean operators such as SoftBank and SK Telecom, which were among the pioneers in implementing these business strategies.

Telcos operating AI data centers in Tier 2 countries, such as Vietnam, Singapore, India, and the Gulf States, will face limits on the quantities of GPUs they can acquire under the BIS AI Diffusion Rules once they come into effect. While their current deployments

Table 1 Telecom service providers engaged with Nvidia on AI infrastructure projects.

Service Provider	Country	BIS Tier	GPU
FPT	Vietnam	2	Hopper
Indosat (IOH)	Indonesia	2	Hopper
KDDI	Japan	1	Blackwell
Ooredoo	Qatar	2	N/A
Reliance Jio	India	2	Hopper
Scaleway (Iliad)	France	1	Hopper
Singtel	Singapore	2	Hopper
SK Telecom	South Korea	1	Hopper, Blackwell (2025)
SoftBank	Japan	1	Blackwell
Telenor	Norway	1	N/A
Telus	Canada	1	Hopper, Blackwell
Viettel	Vietnam	2	N/A
YTL	Malaysia	2	Blackwell

¹ Benchmark GPU is the Nvidia H100 SXM with peak FP8 Tensor Core processing power of 1,978.9 TFLOPS.

are largely below the thresholds allowed, they face a ceiling that restricts the computational power they can deploy in the future (many such AI Factories start with a pilot of ~1,000 GPUs). To illustrate, consider FPT, which announced a \$200 million financial commitment. If all this was allocated to GPUs at a price of \$25,000 per H100 GPU (est. April 2025 pricing), this would be the equivalent of 8,000 units, which is possible under the country allocation for Vietnam. On the other hand, while Reliance has not disclosed the scope of its current deployment, it has announced plans to expand its data center capacity to 2,000 MW. This could theoretically support over a million Hopper GPUs, far exceeding the computational power allowed under the AI Diffusion Rules. For contrast, Elon Musk's xAI cluster, comprising 100,000 H200 GPUs, is currently the largest GPU cluster in the world.

The AI Diffusion rules, while not restricting current deployments, are expected to limit the ability of Tier 2 countries to train large language models. Effectively, these rules aim to preserve the lead that the United States holds in training LLMs. Currently, most AI processing power is directed toward model training, which would widen the divide between Tier 1 countries and all others over time.

Keeping Pace with Computational Power.

While AI Factories currently operate within the limits set by BIS Rules, future constraints are likely to pose challenges. The computational power of GPUs is advancing rapidly, with Blackwell GPUs delivering approximately 5x the Total Processing Power (TPP) of an H100 GPU as defined under BIS rules. As a result, export quotas are expected to shrink from the equivalent of ~50,000 H100 units to fewer than 10,000 Blackwell GPUs. This tightening will widen the gap between Tier 2-based companies and Tier 1-based companies operating in Tier 2 countries.

Investment Risks and Opportunities

Export controls are not the only risk facing AI Factories. One of the most significant risks is financial, stemming from the unpredictable timeline for enterprise AI adoption and the challenge of monetizing consumer AI services and applications. At this early stage of AI adoption, revenue streams remain uncertain. Coupled with high capital expenditure requirements, AI Factories face substantial financial risks.

AI Factories carve out a niche by addressing data sovereignty and security requirements for national governments and large enterprises. Currently, model training serves as the primary revenue driver, making it particularly susceptible to the impact of export controls. Inference, while still underdeveloped, holds significant potential for telcos in particular to capitalize on as the market matures.

To mitigate financial risks, investments are approached cautiously, with many AI Factories starting with a modest deployment of GPUs—typically between 1,000 and 2,000 units. SoftBank, which boasts one of the highest computational capacities with approximately 20,000 GPUs (including H100 and some B100 units), primarily utilizes this capacity for internal use cases, such as training its telco LLM. Telenor, an early adopter of the concept, has secured four customers within a year, while Telus, which entered the market just six weeks after exploring the idea, is still working to establish its position.

Another significant risk is the rapid evolution of AI chip technology, which drastically increases computational capacity while quickly rendering older GPUs obsolete. This steep depreciation rate for AI hardware necessitates a rapid revenue ramp-up to offset costs. However, accelerating revenues presents its own challenges. For instance, in the enterprise segment, one critical gating factor is accessing, cleaning, and organizing data for AI processes—a time-consuming and expensive endeavor.




Figure 2 The rise in computational power across successive GPU generations.

This makes it critical to monitor the future evolution of the BIS Rules for their impact on global competitive AI dynamics. The Rules are currently under consultation and are expected to come into effect in mid-May 2025. Eventually, they will need to address limits that extend beyond 2027. Meanwhile, technology in this domain is evolving rapidly. For example, Nvidia's next-generation VR300 NVL144 system is projected to deliver 5× the performance of the current GB200 NVL72 system—even while using the same number of physical GPU and CPU dies.

Critical Questions for Investors and Operators of AI Factories

The rise of neoclouds and AI Factories has sparked critical business and technical questions that operators and investors in these ventures must address, such as:

- **Defining Strategy:** *What is the optimal business approach for each layer of the AI stack?*
- **Building Differentiation:** *How can the ecosystem be designed to stand out against competitors while aligning with specific operating models?*
- **Ecosystem Monetization:** *How can operators monetize their ecosystem without conflicting with the interests of partners and direct customers?*
- **Future-proofing:** *How can evolving requirements over a 5+ year horizon be anticipated and incorporated into adaptable architecture?*
- **Deployment Models:** *Which AI and data center deployment model—hyperscale availability zones, local zones, edge computing, AI training DCs, inference DCs, AI Factory DCs, or hybrid models—best suits the business objectives?*
- **Navigating Disruption:** *How will ongoing disruption factors like geopolitics, energy dynamics, resiliency policies, and semiconductor competition impact these ventures?*

Key Takeaways

- **Regulatory Quantity Constraints:** AI Factories—including those run by telcos in Tier 2 countries (e.g., Vietnam, Singapore, India, and the Gulf States) as defined under BIS AI Diffusion Rules—face a cap on the quantities of GPUs they can acquire. While this cap does not impact current operations, it will limit their future computational expansion and ability to train large language models.
- **Shrinking Quantities and Growing Divide:** As GPU performance scales up, export quotas could shrink dramatically. For example, Blackwell GPUs, which offer roughly 5x the total processing power (TPP) of H100 GPUs, may reduce the export cap from the equivalent of around 50,000 H100 units to fewer than 10,000 Blackwell GPUs, creating significant constraints for future deployments. Therefore, a growing divide is expected between Tier 1 and Tier 2 countries.
- **Risks in Strategic Expansion Plans:** While current AI Factories, including telco ones, in Tier 2 countries remain within allowed thresholds, ambitious projects like Reliance's planned 2,000 MW expansion have the potential to support computational capacities far exceeding the limits imposed by BIS rules.
- **Opportunity in AI Inference:** Model training remains the critical activity for installed GPU clouds, with AI Factories leveraging data privacy, residency, and security to serve national governments and enterprises. Although AI inference remains lagging, telcos are well positioned to play a role in this segment as the market evolves.
- **Global Competition and Opportunity:** The tightening regulatory environment highlights shifts in global competitive dynamics, where unrestricted access for Tier 1 entities could further challenge Tier 2 ventures in building locally sovereign AI infrastructures.

About Xona Partners

Xona Partners (Xona) is a boutique advisory services firm specializing in technology, media, and telecommunications (TMT). Established in 2012 by a team of seasoned technologists, startup founders, managing directors in global ventures, and investment advisors, Xona leverages its founders' cross-functional expertise to offer a unique, multidisciplinary approach to technology and investment advisory services. Our clientele includes private equity and venture funds, technology corporations, regulators, and public sector organizations. We assist our clients with pre-investment due diligence, post-investment lifecycle management, and strategic technology management, helping them identify new revenue streams and navigate the complex landscape of the TMT sector.

E-mail: advisors@xonapartners.com | Web: xonapartners.com

AI Infrastructure at a Crossroads: Navigating Volatility and Opportunity

Overview. Markets and investors are growing increasingly uneasy about AI investments, as many now associate it with a bubble. Several prominent industry figures, including [Sam Altman](#), [Jeff Bezos](#) and [Mark Zuckerberg](#), have spoken publicly about this possibility. Public market valuations remain elevated, driven by the performance of the Magnificent 7. Private markets show a similar trend, with companies like OpenAI attracting historically high valuations. Economist Steve Hanke presented [Dr. X's bubble detector](#), which now stands at an all-time high. Torsten Sløk, Chief Economist at Apollo Global Management, argued that the [AI bubble](#) exceeds the dotcom bubble, noting that the top 10 stocks in the S&P 500 now account for a much larger share of the index's total market capitalization than they did in 2000.

Against this backdrop, we wrote this Insight Note to share our perspective on the AI infrastructure landscape. Given the complexity of this space, where technology, finance, and geopolitics converge, we aim to offer a high-level map that invites deeper exploration of specific themes. A few observations to set the landscape and outline key topics:

1. The rush to build AI infrastructure has fueled the rise of neoclouds offering GPU-as-a-Service. Hyperscalers such as Microsoft, Google, and Meta, along with AI-first firms like OpenAI, xAI, and Anthropic, are racing to develop more powerful Large Language Models for Generative AI. To meet growing compute needs, they increasingly offload demand to neoclouds, which benefit from current spending trends.

2. Despite strong momentum, neoclouds remain a fragile link in the AI value chain. Any slowdown in infrastructure demand is likely to surface first in neocloud services, exposing them to early signs of market contraction.
3. The holy grail of AI revenue depends on enterprise adoption, which takes time to materialize. Historical patterns from past technology cycles, such as the Internet buildout of the 1990s and early 2000s, suggest a mismatch between capital investment and revenue realization. While consumer AI is active, its monetization follows a different path, often through advertising or subscription models.
4. Steep depreciation of compute assets amplifies sensitivity to the revenue ramp.




Figure 1 AI vs. dotcom market. [Source: Bloomberg, Apollo Chief Economist]




Figure 2 S&P 500 and Dr. X's Bubble Detector. [Source: Prof. Steve Hanke, Fortune Magazine]

This creates a challenge for investors whose long-term objectives may not align with the short investment cycle of AI compute infrastructure, even as broader infrastructure elements such as fiber connectivity, power, and real estate follow longer timelines.

5. AI firms feel compelled to invest in infrastructure given the strategic stakes. This urgency reinforces concerns about a potential bubble. A tense geopolitical climate, where AI is treated as a national priority, further extends the runway for continued capital deployment.

In this Insight Note, we explore foundational issues to offer a perspective on the evolving AI ecosystem and to highlight both risks and potential mitigation strategies. While some concerns are valid, the outlook is far from uniform. Strategic investment opportunities continue to emerge across the AI value chain. This note presents an initial view that we will build upon in future Insight Notes.

Neoclouds and risk posture. Companies like CoreWeave, Nebius, Lambda, and Crusoe have become key pillars of the AI infrastructure landscape by providing compute capacity to absorb overflow demand from hyperscalers.

Take CoreWeave, for example. As a public company, it reports rising revenues and anticipates higher capital expenditures as operations scale. A closer look at its revenue reveals a high level of concentration, with significant contributions from hyperscalers such as Microsoft, Meta, Google, AWS, and large AI-first companies such as OpenAI. This concentration is not easing over time; it is intensifying. CoreWeave's 2024 revenue was approximately 87% concentrated among its top three customers, with Microsoft alone contributing around 62%.

To stabilize revenue, CoreWeave and other neoclouds enter into multi-year, customized take or pay agreements with their clients. These contracts typically include reserved




Figure 3 CoreWeave's customer concentration. [Source: Company data, S1 SEC filing]

Table 1 CoreWeave financial highlights (\$, millions)

Year	2022	2023	2024	2025 (1H)
Revenue	15.8	228.9	1,915.4	2,194
Capex	72	2,943	8,702	3,860
Net Loss	31	593.7	863.4	605

capacity, minimum commitments, termination clauses, and other provisions designed to mitigate revenue volatility. While the terms vary by customer, the primary goal is to secure predictable revenue over a period that aligns with the useful life of the compute infrastructure, often between two and four years. These agreements often include opt-out clauses that allow hyperscalers to reduce utilization, though usually with penalties. In effect, these structures shift the risk of a demand slowdown onto the neoclouds.

It is no surprise, then, that many now refer to neoclouds as the “shock absorbers” of the AI ecosystem. Hyperscalers continue to offload excess demand onto neocloud infrastructure while expanding their own capacity. If demand weakens, they are likely to reclaim that overflow, passing the impact of the downturn to neoclouds in a way that preserves their own flexibility while exposing their partners to greater risk.

Financing structures. Whereas equity dominated early AI infrastructure funding, recent projects rely more heavily on debt. Recent examples include capital raised by Oracle, xAI, CoreWeave and Vantage Data Centers, among others. To finance its massive 4 million square foot, 2 gigawatt Hyperion data center, Meta formed a joint venture with Blue Owl Capital, backed by \$27 billion in equity and debt. Meta will operate the facility and retain a 20% ownership stake. This structure helps Meta protect its balance sheet and credit rating, enabling continued access to lower cost debt. It also shifts a significant portion of the risk to investors.

The debt tied to the Hyperion project reportedly matures in 2049, a timeline designed to appeal to long term institutional lenders. However, it exposes investors to revenue and cash flow risks tied to fluctuations in AI demand, along with other challenges such as construction delays and power provisioning issues.

AI compute infrastructure, which can account for up to half the total infrastructure spending, depreciates quickly, often over a period of three to five years. Other IT components, such as networking infrastructure, depreciate more slowly, as do power plants and transmission systems. Buildings depreciate over several decades. The structure of financing, combined with the timing of expenditures, depreciation, and revenue realization, will directly affect return on investment.

High debt loads, customer concentration, slow revenue growth, and fast-moving tech cycles weaken the foundation of the AI infrastructure ecosystem. Investors must challenge assumptions and validate resilience before committing capital.

Circular capital and contagion risk. AI companies increasingly invest in one another, creating interdependencies that heighten risk across the ecosystem. A

Critical questions in AI infrastructure

Elements to highlight in AI infrastructure due diligence include:

- **Market and competitive assessment:** Understand AI-specific application requirements, evaluate supply dynamics over time, and identify partnership opportunities.
- **Technical and infrastructure assessment:** Evaluate asset-level factors including connectivity, power, cooling, and M&E systems.
- **Operational readiness:** Review the organizational structure across operations, engineering, and client support. Examine onboarding processes to identify gaps that may impact uptime and scalability. Assess differentiation in execution, especially in deploying the various layers of the AI stack effectively.
- **Commercial assessment:** Analyze pricing strategy, unit economics, capex and opex assumptions, revenue ramp and sensitivity, and the sales and go-to-market model.
- **Regulatory and permitting assessment:** Examine government incentives for land, power, and tax. Assess data residency, privacy requirements, and sector-specific regulatory constraints.
- **Risk focus and gap analysis:** Conduct a consolidated gap analysis between market needs and service capabilities. Identify high-impact risks across technical, operational, commercial, and regulatory dimensions.
- **Interconnection and ecosystem integration:** Evaluate the prospects for dense interconnection; access to carrier-neutral facilities; cloud on-ramps; and AI workload exchanges. Consider strategic partners, peering arrangements, and

prominent example is Nvidia's plan to invest \$100 billion in OpenAI, staged across deployments of at least 10 gigawatts of AI data centers to support millions of Nvidia GPUs. Other notable moves include Nvidia's 7% equity stake in CoreWeave and a \$2




Figure 4 Map of recent AI investments. [Source: Morgan Stanley]

billion investment in xAI. Morgan Stanley's chart on this topic now circulates widely.

Circular financing offers some benefits. It channels capital into critical players that need funding and helps accelerate infrastructure development. However, it also introduces serious risks often linked to market bubbles. This type of financing can obscure true demand for AI by creating the illusion of greater market traction than actually exists. It also weakens ecosystem resilience and amplifies contagion risk due to the tight interconnections among participants.

One additional point on Nvidia, though other firms follow a similar strategy. Hyperscalers such as Google, Microsoft, and AWS continue to design their own AI chips. Google's TPU is a clear example, and OpenAI's collaboration with Broadcom reinforces the trend. Nvidia benefits when companies like OpenAI and CoreWeave succeed and maintain influence. Their growth helps protect Nvidia's future revenue stream from potential declines in demand from the hyperscalers.

Enterprise adoption and revenue timing. AI revenues depend heavily on enterprise adoption, a process that many expect will take years to fully materialize based on historical comparisons. AI introduces a fundamental shift in how enterprises manage operations and conduct business. It demands proprietary solutions that involve curating internal data and building tailored

AI compute power and efficiency

AI hardware has seen a roughly 1,000-fold increase in computational power over the eight years between Nvidia's Pascal GPU in 2016 and Blackwell in 2024. Nvidia plans to launch the Rubin GPU in the second half of 2026, delivering up to 50 petaflops of FP4 performance. This doubles the performance of Blackwell and reflects an average annual increase of about 200 times in compute capability.

This leap in computational power comes with a rise in absolute power consumption per chip, although overall power efficiency has improved. For example, the Blackwell GPU consumes between 1,200 and 1,400 watts, compared to approximately 700 watts for Hopper. This results in higher power density per rack and drives the need for more advanced cooling systems to manage the heat generated by the GPUs.




Figure 5 Nvidia compute and efficiency trends. [Source: Nvidia]

models. This requires significant investment and time, and most enterprises remain in the early stages of AI deployment.

A July 2025 study by MIT's Project NADA (Networked Agents and Decentralized Architecture) highlighted this gap. Enterprises spent between \$30 and \$40 billion on generative AI, yet 95% reported no measurable return. While 80% piloted tools like ChatGPT and Copilot, and 40% moved to deployment, these tools improved individual productivity but failed to deliver meaningful impact on profit and loss. This does not suggest AI lacks value, but rather that returns will take more time to materialize.

Still, AI-first organizations are growing their revenues. The Information reports that AI-Native applications reached \$18.5 billion in annualized revenue, with OpenAI and Anthropic capturing 88% of that. These companies have accelerated early revenue growth faster than their SaaS counterparts. However, the critical issue remains the pace of enterprise revenue growth compared to infrastructure capital expenditure. OpenAI, for instance, is not expected to turn a profit until 2029, when annual revenue is projected to reach \$125 billion according to internal forecasts shared with investors and reported by Bloomberg and The Information.

Strategic pressure and geopolitics. Despite the risks tied to AI investments, ecosystem players appear to have little choice but to engage. The foundational transformation that AI promises offers major gains for the winners and threatens the losers with irrelevance (i.e. winner takes all model).

Adding Up AI Revenues

OpenAI, Anthropic and 16 other 'AI native' startups are cumulatively generating \$18.5 billion in annualized revenue, with the top two firms capturing 88% of that amount.




Figure 7 AI LLM revenues. [Source: The Information]

Jevon's paradox and persistent GPU demand

In the 19th century, economist William Jevons observed that coal consumption increased even as technology made coal use more efficient. A similar paradox now applies to AI infrastructure, where demand for the latest generation of GPUs remains insatiable. This stems from a relentless pursuit of performance in software applications.

First, model training costs have climbed over time as developers train larger models using more advanced hardware.

Second, AI application builders consistently choose frontier models over cheaper and faster alternatives. They prioritize performance and willingly pay for it, according to data from Menlo Ventures. Even when individual models drop tenfold in price, developers do not capture savings by reverting to older models. Instead, they shift collectively to the highest performing option available.




Figure 6 Estimated training cost of select AI models 2016-2024. [Source: Epoch AI, Stanford University Artificial Intelligence Index Report 2025]

Consider the competition between ChatGPT and Google Search. By August 2025, around 700 million weekly users gave ChatGPT approximately 63% of the AI search market, according to Semrush data. While market share estimates vary by methodology, Semrush also reported that Google's global search traffic, though still holding 90% of the market, declined by 7.9% between 2023, the

year after ChatGPT launched, and 2024, making the first meaningful drop in decades.

Moreover, studies show that AI Overviews and AI Mode within Google Search reduce overall click-through rates on traditional search results. A [Pew Research](#) study found that when users encountered an AI summary, clicks on traditional search links fell from 15% to 8%, which directly impacts ad performance. The rivalry between OpenAI and Google reached a peak just days before this Note was published, when OpenAI launched its [Atlas web browser](#) with ChatGPT at its core.

AI has become a necessary arena for hyperscalers, cloud service providers, and technology companies. At the national level, governments now treat AI as a matter of national security in a world shaped by geopolitical tension. With power emerging as a key bottleneck in AI development, authorities have begun easing environmental restrictions and accelerating the permitting process to fast-track data center construction.

Summary and risk outlook. AI infrastructure presents both a generational opportunity and a classic capital cycle challenge. These buildouts require large, uneven investments that must anticipate demand still in flux. The immediate risk does not stem from a single point of failure but from a convergence of uncertainties. These include unpredictable enterprise demand, a concentrated customer base, a fast-moving innovation cycle, rapid hardware obsolescence, customized financing structures, and circular capital flows that can magnify shocks.

Investors should recognize that AI infrastructure investments fundamentally involve data center real estate, compute, power, and connectivity. A market bubble would affect each component in distinct ways. To navigate this dynamic landscape, investors must apply rigorous due diligence, build resilient scenarios, and maintain

The global-local divide in AI

AI dynamics differ depending on where in the world one looks. The foundational nature of AI and its impact of national security and economic competitiveness led to "local" imperatives shaping its trajectory, in contrast to the global nature embodied in the expansive reach of the hyperscalers and AI-first companies. This duality has fuelled the ascent of sovereign AI initiatives, most vividly illustrated by the European Union's €200 billion InventAI Joint Undertaking, which channel billions into bolstering Europe's AI ecosystem and infrastructure against external dependencies.

China stands out as a paragon of this localization, having cultivated a sovereign cloud ecosystem through its tech titans—Alibaba, Baidu, ByteDance, Tencent and Huawei. While Huawei positions itself as a direct counterpart to Nvidia in the U.S. arena, cloud leaders like Alibaba and Baidu are forging their own AI compute pathways, mirroring the custom silicon pursuits of American peers such as Google with its TPUs. Complementing this, a cadre of Chinese firms, including SenseTime and Moonshot AI, are pioneering large language models like DeepSeek, Baidu's Ernie Bot and Alibaba's Qwen, which predominantly build on open-source foundations—a sharp contrast to the U.S. emphasis on proprietary architectures.

To encapsulate the landscape, global investors such as sovereign wealth funds see opportunity in bridging global and local dynamics. Singapore's Temasek's stake in Alibaba, Norway's Norges Bank's AI consortia push, and Gulf investments in data centers reflect this. Yet, localization efforts through mandates and incentives have yielded limited results.

In short, while the world seeks to democratize AI across countries and ecosystem players, control remains concentrated in the hands of a few large entities. This concentration significantly increases risk.

discipline. The major winners will be those who deliver clear differentiation through

superior execution, phased scaling, and the ability to adapt as AI evolves.

Key Takeaways

- **Rising bubble talk warrants caution, but investors should balance risk management with opportunity.** Apply disciplined due diligence and scenario planning to protect capital while acknowledging durable demand for AI workloads, structural advantages for best-in-class operators, and near-term chances to capture market share through superior execution, phased scaling, and targeted partnerships.
- **AI infrastructure is a composite investment in data center real estate, compute, power, and connectivity, with each component exposed differently in a market correction.** Investors must run rigorous due diligence, stress-test downside scenarios, and enforce capital discipline; advantage will go to operators who execute cleanly, scale in phases, and adapt as AI workloads evolve.
- **Strategic imperative fuels continued investment.** AI firms and governments treat AI as mission critical, so firms feel compelled to invest heavily in infrastructure; that urgency reinforces bubble narratives while geopolitical priorities extend the runway for ongoing capital deployment.
- **Neoclouds are the potential weak link.** Neoclouds offering GPU-as-a-Service capture current AI spend but will likely show demand pullbacks first. They act as a buffer for hyperscalers and large AI-first companies, concentrating downside risk for investors.
- **Enterprise adoption lag creates a capex versus revenue mismatch.** Enterprises face long, costly implementation cycles to curate data and build models, so infrastructure capex outpaces realized enterprise revenue today; the problem is amplified by the short depreciation cycle of AI compute assets.
- **Compute advances drive infrastructure stress and demand persistence.** Rapid GPU performance gains raise rack power density and cooling needs, and app builders continually choose frontier models, sustaining demand despite improved efficiency (Jevon's Paradox).
- **Global-local divide shapes AI infrastructure strategy.** National security and economic imperatives drive sovereign initiatives that contrast with the global dominance of hyperscalers and AI-first firms. Despite efforts to democratize AI, these initiatives have yet to deliver meaningful results.

About Xona Partners

Xona Partners (Xona) is a boutique advisory services firm specializing in technology, media, and telecommunications (TMT). Established in 2012 by a team of seasoned technologists, startup founders, managing directors in global ventures, and investment advisors, Xona leverages its founders' cross-functional expertise to offer a unique, multidisciplinary approach to technology and investment advisory services. Our clientele includes private equity and venture funds, technology corporations, regulators, and public sector organizations. We assist our clients with pre-investment due diligence, post-investment lifecycle management, and strategic technology management, helping them identify new revenue streams and navigate the complex landscape of the TMT sector.

E-mail: advisors@xonapartners.com | Web: xonapartners.com

Part 2:

Non-Terrestrial Networks

Orbiting on Demand: How GSaaS is Rewiring the Satellite Industry's Ground Game

Overview. Strategic interest in Ground Stations as a Service (GSaaS) is rising, driven by two counteracting trends in the satellite industry. The rapid proliferation of low Earth orbit (LEO) satellites is fueling demand for flexible, distributed ground station networks that can support dynamic, high-throughput connectivity. At the same time, financial pressures on geostationary orbit (GEO) operators are prompting them to monetize existing infrastructure through GSaaS offerings. This Insight Note outlines the GSaaS landscape and highlights key considerations for investors evaluating opportunities in this evolving segment.

Origins and Market Formation. The GSaaS market has evolved through distinct phases, shaped by satellite deployment trends and advances in infrastructure and software. Commercial ground station services began to emerge in the late 1990s, driven by a surge in planned LEO satellite constellations that required scalable global infrastructure. Universal Space Network (USN), founded in 1996, was among the earliest to offer flexible pay as you go telemetry, tracking, and command services, launching operations in 1998 with stations in Alaska and Hawaii. This model laid the foundation for GSaaS. In 2009, the Swedish Space Corporation (SSC) acquired USN and expanded its managed ground station services globally. Kongsberg Satellite Services (KSAT), established in 2002 as a joint venture in Norway, also scaled its offerings to support hundreds of low Earth orbit missions.

The mid 2010s marked a new phase with the rise of software-defined and cloud-integrated GSaaS startups. These providers differentiated themselves through platforms

The GSaaS Investment Case

GSaaS allows satellite operators to access ground infrastructure on a flexible and on-demand basis. Instead of building and maintaining proprietary stations, operators lease antenna time to communicate with satellites over designated frequency bands. These services typically include fiber connectivity for data transport and are governed by service level agreements that ensure reliability and uptime.

The GSaaS model resembles the telecom tower business in its structure and economics. Providers invest capital upfront to deploy antennas, radio systems, and connectivity, then monetize those assets through shared usage. Revenue is generated through recurring contracts, often priced per satellite pass, per gigabyte, or through subscription tiers. This creates predictable cash flow anchored in physical infrastructure.

Private equity and infrastructure funds are increasingly drawn to GSaaS because it combines capital intensity with recurring revenue. The model supports high utilization through dynamic scheduling and multi-mission support. It benefits from elastic demand across Earth observation, IoT, and remote sensing missions. Relatively long satellite lifecycles and mission-critical data flows reduce churn risk and strengthen contract durability. Given the scale of these investors, carve-outs from incumbents are likely the most suitable entry point, alongside other potential funding opportunities.

that emphasized API driven scheduling, dynamic resource allocation, and multi-

tenant architectures. Leaf Space, ATLAS Space Operations, RBC Signals, CONTEC, and Infostellar, all launched between 2014 and 2016, introduced innovative service models. These firms generally fall into two categories: operators of proprietary ground station networks such as Leaf Space and ATLAS, and aggregators that broker partner capacity such as Infostellar and RBC Signals. As market demand for scalability grows, many of these companies are evolving into hybrid models that combine owned and partner ground stations.

By the end of the decade, hyperscalers entered the GSaaS market. AWS launched Ground Station in 2019 to support its Kuiper constellation, integrating satellite downlink with its global cloud infrastructure. Microsoft followed with Azure Orbital in 2020, leveraging its data centers and fiber backbone. Azure Orbital was retired in 2024, the same year GSaaS interest accelerated following EQT's acquisition of Eutelsat's ground station assets. Microsoft sold its antennas to Space Leasing International, which transferred them to RBC Signals for operational use.

In recent years, the rise of cellular non-terrestrial networks (NTNs) and Direct-to-Device (D2D) services has further increased interest in GSaaS. New constellations such as AST SpaceMobile and the Chinese megaconstellations, along with established players like Globalstar whose satellites support Apple's Emergency SOS via Satellite service, require access to ground facilities that exceed those used for legacy services, if any were provided. These developments signal a shift toward cloud native GSaaS platforms that integrate seamlessly with downstream data processing, storage, and analytics. They also create new opportunities for telecom service providers delivering connectivity to these facilities, especially in developing markets.

Together, these phases reflect the GSaaS market's transition from mission specific ground stations to flexible, distributed, and

GSaaS Market Drivers

Current estimates place GSaaS market revenue at approximately \$250 million, with rapid growth driven by the surge in LEO satellites supporting Earth observation, remote sensing, navigation, IoT connectivity, space science, and early-stage missions. With roughly 12,000 small satellites projected by 2030, demand for scalable and cost-efficient ground access is accelerating. Commercial operators increasingly rely on GSaaS providers to avoid the capital and operational burden of building dedicated infrastructure. In contrast, major communication constellations such as Starlink construct proprietary ground stations to maximize utilization and performance, while classified government and defense missions continue to rely on dedicated facilities to ensure security and sovereignty.




Figure 1 Active satellites in orbit. [Source: Ofcom]

software-defined infrastructure. This transformation continues to reshape how satellite operators, data users, and investors engage with space based assets.

Strategic Pivot by GEO Operators. In recent years, incumbent GEO satellite operators have pursued strategic partnerships with GSaaS providers to monetize their ground station networks. This includes unlocking capital through leasing arrangements and asset carve-outs. The GEO satellite industry, facing competitive pressure from LEO operators such as SpaceX and declining demand for broadcast and video services, is grappling with falling revenues and elevated

Table 1 Select GSaaS providers.

Provider	Year Founded	Capital Raised (M)	Revenue ¹ (M)	Number of Ground Stations ²
Swedish Space Corporation (SSC) (Sweden)	1972	N/A	N/A	10
KSAT (Kongsberg Satellite Services) (Norway)	2002	N/A	556 NOK (1Q2025) (~\$53)	20 (+4)
Leaf Space (Italy)	2014	\$48.5	N/A	35
ATLAS Space Operations ³ (USA)	2015	\$39.8	N/A	11 (+24)
CONTEC (South Korea)	2015	N/A	\$71.4 (TTM; Jun '25)	12
RBC Signals (USA)	2015	\$3.2	N/A	50+ (aggregated)
Infostellar (Japan)	2016	\$21.8	N/A	26 (own & partners)
AWS Ground Station (USA)	2019	N/A	N/A	12

1 GSaaS operates as a business unit within AWS, KSAT, and SSC, none of which typically disclose segment-level revenues. Financial data is also unavailable for privately held GSaaS providers.

2 Ground stations may be owned or shared from partners. Partner stations are listed in parentheses. Additional indicators of network scale include the number of antennas and daily satellite passes. Stations differ in their supported frequency bands.

3 Acquired by York Space Systems in July 2025 for an undisclosed amount.

debt. Consolidation, including Viasat's acquisition of Inmarsat and the merger between SES and Intelsat, reflects one strategic response. With extensive ground infrastructure, GEO operators have sought to extract value by partnering with GSaaS providers or pursuing full asset divestitures.

One of the most visible examples of this trend came in 2024, when Eutelsat sold its ground station assets to EQT in a transaction valued at €790 million, or approximately \$830 million. The portfolio included around 1,400 antennas across more than 100 locations worldwide. A new entity was created to hold the passive ground infrastructure assets. These include land, buildings, support systems, antennas, and fiber connectivity circuits across the combined portfolio of teleports and satellite network portals. Eutelsat, serving as the anchor tenant, retained a 20% ownership stake in the new entity. This transaction opened the door for similar deals involving other incumbent satellite operators that are under financial pressure due to rising debt and declining revenues.

Beyond asset sales, GEO operators and GSaaS providers are also pursuing operational synergies through hosted infrastructure and shared service models. These arrangements allow incumbents to expand their satellite service without divesting ownership of ground stations. For example, Viasat, an integrated satellite operator and developer, partnered with Infostellar in 2020 in a collaboration where each company acts as a market channel for the other's services in their respective regions. In 2023, Viasat opened a Real-Time Earth (RTE) ground station in Japan, hosted at Infostellar's site in Hokkaido. Most recently, in September 2025, it announced a joint venture with Space42 called Equatys to operate a D2D service over a constellation that leverages shared infrastructure.

Ground Station Innovation Trends. Ground stations today differ significantly from the first generation that emerged in the late 1990s. Early deployments were dedicated and mission-specific, relying on proprietary hardware, custom protocols, and semi-automated operations. These early GSaaS models resembled colocation data centers,

offering power, space, and connectivity to satellite operators.

Modern GSaaS platforms differentiate through both technology and business model. Technologically, they enable greater flexibility, dynamic scheduling, and highly automated operations, which improve efficiency and reduce costs. From a business model perspective, GSaaS aggregators act as capacity brokers, selling antenna time on third-party teleports. Both aggregators and owner-operators offer opex-based models to customers, who are primarily private enterprises in the NewSpace era. Governments remain an important segment for non-classified missions, in contrast to the government-led demand of earlier decades.

Technological innovation in GSaaS spans both hardware and software domains, driving operational efficiency and enhancing profitability. Key advancements include:

- **Software-defined radios (SDR):** Modern GSaaS platforms leverage SDRs to virtualize the communication stack, especially the baseband processing. This enables dynamic reconfiguration of waveforms and modulation schemes on a per-mission basis. The flexibility allows a single ground station to support diverse customers with tailored protocols, creating a differentiated multi-tenant service layer beyond basic connectivity.
- **Cloud integration and APIs:** Seamless integration with major cloud providers such as AWS, Azure, and GCP supports automated data ingestion, storage, and processing pipelines. APIs allow satellite operators to programmatically schedule passes, retrieve telemetry, and trigger analytics.
- **Antenna technologies:** The industry is actively adopting electronically steered phased array antennas. These enable simultaneous multi-beam tracking, faster satellite handovers, and reduced mechanical maintenance, which is critical

Strategic Differentiation in GSaaS Platforms

While GSaaS providers differentiate along the innovation trends described earlier, there are additional critical parameters to consider. Three key aspects include:

1. **Frequency band support:** The range of supported satellite frequency bands for telemetry, tracking and command (TT&C) and data payloads includes VHF, UHF, S, X, C, Ku and Ka bands, among others. These bands correspond to specific antenna systems that influence cost and affect utilization ratios.
2. **Geographic presence:** The global distribution of a provider's ground assets is essential for maximizing data throughput, minimizing latency, ensuring redundancy, enabling global service reach, and supporting complex constellation operations.
3. **Daily pass capacity:** The number of satellite passes per day is determined by a provider's footprint and asset density. This metric reflects a network's ability to support large constellations and deliver high-value, time-sensitive data. Scheduling and antenna technology play a critical role in contact capacity.




Figure 2 High Level Architecture of a GSaaS.
[Source: MilsatMagazine]

for high-cadence LEO constellations. Though not yet dominant, phased arrays represent the next evolution in ground segment scalability.

- **Edge computing:** Processing data directly at or near the ground station, such as

decrypting, decompressing, or running AI models for anomaly detection or feature extraction, reduces upstream bandwidth costs and delivers low-latency insights for time-sensitive applications. Edge computing optimizes total cost of ownership and enhances service value.

- Scheduling and analytics leveraging AI/ML: Orchestrating passes across

multiple satellites, customers, and global antenna sites under dynamic constraints (weather, priority SLAs, orbital mechanics) is a complex optimization challenge.

AI/ML models enable predictive scheduling, interference mitigation, and real-time resource allocation. These capabilities improve antenna utilization, ensure reliability, and directly enhance operations and service quality.

Key Takeaways

- The GSaaS business model is gaining momentum from two opposing trends: rapid growth in LEO satellites and financial pressure on geostationary operators. These forces are driving demand for GSaaS offerings that combine capital intensity with recurring revenue, which appeals to private equity and infrastructure investors.
- GSaaS mirrors the data center model by decoupling infrastructure from service delivery, enabling multi-tenant access, dynamic scaling, and pay-as-you-go pricing integrated with cloud-based compute, storage, and analytics. For investors in data centers, tower companies and shared infrastructure, GSaaS presents a familiar model with new growth vectors tied to satellite service demand.
- The GSaaS model today is fundamentally different from its predecessor of the late 1990s, shaped by advances in technology and business strategy. Key innovations include software defined radios, cloud integration, phased array antennas, edge computing, AI driven scheduling, and a pay as you go model.
- Estimated GSaaS market annual revenue is around \$250 million, driven by growth in LEO satellite deployments. Commercial operators prefer GSaaS to avoid the cost and complexity of building proprietary infrastructure. Key segments include earth observation, remote sensing, navigation, technical and scientific missions, and selectively some communication satellite constellations. Government and defense missions maintain proprietary infrastructure for security and sovereignty.
- GSaaS providers fall into two categories: operators of proprietary networks and aggregators of partner capacity. Hybrid models combining elements of both categories are emerging quickly as demand for scalability grows across geographic coverage and frequency band support.
- GSaaS performance depends on frequency band support, global footprint, and network contact capacity, along with the infrastructure differentiators outlined above. Providers are increasingly embedding orchestration, analytics, and cloud native interfaces to support constellation management and time sensitive data delivery.
- High utilization is enabled by dynamic scheduling and multi mission support across diverse satellite applications. This operational efficiency allows GSaaS providers to maximize asset use across multiple customers and missions, improving revenue per ground station. By reducing idle time and enabling flexible access windows, providers can achieve stronger margins and faster payback periods on infrastructure investments.
- Investor entry points favor carve outs and infrastructure roll ups. Given the scale and capital profile of private equity and infrastructure funds, strategic carve outs from incumbents or aggregation of regional assets offer viable paths.

The Disruptive Force of Starlink: Reshaping Rural Connectivity, Challenging Incumbents

Overview. As Starlink approaches 5 million subscribers—a milestone achieved in just over four years since its beta launch in November 2020—its disruptive impact is evident across multiple segments of the telecom ecosystem. This Insight Note examines Starlink’s influence on geostationary orbit (GEO) satellite operators, rural wireless Internet service providers (WISPs), and subsea fiber cable projects.

Destruction in the GEO broadband market. A significant portion of Starlink’s growth appears to have come at the expense of competing GEO satellite services. For instance, both Viasat and Hughes have experienced notable declines in their subscriber bases during the same period. Specifically, Viasat’s subscriber count dropped from 596,000 at the end of December 2020—just as Starlink launched—to 205,000 by the end of December 2024. This represents a loss of nearly two-thirds of its subscribers in just over four years.

Hughes has also seen a significant decline, though the impact appears less stark because it reports global subscriber numbers. At the end of December 2020, Hughes had 1,564,000 subscribers. By September 2024, this number had dropped to 912,000 broadband subscribers—a decline of 42%, or roughly 2 out of every 5 subscribers.

These numbers highlight the significant threat posed by LEO satellite services to traditional GEO service providers. While this impact is currently most evident in the satellite broadband/fixed access market, broadcast services are also experiencing a decline as alternative options become more

widely available. These trends underscore the urgency for GEO providers to develop robust and effective strategies to address the LEO challenge. Many leading GEO players have already adopted such strategies, including LEO-GEO partnerships, acquisitions, and even building their own LEO constellations.




Figure 1 Satellite broadband subscribers.
[Source: Xona Partners, company data]

Pressuring the WISPs. The example of the Canadian market applies to other markets as well. In mid-2024, Starlink announced that it had reached 400,000 Canadian subscribers. Considering there are approximately 1.8 million dwellings in Canada classified as “moderately remote”, “more remote”, or “most remote”, this translates to a penetration rate of just over 22%, assuming all of Starlink’s customers are in these areas. In other words, more than 1 in 5 rural or remote Canadian households now have a Starlink connection.

Starlink’s growth is occurring, in part, at the expense of rural WISPs. Many rural wireless networks are patchy and stretched thin, offering low data rates at high prices. While Starlink pricing is not cheap, some will buy it

for its higher offered data rate and potential improvement in quality of service compared to typical WISPs.




Figure 2 Classification of Canadian municipalities according to the index of remoteness, 2021. [Source: Statistics Canada]

However, Starlink is not the only challenge facing WISPs. They are also facing significant competition from government-subsidized fiber development programs. These programs have accelerated in response to the COVID-19 pandemic, and there is still substantial funding available for further development. WISPs are squeezed between competition from LEO satellites and fiber deployments. For this reason, some WISPs have pivoted to seek grants for fiber deployments, which is a wise strategy. Other WISPs have opted to exit the market and sell to larger entities that can build economies of scale. For the remaining WISPs, it will be imperative to carefully assess their market and devise and execute a clear strategy to mitigate such destructive market dynamics.

Starlink's penetration impacts the sales of fixed wireless equipment vendors, particularly those offering unlicensed-band solutions. Assuming an average of 20 rural subscribers per typical unlicensed-band Fixed Wireless Access (FWA) access point, Starlink's growth could have displaced up to 250,000 FWA units from the market. In reality, the number should be lower because many Starlink subscribers were previously served by other technologies, including GEO satellites, or had no service at all.

Dealing with the threat of Starlink.

Starlink disrupted the traditional GEO satellite market, prompting incumbent operators to adopt various strategies to counter the rise of LEO satellites. These strategies include:

- **Acquiring LEO constellations** – The most notable example is Eutelsat's acquisition of OneWeb, positioning itself as a multi-orbit operator.
- **Building proprietary LEO constellations** – Telesat opted to develop its Lightspeed LEO network with 188 satellites, targeting enterprise broadband, particularly in maritime and aviation. Hughes also announced plans to develop its own LEO constellation.
- **Consolidating GEO operators** – Viasat's 2023 acquisition of Inmarsat exemplifies this approach. Notably, Inmarsat had previously planned its own LEO constellation, Orchestra.
- **Developing multi-orbit operations** – SES acquired the O3b MEO constellation in 2016, before Starlink's launch. However, the rapid rise of LEO reinforced SES's commitment to a multi-orbit strategy.
- **Forming strategic partnerships** – Intelsat partnered with OneWeb early on, a collaboration that continues even as OneWeb is now part of Eutelsat, reflecting the shared challenges GEO operators face in the LEO era.

These responses highlight the shifting dynamics in satellite communications as legacy operators adapt to remain competitive in an increasingly LEO-dominated market.

Nevertheless, Starlink's impact on the FWA equipment market exceeds a few tens of thousands of access points. Companies such as Ubiquiti, Cambium, Tarana, Mikrotik, Mimosa, and others are particularly at risk.

Encroaching on Subsea Fiber Cables. There is a potential overlap between satellite services and subsea fiber cables in certain market segments, particularly in northern coastal communities such as those in Alaska, Canada, and parts of Northern Europe. For remote communities, Starlink introduced its Community Gateway service, which was first deployed in Alaska. This service can scale up to 10 Gbps/month, with pricing set at 75,000/Gbps/month and an upfront capital expenditure charge of \$1.25 million.

The challenge with serving these northern communities lies in their geographic spread and low population density, which make the business case for subsea fiber cables difficult without significant government funding. For example, the planned FISH South subsea cable in Alaska aims to connect several communities, including Cordova (population ~2,400), Yakutat (~500), Gustavus (~600), Hoonah (~1,000), Juneau (~32,000), and Pelican (~100).

Despite challenging economics, Starlink is unlikely to fully replace subsea fiber in such scenarios for several reasons. Starlink's opex-based model contrasts with the capex-based model of subsea fiber cables, which aligns better with equity investing and financing from both government and private sources. As a result, while Starlink may serve as a viable solution for inland remote communities, its use in shoreline remote communities is likely to be a temporary gap-fill solution.

Another example of overlap is in connecting oil and gas (O&G) platforms. In many cases, Starlink is a viable solution. However, large O&G platforms with long operational lifespans and proximity to dry-plant infrastructure remain more likely to be served by subsea fiber due to its reliability and high capacity.

Persisting profitability challenge. The disruptive nature of Starlink is poised to extend into other sectors, such as maritime

Achieving economies of scale for user terminals.

The cost of the user terminal is critical to the profitability of wireless access technologies, including LEO satellite systems.

Starlink user terminals are complex products, featuring adaptive beam-steering antennas designed to track and hand over between satellites. They also operate with a relatively high RF power output of approximately 4 watts, utilizing the Ku and Ka microwave bands (13 GHz and 27 GHz), which are not commonly used in other telecommunication systems. As a result, the only economies of scale for these terminals come from Starlink's own subscriber base.




Figure 3 Cost of Starlink User Terminal.
[Source: ARK Invest]

To achieve these economies of scale, it is imperative to build a large subscriber base. Without sufficient scale, SpaceX would need to subsidize the cost of user terminals, which could strain profitability. This explains SpaceX's push for inclusion in government subsidy programs, such as rural broadband initiatives, as a key strategy to expand its subscriber base and achieve the necessary scale for cost efficiency.

and aviation, as the company aims to broaden its service offerings. Diversifying geographically and across different service types is essential to drive higher network

efficiency. LEO satellite networks are inherently low-utilization networks due to the concentration of traffic in specific regions. By expanding its geographic reach and service portfolio, Starlink can increase

network utilization, which directly impacts profitability. However, the question remains: Will Starlink overcome the ever-moving goalpost of profitability?

Key Takeaways

- Starlink's meteoric rise has exposed a competitive weakness in GEO satellite operators, leading to a significant loss of subscribers since Starlink's launch in late 2020.
- GEO satellite providers have adopted several strategies to counter Starlink's threat, including acquiring LEO constellations, developing their own LEO systems, and forming strategic partnerships.
- While Starlink's growth to 5 million subscribers has improved the cost efficiency of its user terminals through economies of scale, the question of profitability remains unresolved.
- Rural WISPs are under pressure from both Starlink and government-subsidized fiber projects. To compete, WISPs are developing strategies such as participating in fiber network development or pursuing mergers and acquisitions.
- Starlink's opex-based model provides a temporary solution for remote coastal communities, but it does not eliminate the need for subsea fiber cable projects, whose capex-based model aligns better with government and private investment frameworks.

About Xona Partners

Xona Partners (Xona) is a boutique advisory services firm specializing in technology, media, and telecommunications (TMT). Established in 2012 by a team of seasoned technologists, startup founders, managing directors in global ventures, and investment advisors, Xona leverages its founders' cross-functional expertise to offer a unique, multidisciplinary approach to technology and investment advisory services. Our clientele includes private equity and venture funds, technology corporations, regulators, and public sector organizations. We assist our clients with pre-investment due diligence, post-investment lifecycle management, and strategic technology management, helping them identify new revenue streams and navigate the complex landscape of the TMT sector.

E-mail: advisors@xonapartners.com | Web: xonapartners.com

Friend or Foe: Rethinking Mobile Strategy in the Age of Non-Terrestrial Networks

Overview. Low Earth Orbit (LEO) satellite constellations providing direct to device (D2D) services rely on terrestrial mobile spectrum to operate. This is only possible through an agreement with a mobile network operator, as regulated by national frameworks governing such services. This was the case until recently, when AST SpaceMobile and SpaceX acquired Mobile Satellite Spectrum (MSS) with the intention to operate independently of mobile network operators using an over the top model. They join Globalstar, which already operates D2D services in partnership with Apple.

This development carries important implications for competitive dynamics in the wireless industry. It allows the Satellite Network Operator (SNO) to operate without relying on a partnership with a mobile network operator or being tied to the mobile operator for revenue generation. We explore these shifts in this Insight Note, along with the potential consequences and new opportunities emerging in the market.

NTN Operational Challenges. The operation of Non Terrestrial Networks (NTN) offering cellular services through LEO or other satellites is governed by national regulatory frameworks for accessing mobile spectrum. This includes the Supplemental Coverage from Space (SCS) framework in the United States and the Supplemental Mobile Coverage by Satellite (SMCS) framework in Canada. These frameworks define the frequency bands where satellite services are permitted to operate, the geographic area requirements, and the technical specifications needed to protect terrestrial systems from interference.

EchoStar's Spectrum Unwind

In the spring of 2025, SpaceX applied pressure on Echostar by submitting letters to the Federal Communications Commission (FCC) alleging that Echostar was not meeting its spectrum license obligations due to underutilization of its assets. This led the FCC to open an inquiry into Echostar's spectrum usage. Facing regulatory scrutiny, Echostar agreed to sell its MSS spectrum to SpaceX for \$19 billion. The deal included a combination of cash, stock, and debt payments.

Echostar had originally acquired this spectrum for \$4.42 billion through two transactions. The first involved purchasing 40 MHz of MSS S-band spectrum from two bankrupt companies in 2011 for \$2.86 billion. The second was the acquisition of 10 MHz of H-block spectrum for \$1.564 billion.

In a separate transaction, Echostar sold its valuable 600 MHz and 3.45 GHz assets to AT&T for \$23 billion, marking its exit from the mobile services market. These assets had been acquired for approximately \$13.54 billion across two auctions. As of the time of writing, Echostar has raised \$42 billion from assets it originally purchased for just under \$18 billion, representing a gain of approximately 134%. For comparison, T-Mobile acquired Sprint for \$26 billion in a deal announced in 2018 that secured a nationwide average of 150 MHz of 2.5 GHz spectrum, along with PCS and 800 MHz holdings.

Echostar continues to hold AWS-3 spectrum with a book value of around \$10 billion. It also retains other spectrum assets across various bands, including 700 MHz, CBRS, millimeter wave, and satellite spectrum, with a combined book value of approximately \$2 billion.

The challenge with this operating model include:

1. **Global Regulatory Adoptions.** The D2D service concept is relatively new. Only a few countries, including the United States with FCC's SCS, China, Canada with ISED's SMCS, and Australia, have defined regulatory frameworks for D2D services. Australia's ACMA issued guidance in 2024 to enable D2D under existing spectrum licenses, supporting partnerships like Optus and Telstra with Starlink. Ofcom in the United Kingdom is developing a framework, aiming to be the first in Europe to enable satellite-to-smartphone connectivity. Most other jurisdictions lag in defining specific D2D frameworks, limiting global proliferation due to regulatory and technical challenges.
2. **Stringent Technical Co-Existence Requirements.** Coexistence of satellite and mobile services in the same frequency spectrum requires close coordination to avoid both co channel and adjacent channel interference. This has led to requirements such as the Geographically Independent Area (GIA), which defines the region where a SNO can provide service, and the power flux density threshold needed to prevent interference. These requirements have a significant impact on satellite design considerations, including antenna configuration and radio frequency power, as well as operational behavior such as service coverage and performance. These limitations affect both the potential service revenue and the cost of deploying a satellite constellation, with far reaching consequences.
3. **Business Model Constraints.** SNOs need to partner with MNOs to provide D2D services due to spectrum access requirements. The most common model involves revenue sharing between the satellite and mobile operators. Since MNOs hold the spectrum licenses, they may be seen as having leverage over the SNOs, who must offer real or perceived exceptional value to gain bargaining

The Business Case for NTN

Market research firms are optimistic about NTNs. Novaspacespace forecasts 425 million D2D users by 2033 and \$42 billion in service revenue over the next decade. ABI projects \$17 billion by 2032.

Still, the business case for D2D services is risky, and profitability is uncertain. A global constellation costs \$5–6 billion to deploy, and 15–20% of satellites must be replaced annually, resulting in higher opex than traditional geosynchronous systems.

Global coverage is needed to scale revenue. Terrestrial networks, while not covering all landmass, reach most of the population, especially in developed markets with higher willingness to pay. In these regions, regulations often require over 95% population coverage. For example, OpenSignal data shows U.S. mobile users are disconnected only 1% of the time. In emerging markets, lack of connectivity is often due to affordability or other non-coverage factors.




Figure 1 Proportion of time spent without cellular signal. [Source: OpenSignal]

For context, Globalstar generated \$237 million in revenue in 2024, including fees from Apple, which we expect to be a significant share. For these reasons, we believe the success of the D2D market will depend on additional applications beyond personal connectivity.

power. Many mobile operators remain uncertain about the value of partnering with a D2D satellite operator, despite several incentives. These include

improving network reliability and resiliency during outages, gaining a competitive edge over other mobile operators, and earning favor with regulators who are focused on enhancing the reliability of emergency services. These factors are especially important in a saturated mobile market that increasingly resembles a zero sum game in terms of customer acquisition and revenue growth.

Advantages of MSS Spectrum. The regulatory framework for providing mobile satellite services in MSS frequency bands is well established across many jurisdictions. One advantage in the context of D2D services is that it allows the SNO to operate without entering into a special business partnership with a MNO, since spectrum subordination is not required. In this arrangement, the satellite operator can act as a roaming partner, which simplifies service testing, integration, and the business model, which is already well understood.

Globalstar, in partnership with Apple, operates under this model and was the first to launch services in November 2022 with the introduction of the iPhone 14. In recent months, both AST SpaceMobile and SpaceX have acquired MSS spectrum.

MNOs in the Crosshair? There is a common misconception that by acquiring MSS

Table 1 MSS Spectrum holding of select D2D SNOs.

	User Downlink (MHz)	User Uplink (MHz)
Globalstar	2483.5–2500	1610–1618.725
SpaceX	2180–2200	2000–2020
AST L-Band US+Canada (45 MHz)	1525–1559	1626.5–1660.5 and 1670–1675
S-Band Global Priority Rights	2170–2200	1980–2010

A Perspective on MSS Spectrum Value

The amount of money SpaceX paid EchoStar for a combined 50 MHz of S-band MSS spectrum and nationwide PCS H-block spectrum—\$19 billion—is typically associated with terrestrial spectrum valuations. In contrast, EchoStar originally acquired these assets for just \$4.42 billion about 14 years earlier: \$2.86 billion for assets from bankrupt satellite companies, including 40 MHz of S-band MSS spectrum, and \$1.564 billion for 10 MHz of nationwide PCS H-block spectrum.

This contrasts with a recent transaction by AST SpaceMobile, which acquired 45 MHz of L-band spectrum from Ligado and Inmarsat for \$550 million.

Table 2 MSS Spectrum Value.

	MHz	\$, million	\$/MHz-PoP
SpaceX: S-Band + PCS H-Block	50	19,000	1.12
AST: S-Band	60	64.5	Global priority rights
AST: L-Band	45	550	0.0321

Regardless of the relative quality of the spectrum assets, the price disparity remains significant. This raises important questions about SpaceX's future ambitions and the strategic value it assigns to these holdings. Notably, SpaceX's MSS spectrum permits terrestrial operations (AWS-4; 3GPP Band 23), opening the door to a broad range of potential services beyond traditional satellite connectivity. This echoes Globalstar's efforts to enable terrestrial use of its S-band MSS spectrum, reinforcing the trend toward hybrid deployments that blur the lines between satellite and mobile networks.

spectrum, SpaceX would offer direct to device services that compete directly with MNOs. This belief overlooks fundamental laws of physics that govern wireless communications. The distance between a

satellite and a mobile device, which ranges from 300 to 1500 kilometers for LEO satellite constellations, is far too great to deliver a strong and reliable signal compared to terrestrial towers, which typically serve cells with a radius of less than 20 kilometers. Satellite communication also relies on a single antenna system. The line-of-sight nature of satellite links makes multi antenna techniques, which drive large capacity gains in terrestrial networks, ineffective. In addition, satellite services are limited to open and unobstructed environments and generally do not function indoors.

Given the high price SpaceX paid for S-band MSS spectrum, it is likely that the company has specific applications in mind that justify the valuation. These applications are not only focused on personal D2D services. Instead, SpaceX may intend to use these frequencies for automotive use cases such as telemetry for autonomous vehicles or for advanced future applications in robotics and

automation. High-value Internet of Things (IoT) connectivity and emergency services round out the potential use cases. The key point is that MSS spectrum would be deployed in high value use cases that offer strategic advantage in emerging markets rather than in consumer applications that already exist in some form. For example, Iridium and Globalstar already offer personal communication services via LEO satellites, while Orbcomm provides IoT connectivity solutions.

We emphasize this to encourage MNOs, who continue to face challenges in growing revenue, to think beyond traditional models and consider how emerging NTN services could reshape future business opportunities. While there are areas of overlap, mobile and satellite networks are largely complementary across cost and performance parameters. This creates fertile ground for collaboration in developing new services and revenue streams.

Key Takeaways

- **Shift Toward Independent Satellite Operations.** The acquisition of MSS spectrum enables satellite operators like SpaceX and AST SpaceMobile to bypass traditional MNO partnerships, operating via an over-the-top model under existing regulatory frameworks.
- **Regulatory Frameworks Are Critical but Limited.** Only the U.S. and Canada have established formal frameworks (SCS and SMCS) for D2D satellite services; others like the UK and Australia are still developing theirs. This limits the global proliferation of D2D services.
- **Technical Coexistence and Business Challenges for D2D Services.** Operating satellite and mobile services in the same spectrum requires strict coordination to avoid interference. On the business side, SNOs partner with MNOs for spectrum access, through revenue-sharing models where MNOs hold leverage. Many remain skeptical about the value of D2D partnerships despite benefits like improved network reliability, competitive differentiation, and regulatory favor for emergency services.
- **High Valuation Reflects Strategic Intent.** SpaceX paid approximately \$1.12 per MHz-pop for its MSS spectrum, far exceeding typical satellite valuations. This suggests ambitions beyond consumer connectivity, potentially targeting automotive, robotics, or IoT applications.
- **MNOs Not Directly Threatened.** Despite concerns, MSS-based D2D services are unlikely to compete directly with MNOs due to physical limitations of satellite communication, including low throughput from long distances, line-of-sight constraints, and the inability to leverage multi-antenna techniques. Satellite connectivity is better suited for open environments and specialized use cases such as automotive telemetry or robotics, rather than indoor consumer applications.
- **MNOs Should View NTN as Opportunity, Not Threat.** Rather than seeing NTNs as disruptive, mobile operators should explore synergies in resilience, emergency services, and new verticals (e.g., autonomous vehicles, automation), especially in saturated markets where growth is stagnant.

Navigating the Future of HAPS: Challenges, Investments, and Innovations

Overview. High Altitude Platform Systems (HAPS) have gained increasing attention as a tactical tool in the defense sector; however, their adoption as a communication solution has lagged, with several attempts failing due to high costs. With a new generation of active startups and projects in this area, we outline some of the critical hurdles that HAPS—as a communication platform, which is the focus of this paper—must overcome, as well as some of the opportunities for their deployments. In our view, HAPS complement low Earth orbit (LEO) satellite systems; therefore, it is not beyond the realm of possibility that the future will see greater investment interest from the satellite sector.

Evolution of HAPS. HAPS refer to platforms operating in the stratosphere at altitudes between 20–50 km above Earth's surface, with most positioned between 20–23 km to avoid the stronger stratospheric winds found at higher elevations. This is relatively low compared to LEO satellites (350–2,000 km), but significantly higher than commercial airplanes, which typically cruise at around 10 km altitude.

The concept of HAPS dates back to the 1990s in the context of mobile communications, when manned fixed-wing aircraft with an 8-hour flight time were proposed. However, it quickly became evident that such a proposition would not be seriously considered.

Recent advancements in materials science, manufacturing, solar cells, power storage, AI, and automation have enabled the development of a new generation of HAPS, including fixed-wing aircraft (e.g., Facebook Aquila), airships, and balloons (e.g., Google

Investments in HAPS

Investments in HAPS have experienced ebbs and flows over the past decade. Cloud service providers were among the earliest and most visible investors, with Google's Loon and Meta's Aquila serving as prominent examples during the 2010s. Their involvement helped stimulate a broader wave of HAPS-related activity, which has since included aerospace offshoots, defense initiatives, and national space agencies.

One example from the telecom sector is SoftBank, which launched HAPSMobile in 2017 to explore high-altitude connectivity. Partnering with AeroVironment, the company invested around \$65 million to develop the SunGlider platform. It also acquired approximately 200 patents from Google Loon and played an active role in advocating for global HAPS spectrum at WRC-23. In 2023, SoftBank reintegrated HAPSMobile into its core operations.

NTT Docomo, along with Space Compass, Mizuho Bank, and the Development Bank of Japan, collectively committed \$100 million to Aalto, an Airbus subsidiary spun out in 2022 to commercialize the Zephyr platform across Asia. A similar telecom-driven investment includes Deutsche Telekom's ongoing support for the UK's Stratospheric Platforms, which dates back to 2016.

While defense funding remains strong, commercial HAPS success depends on financial viability. Multi-mission models—blending telecom, sensing, and surveillance—offer promise but pose added risk and complexity in practice.

Table 1 Comparative Analysis of Fixed-Wing, Airship, and Balloon HAPS Platforms.

Criteria	Fixed-Wing HAPS	Airship HAPS	Balloon HAPS
Endurance	Weeks to months (solar).	Weeks to months (up to 1 year; buoyancy-driven)	Weeks to months (up to 6 months; buoyancy-driven)
Maneuverability	High (dynamic lift, precise navigation)	Moderate (limited propulsion)	Low (wind-dependent)
Flight Stability	Moderate (susceptible to gusts, robust control systems)	High (station-keeping capability, wind-resistant envelopes)	Moderate (stable but wind-sensitive)
Payload Capacity	Low to moderate (15–50 kg, limited by wing loading)	High (50–250 kg, large volume)	Moderate to high (up to 100 kg, balloon size-dependent)
Deployability	Moderate (requires runway/catapult, ground infrastructure)	Low (complex launch, large ground facilities)	High (simple launch, minimal infrastructure)
Operational Team Requirements	High (pilots, engineers, maintenance crew)	Moderate (ground crew, fewer flight control needs)	Low (small team, automated systems)

Loon). Although these ambitious projects were ultimately retired—Aquila in 2018 and Loon in 2021—they helped shape the development of modern HAPS, paving the way for further innovation.

Type of HAPS. HAPS fall into two main categories with three platform types: (a) heavier-than-air (HTA), such as fixed-wing aircraft; and (b) lighter-than-air (LTA), including airships and balloons. HTA systems rely on engines, onboard power generation, and storage to produce lift—factors that directly impact endurance. In contrast, LTA platforms use buoyant gases like helium to stay aloft. These fundamental operational differences lead to critical trade-offs affecting performance and cost parameters:

Endurance – The ability to remain aloft for extended periods.

Maneuverability – The ability to follow precise paths, maintain stability, and adjust positioning or course efficiently.

Deployability – Refers to the ease of launch, recovery, and integration into existing networks.

A Lesson from Google Loon

One of the main drivers of Google Loon’s high costs was the impact of stratospheric winds and the difficulty of keeping balloons stationary over a specific location. While Loon’s engineers developed advanced wind prediction models and implemented techniques to maintain positioning, ultimately, they had to overprovision balloons to achieve coverage, which significantly increased operational costs.

Flight Stability – The ability to maintain a steady position despite stratospheric winds and turbulence (i.e. station-keeping).

Payload Capacity – The payload determines the applications a HAPS can support, such as communications, surveillance, or scientific missions. While heavier payloads expand utility and diversify use cases or tenants, they may also reduce endurance and maneuverability.

Operational Team – Teams are required for pre-launch preparation, flight operations, and post-flight maintenance and recovery.

Considering the various trade-offs, we can summarize the differentiation among type of HAPS:

- Fixed-Wing: Best for maneuverability, but limited by payload capacity.
- Airships: Excel in endurance and payload capacity, but complex to deploy.
- Balloons: Cost-effective and rapidly deployable, but lack maneuverability.

Examples of HAPS include:

- Fixed-wing HAPS – Aalto Zephyr (Airbus spin-out), SoftBank-AeroVironment SunGlider and Prismatics’ Phasa-35, Swift Engineering, and Mira Aerospace.
- Airship HAPS – Sceye, Thales Stratobus, and Stratosyst’s Skyrider.
- Balloon HAPS – Stratollite by World View Enterprises, Stratostats by Involve, and Aerostar.

Use Cases in Telecom Services. There are several notable use cases for HAPS in telecoms. Some widely cited examples include:

1. **Greenfield Deployments** – Extending connectivity to unserved areas, including remote regions in developed markets and emerging markets where terrestrial coverage is cost-prohibitive.
2. **Network Resiliency & Emergency Services** – HAPS provide rapid-deployment connectivity, supporting disaster recovery efforts and complementing temporary terrestrial solutions, such as cell-on-wheels with satellite backhaul.
3. **Private Networks & IoT** – Used in specialized deployments, though cost sensitivity is a key factor in determining feasibility.
4. **Backhaul Services** – Competing with satellite-based solutions, some HAPS

offer advantages by maintaining a relatively static position over a given location. Although HAPS theoretically provide lower latency due to their proximity to Earth, the end-to-end delay may not always be significantly lower than LEO satellites—especially if the transport network relies on satellite links, negating any latency benefits.

5. **Extended Coverage Over the Sea** – Functions similarly to greenfield deployments, delivering connectivity to maritime operations and offshore sites.
6. **Satellite Capacity Augmentation** – A less frequently considered use case, primarily due to the relative novelty of modern LEO satellite networks. However, this approach has the potential to drive financial justification for HAPS




Figure 1 From top: Aalto Zephyr fixed-wing aircraft; Sceye airship; and Raven Aerostar balloon.

deployment more effectively than other use cases, as explored further below.

Complementing Satellites. LEO satellites provide blanket coverage, distributing traffic evenly across the Earth—irrespective of where demand is concentrated. While this ensures global connectivity, it also means that increasing capacity in high-demand areas requires scaling the entire constellation, an expensive and often inefficient approach.

HAPS, on the other hand, deliver targeted, localized capacity without the need for a massive fleet, keeping costs under control. For LEO satellite networks, HAPS serve as a high-impact complement, efficiently reinforcing capacity where it’s needed most.

Similarly, GEO satellite operators can strategically integrate HAPS to boost coverage in key locations, enhancing their multi-orbital strategies for greater flexibility and performance.

Market Acceptance. Several key challenges must be addressed before HAPS achieve wider market acceptance.

One major regulatory hurdle is spectrum allocation, as described in the Telecommunications Regulations side box. Beyond this, a fundamental question arises: Should HAPS require dedicated spectrum? Sharing spectrum with satellites or terrestrial networks can lead to significant interference at the boundaries between coverage layers. This issue has historically hindered scalable deployments of small cells in terrestrial networks.

A second challenge is spectral efficiency—just as with satellites, HAPS must improve their ability to maximize bandwidth usage. Unlike terrestrial systems, which leverage multiple-antenna techniques to increase capacity, the typical line-of-sight channel between HAPS and user devices does not support such methods. Advancing technical

Telecommunications Regulations

The ITU has been studying HAPS since the 1990s and was the first to allocate spectrum for HAPS at the World Radiocommunication Conference (WRC) in 1997. Since then, the list of spectrum bands for data backhaul applications (fixed use case) has expanded significantly. More recently, at WRC-23, mobile frequency spectrum was allocated to High Altitude IMT Base Stations (HIBS), further advancing HAPS-based communication capabilities.

Table 2 ITU Allocated Frequency Bands for HAPS and HIBS.

Service	Frequency Band
Mobile	698–960 MHz (Region 2)
	694–960 MHz (Region 1)
	1710–1980 MHz, 2010–2025 MHz and 2110–2170 MHz (Regions 1, 3)
	1710–1980 MHz and 2110–2160 MHz (Region 2)
	2500–2690 MHz (Regions 1, 2)
	2500–2655 MHz (Region 3)
Fixed	6440–6520 MHz (Specific countries; HAPS-to-ground)
	6560–6640 MHz (Specific countries; ground-to-HAPS)
	21.4–22 GHz (Region 2)
	24.25–27.5 GHz (Region 2)
	27.5–28.35 GHz (Region 2)
	31.0–31.3 GHz (Region 2)
	47.2–47.5 GHz (Global)
	47.9–48.2 GHz (Global)

National regulators have yet to establish frameworks for HAPS spectrum licensing and technical requirements necessary for commercial operations. Telecommunication regulations, along with the critical aviation regulations—which are beyond the scope of this discussion—present significant barriers that must be addressed to enable the future commercial growth of HAPS.

innovations to address this limitation will be crucial in enhancing HAPS’ ability to deliver high-speed broadband effectively.

Key Takeaways

- **HAPS provide localized, high-altitude connectivity** that bridges the gap between terrestrial networks and satellites—ideal for boosting capacity in high-demand zones without the costly scale of full LEO satellite constellations.
- **Fixed-wing aircraft, airships, and balloons** each involve trade-offs between payload characteristics (size, weight, power) and operational factors (endurance, maneuverability, deployability), as well as cost. Optimizing these factors to meet specific telecom demands remains a central challenge for future innovation.
- **Advancements in solar power, materials, and energy storage** have renewed momentum in HAPS, despite the shutdown of earlier ventures like Facebook Aquila and Google Loon, which exposed the cost-ambition mismatch.
- **Telecom applications range from remote coverage and emergency response to IoT**, though spectrum efficiency, affordability, and integration into broader networks remain critical hurdles.
- **Global spectrum allocations from the ITU exists**, but national-level regulatory clarity—including spectrum licensing, airspace permissions, and technical standards—is still evolving and poses a barrier to scale.
- **HAPS complement LEO satellite networks** by offering targeted coverage in bandwidth-constrained areas, opening the door to hybrid deployments and future investment opportunities from satellite providers.
- **Industry proposals envision HAPS supporting telecom, sensing, and surveillance applications**, but translating this into practice poses challenges for mobile operators aiming to ensure high service availability and reliability.

About Xona Partners

Xona Partners (Xona) is a boutique advisory services firm specializing in technology, media, and telecommunications (TMT). Established in 2012 by a team of seasoned technologists, startup founders, managing directors in global ventures, and investment advisors, Xona leverages its founders' cross-functional expertise to offer a unique, multidisciplinary approach to technology and investment advisory services. Our clientele includes private equity and venture funds, technology corporations, regulators, and public sector organizations. We assist our clients with pre-investment due diligence, post-investment lifecycle management, and strategic technology management, helping them identify new revenue streams and navigate the complex landscape of the TMT sector.

E-mail: advisors@xonapartners.com | Web: xonapartners.com

Optical Space Communications: Enabling the Next Leap in Connetivity

Overview. Using optical links for communications between space and Earth sits at the frontier of technology and presents significant challenges. While optical inter-satellite links (ISLs) have surpassed RF ISLs in deployments and applicability, RF remains the primary mode for feeder links connecting satellites to ground stations. Still, interest in optical feeder link (OFL) technologies persists and mirrors that in optical ISLs, driven by the security and high throughput advantages of optical systems. Some of this interest is directed toward research and development efforts or a few pioneering innovators who face the difficult task of making their technologies both technically and economically viable. Although this remains a niche segment within telecom, it represents a leading edge of innovation.

Advantages of OFLs. OFLs offer three key advantages over RF feeder links:

- **High throughput:** Optical links deliver gigabit-grade throughput, starting at 10 to 100 Gbps and scaling higher. This exceeds



Figure 1 Cailabs' TILBA-OGS optical ground station capable of 10+Gbps bidirectional throughput. [Source: Cailabs]

Commercial Feasibility of OGS

Government, defense and even commercial use cases tied to Earth observation and other high data volume missions drive demand for optical ground stations. The open question is whether telecom operators can deploy optical ground stations at commercial scale for bidirectional data transport services. Commercial scaling depends on technical, operational, and economic factors. OFLs must deliver robust pointing acquisition and tracking, real time uplink compensation, and automated handover between geographically diverse sites at scale. They must also meet performance requirements such as latency while achieving a cost per gigabyte that competes with RF alternatives. The trade-offs differ materially between commercial and government markets, and among the different use cases.

RF capabilities, which are constrained by bandwidth and typically offer throughput in the range of hundreds of Mbps to a few Gbps. Throughput is especially critical in Earth observation satellite applications, among other use cases.

- **Security:** RF waves propagate over wide areas. Even the narrowest RF beams cover large footprints, making them relatively easy to detect, intercept, or jam. In contrast, laser beams are highly directive and extremely difficult to detect or intercept. They are also resistant to jamming. These characteristics make OFLs particularly attractive for government and defense applications.
- **No licensing:** OFLs are not subject to licensing requirements. RF feeder links, by comparison, must comply with

international and national spectrum licensing frameworks.

In addition to the above, optical space-to-ground links enable quantum key distribution (QKD) through satellite systems, offering a pathway to secure global communications.

Cloud Cover: A Manageable Constraint. OFLs face a known limitation: optical signals cannot penetrate clouds. This challenge is addressed through diversity of ground stations, provided they are strategically geolocated so that at least one station maintains a clear view of the sky. While this adds cost to the overall network, multiple ground stations are already integral to the architecture of low Earth orbit (LEO) constellations. The strategy becomes even more viable as optical ISLs are deployed on the satellites.

A few observations on ground station locations:

- Dry regions are more suitable for optical ground stations (OGS), while polar sites are less ideal. This contrasts with RF ground stations, where low radio noise, not weather, is the primary driver of location.
- A network of roughly a dozen ground stations can deliver close to 3-9 availability. That is a manageable number given the overall architecture and operational needs of a LEO satellite constellation.

Environmental factors once limited large-scale deployment of terrestrial free-space optical links (FSO), but ground station diversity helps overcome those constraints in the OFL context.

Atmospheric Turbulence: The Hard Problem. Atmospheric turbulence, a well-known effect in astronomy, is the single greatest challenge in OFLs. Even under clear skies, the optical signal experiences distortions as

Emerging Trends in Optical Space Connectivity

While efforts to mitigate atmospheric turbulence remain central to improving throughput and reliability in OFLs, several adjacent developments are gaining momentum:

- **Space-to-Aircraft and HAPS Links:** Satellite-to-aircraft and High Altitude Platform Station connectivity is advancing. Kepler Communications and General Atomics have demonstrated such links, while SoftBank and Japanese partners including ArkEdge Space and Kiyohara Optics plan to test satellite-to-HAPS optical links.
- **OGS Miniaturization:** Portable, lower-cost optical ground stations are expanding deployment options for LEO and multi-orbit networks. These compact terminals support flexible site selection and enable performance scaling across diverse use cases.
- **QKD Integration:** Quantum key distribution via satellite has been demonstrated by China. The next step is assessing commercial viability for secure government and enterprise applications, particularly in multi-orbit and hybrid network environments.



Figure 2 General Atomics optical terminal established communications with a Tesat optical terminal on a Kepler Communications satellite. [Source: General Atomics]



Figure 4 TeraNet TN-3 mobile OGS. [Source: TeraNet]

it passes through layers of the atmosphere that vary in temperature, pressure, and density. These variations introduce several artifacts and signal degradations:

- **Scintillation:** This resembles fading in RF signals. The intensity of the optical signal fluctuates over time. A deep fade can result in loss of information when the received signal power falls below the sensitivity threshold of the receiver.
- **Phase aberration:** The wavefront of the optical signal becomes distorted upon entering the atmosphere. Different portions of the wavefront reach the receiver at different times. The impact depends on the nature of the optical signal. For example, it can blur the focal spot, enlarge it, and produce a speckled pattern. The result is reduced power density and a degraded signal-to-noise ratio.
- **Beam wander and instability:** Turbulence can deflect the beam from its intended path. In some cases, the beam may miss the receiving telescope entirely. A Pointing, Acquisition, and Tracking (PAT) system is needed to maintain alignment.
- **Beam spreading (tip tilt):** Scattering causes the signal to spread more widely. As a result, only a fraction of the beam's power is captured by the receiving telescope.

In summary, atmospheric turbulence significantly impacts the quality of the received optical signal by altering both its intensity and phase.

Mitigation Strategies. Atmospheric turbulence affects the optical signal differently in the downlink and uplink:

- In the **downlink**, the signal travels through vacuum before entering the atmosphere, resulting in spatially correlated distortions across the wavefront. A large telescope can apply corrective measures to recover the signal.
- In the **uplink**, the beam experiences immediate distortion upon transmission, producing a speckled pattern that diverges as it propagates through vacuum. The receiver sees shifting speckles and fluctuating power.

Mitigation techniques for uplink distortions include uplink pre-compensation, advanced forward error correction codes with interleaving, and spatial diversity. Each method has strengths and limitations and must be evaluated based on the use case, such as LEO satellite systems. These techniques differentiate commercial solutions and represent areas of valuable intellectual property.



Figure 3 Starlink O-ISL performance. [Source: PCMag/Michael Kan]

Commercialization of Space Optical Links.

The space optical market, especially OFL, remains in its early stages. Government agencies such as SDA and NASA in the United States and ESA in Europe have led research and development efforts. As LEO satellite constellations expand, commercial players have entered the field and intensified competition.

The most aggressive race is unfolding in the optical ISL segment, which has matured faster than the OFL and OGS segments due to large-scale deployments in LEO satellites. Starlink, for example, equips each satellite with three OISLs that operate at up to 100 Gbps in unidirectional throughput. In January 2024, Starlink reported deploying over 9,000 units. We estimate that the current number exceeds 23,000.

The OISL market continues to attract mergers and acquisitions, but many companies face funding challenges. The capital-intensive nature of optical systems, combined with delays in constellation funding and launch schedules, has strained financial resources. Venture capital surged during the Covid-era boom in 2020 and 2021 but dropped sharply by 2022 and 2023.

Several companies are active in the OISL space, including Tesat, Mynaric, Ball Aerospace, Sony, CACI, Skyloom, BlueHalo, and Blue Marble. Mynaric illustrates the financial pressure in this sector. It filed for protective self-administration to shield itself from creditors and launched cost-cutting measures, including layoffs and the cancellation of research and product lines. Ball, which developed an OISL product with Honeywell, joined BAE through acquisition.

The OFL and OGS segment remains a niche within a niche, shaped largely by its ties to government and defense. Several companies that build optical ISL systems, such as Tesat, Mynaric, and Ball, also develop space-to-ground solutions. Because OGS platforms require integration of multiple subsystems, only a few systems integrators operate in this space. Defense contractors dominate the field since government and defense agencies continue to drive demand. One company that stands out in OGS development is Cailabs, which has raised more than 100 million dollars since its founding in 2013. Startups like BridgeComm and Astrolight are also advancing portable, bidirectional OGS solutions for multi-orbit satellite networks and LEO downlinks.

Key Takeaways

- **High throughput and security.** OFLs offer gigabit-grade throughput and are highly resistant to interception and jamming, making them ideal for defense and Earth observation missions.
- **Cloud cover is manageable.** While OFLs cannot penetrate clouds, ground station diversity through strategic placement of multiple OGS sites can maintain link availability.
- **Atmospheric turbulence is the hard problem.** Turbulence causes signal fading, beam wander, and phase distortion. Mitigation requires adaptive optics, uplink pre-compensation, and robust tracking systems. These techniques help differentiate commercial solutions.
- **Commercial scalability remains uncertain.** Defense demand drives early adoption, but commercial viability depends on automation, cost per gigabyte, and reliable bidirectional performance at scale.
- **Defense remains the anchor customer.** Most OFL and OGS deployments are tied to government and defense programs, which shape technical requirements and procurement cycles.
- **Commercial use cases are emerging.** Earth observation data offload, secure enterprise backhaul, and multi-orbit network support are promising early applications for OFLs and portable OGSs.
- **Market momentum is strongest in OISLs.** Optical intersatellite links (OISLs) have scaled rapidly, especially in LEO constellations. OFLs and OGSs lag behind but are gaining traction through defense-led R&D.
- **Integration complexity limits market entrants.** OGS platforms are complex systems—optics, tracking, and control systems. Only a few firms, often defense contractors, have full-stack capabilities.
- **New frontiers in optical connectivity are emerging.** Portable OGS terminals, satellite-to-aircraft and HAPS links, and QKD integration signal expanding use cases beyond traditional ground networks.
- **Capital intensity is a barrier to scale.** Optical systems require significant upfront investment, and delays in constellation funding or launch schedules strain financial viability.

About Xona Partners

Xona Partners (Xona) is a boutique advisory services firm specializing in technology, media, and telecommunications (TMT). Established in 2012 by a team of seasoned technologists, startup founders, managing directors in global ventures, and investment advisors, Xona leverages its founders' cross-functional expertise to offer a unique, multidisciplinary approach to technology and investment advisory services. Our clientele includes private equity and venture funds, technology corporations, regulators, and public sector organizations. We assist our clients with pre-investment due diligence, post-investment lifecycle management, and strategic technology management, helping them identify new revenue streams and navigate the complex landscape of the TMT sector.

E-mail: advisors@xonapartners.com | Web: xonapartners.com

Part 3:

Telecom Networks

Turning the Page on 5G: Service Providers' Vision for 6G

Overview. Service providers, vendors, cloud service providers, and other stakeholders gathered to define the vision, high-level requirements, and considerations for 6G networks during a workshop in Incheon, South Korea, from March 10–11, 2025. This Insight Note focuses on the most critical requirements voiced by service providers. While the views of other stakeholders are important, we will address them in a separate Note.

If there is one key theme that emerged from the two days of deliberations, which included over 230 presentations, it is simplicity. Service providers, who consider 5G “over-flexible,” called for a technology that builds on the 5G roadmap with practical, focused requirements that are achievable and monetizable. This represents a notable shift from 5G’s grand ambitions, which service providers associate with its high implementation costs, lengthy service delivery cycles, and limited monetization potential. Hence, while 5G focused on achieving impressive technical metrics, such as peak throughput, 6G will prioritize user experience, efficiency, and commercial viability. It is as though 6G aims to succeed where 5G fell short. If this proves to be the case, it would further perpetuate the so-called “curse of the odd-numbered technology”—[2G and 4G are widely considered successful, whereas 3G is regarded as a failure.]

Service providers, both implicitly and explicitly, emphasized limited capital expenditures and strict cost control. This was reflected in calls for prioritizing software over hardware through cloud-native solutions, artificial intelligence, and energy efficiency. Security was also a major concern,

Financial Realities

China Mobile provided a realistic perspective on their operations, highlighting that, despite being among the first to deploy 5G and launching 5G SA in September 2020, deploying 2.4 million 5G base stations (accounting for 33% of the global total), serving 30,000 commercial cases, and serving 544 million 5G mobile users, their revenue has remained stagnant. This financial context has been central to shaping operators’ vision for 6G, emphasizing simplicity, cost control, and a focus on applications with monetization potential.




Figure 1 China Mobile’s mobile data revenue and growth trends. [Source: China Mobile]

with some advocating for the adoption of Quantum Security.

In this Note, we highlight several key areas and examine where service providers aligned or diverged, reflecting their unique operating contexts. Broadly, service providers can be categorized into innovators, such as many in the Asia-Pacific region, who push the boundaries of technology; and pragmatists, like those in Europe, who prioritize cost control. North American operators occupy a middle ground, incorporating elements of both approaches.

It is worth noting that it is impossible to address every aspiration voiced by stakeholders, so we have focused on the highlights that will have the most direct impact on investments. Consequently, we acknowledge that several critical topics were omitted, such as the integration of Wi-Fi, the low-altitude economy, and the use of 6G within data centers.

Defining a new communication standard is a challenging process that takes years to complete. 6G standards are not expected until March 2029. This timeline is significant given the current state of world affairs, where geopolitical tensions introduce unpredictable outcomes. Such tensions could lead to the fragmentation of standards, as occurred before 4G. This remains a serious risk that must be carefully considered.

Simplification as a General Theme. Operators envision 6G with a central focus on simplification, contrasting sharply with what they perceive as the "over-flexibility" of 5G. The array of options provided by 5G prolonged infrastructure deployment timelines and increased costs for design, implementation, testing, and verification. For 6G, the motto is clear: no options! Deutsche Telekom has advocated for cutting all functions lacking commercial success. Similarly, AT&T emphasized a globally aligned "meaningful minimum set of features," reflecting frustration with the fragmentation of 5G standards. Orange and NTT Docomo also criticized 5G for offering "too many options."

Simplification as an overarching principle affects various aspects of 6G, including backward compatibility, core network architecture, protocol stack design, and features like network slicing. This approach will lead to significant changes in how 6G develops and is deployed.

In essence, operators are pushing for a streamlined set of features that are easier to

The Demographic Factor

SK Telecom highlighted declining populations as a significant threat to future revenue, compounding the challenges posed by decreasing ARPUs. Several other countries in the developed world, such as Japan and Italy, are experiencing similar demographic declines. This issue is particularly relevant for future wireless networks, as 6G standardization is not expected to be completed before March 2029, with meaningful and widespread deployments anticipated around 2031.




Figure 2 Demographic forecast for South Korea.

design, implement, and deploy. They aim to prioritize services with clear commercial value that can drive actual revenue.

Backward Compatibility. Supporting legacy technologies contributes to complexity. Therefore, operators' perspectives on backward compatibility are crucial for achieving the desired simplification. Operators align on two key points: 1) ensuring interworking with 5G and positioning 6G as an evolution of 5G technology, and 2) eliminating the need for interworking with 2G and 3G technologies. However, opinions differ regarding 4G, requiring a more nuanced approach. For example, Deutsche Telekom emphasized, "We do not need interworking with legacy 2G/3G/4G networks," while SK Telecom

highlighted the generation replacement timeline and called for a detailed evaluation of the pros and cons of 4G-6G interworking.

Radio Access Network. The radio typically serves as a highlight of any new technology generation, defining advertised capabilities such as peak throughput. It is also the most capex-intensive part of the overall network and, therefore, receives special attention.

For 6G, operators advocate for realistic and achievable targets, being critical of overly ambitious goals that are unachievable in real-world scenarios. Deutsche Telekom emphasized: “No forced 10 years radio cycle,” including, “No forced 10x lower latency; 10x higher peak data rate; 10x more devices; and 10x more accurate positioning,” which characterized 5G. Orange stated, “No need to increase requirements still not achieved by 5G deployments, e.g. peak rates, latency.” AT&T called for “6G performance improvements [to] focus on mean and cell edge user throughput instead of peak throughput,” a sentiment echoed by SK Telecom’s emphasis on prioritizing “user experienced data rate over peak data rate.” A notable outlier is China Mobile, which proposed a 100 Mbps–1 Gbps bit rate and 1–40 msec latency service requirement for a single user, though this was an exception to the general consensus.

Operators largely agree on reusing 5G foundations, building upon 5G NR waveforms, numerology, modulation, and coding to avoid hardware changes, at least in sub-6 GHz bands. However, maintaining this approach may prove difficult amidst calls for significantly improved spectral efficiency gains.

Despite some nuanced differences, there is general agreement on the importance of the upper middle-band spectrum between 6–12 GHz, with variations in the upper bounds depending on the country or region. For example, European operators, such as Deutsche Telekom, do not foresee a future

Demand for Network Sharing

Network sharing was not a major focus for operators but was mentioned by a few, including SK Telecom, China Unicom, and Oppo. We expect network sharing to gain more attention as operators look to reduce capex. SK Telecom explicitly highlighted the potential for network sharing in the upper mid-band spectrum, where coverage areas are expected to shrink further. SK also noted that Korean operators deployed 5G network sharing in rural areas using the C-band, expressing a preference for leveraging existing MOCN and MORAN solutions. They suggested that additional standardization work may be required to adapt these options for 6G.

for FR3 (7–24 GHz) in Europe, instead focusing on the 410 MHz–7.125 GHz range, particularly 6.425–7.125 GHz, for new 6G spectrum. Asian operators advocate prioritizing the upper middle band spectrum, such as 6.525–7.125 GHz in China and Japan and 7.125–8.4 GHz in Korea. They see the upper middle-band as a primary candidate for new 6G deployments, where MIMO capabilities could be scaled and bandwidth extended to 200 MHz. In contrast, support for millimeter waves was rare with a few exceptions like Rakuten, and was largely ignored by most operators, or opposed, particularly among European operators.

Operators also agree on supporting Open RAN and generally align in allowing the O-RAN Alliance to continue defining RAN interfaces. However, exceptions exist, such as Deutsche Telekom, which called for the 3GPP to take on this task (previous attempts by the 3GPP to develop such interfaces failed due to vendor opposition). Greenfield operators, like Rakuten, emphasize disaggregation, open interfaces, and interoperability. SK Telecom specifically called for a review of functional splits to determine optimal configurations for 6G. This will remain a topic to watch as

operators balance cost against performance, likely resulting in regional differences.

Core Network. Operators agree on several fundamental requirements, with the most critical being: 1) the 6G core must be standalone from day one; and 2) the 6G core should evolve from the 5G core rather than adopt a clean-slate design. For example, Verizon advocated using the 5G SA core as the baseline for the 6G core to leverage existing 5G investments. Structural simplification and reducing network functions are common demands from many operators, including SK Telecom and NTT Docomo. However, this does not rule out more ambitious visions, such as China Mobile's call for distributed autonomous subnets. NTT Docomo also proposed dual registration, enabling devices to register independently and simultaneously on both 5G and 6G networks, contrasting with the current dual-connectivity between 4G and 5G.

Additional requirements include reusing the service-based architecture, adopting cloud-native implementation, and ensuring backward compatibility with 4G and 5G. The key takeaway is that 6G will avoid multiple options for the core network, as operators aim to simplify and reduce core network deployment costs. While operators broadly agree on the overall core network architecture, evolution, and strategy, their views differ on the scope and prioritization of enhancements.

Artificial Intelligence. AI was omnipresent across all operators' requirements. While some mentioned AI in passing, others—such as SK Telecom, China Mobile, Docomo, and AT&T—treat it as a cornerstone of 6G. Asian operators appear more aggressive in implementing AI, whereas their European counterparts take a more cautious approach, focusing on specific use cases like sustainability. SK Telecom called for “AI Telco Infrastructure” and proposed various AI applications across the network, including RAN, edge, and core. This aligns with SK's

Operators: Contrasts and Commonalities

Operators in different regions demonstrate distinct priorities and challenges shaped by local market dynamics, regulatory frameworks, and deployment scenarios. Asian operators often emerge as innovators, prioritizing cutting-edge use cases and the rapid adoption of transformative technologies such as AI and ISAC. They tend to propose bold architectural and functional advancements, with a stronger emphasis on revenue generation through new services and capabilities.

In contrast, European operators adopt a more pragmatic approach, focusing on addressing real-world market needs. Their strategies emphasize energy efficiency, cost control, and practical deployment considerations, such as avoiding spectrum bands like mmWave that present significant capex challenges. This reflects a bias toward sustainability and operational simplicity while ensuring meaningful improvements over 5G.

North American operators strike a balance between these perspectives, advocating for practical deployments while exploring monetization opportunities through emerging services like AI. These regional distinctions—ranging from Asia's innovation-driven vision to Europe's cost-conscious pragmatism and North America's balanced approach—shape unique 6G strategies, reflecting diverse priorities and market realities.

strategy to monetize AI services. China Mobile's position stems from its focus on Artificial General Intelligence, which it anticipates becoming a reality beyond 2030, alongside its vision for Agentic AI as a key feature of 6G.

In summary, operators agree on the significance of AI, but questions remain about its deployment—specifically where, how, and why it should be implemented.

Energy Savings. Operators concur that 6G must achieve significant energy savings. Most emphasize the use of AI/ML to optimize power consumption, while several highlight the need for simpler protocols—referred to as lean signaling by Orange. A few operators have outlined target reductions compared to 5G; for instance, Verizon stressed a >50% reduction, and Deutsche Telekom called for >>10% energy efficiency gains. However, no specific targets have been universally established.

Non-Terrestrial Networks (NTN). Many operators emphasized the integration of NTN into 6G networks, considering it critical for extending coverage and supporting emergency services. NTN is widely seen as an extension of terrestrial services rather than a replacement. Key topics related to NTN include interoperability with LTE—which is not currently standardized—and spectrum co-existence. SK Telecom underscored the importance of including NTN in the initial 6G design and highlighted technologies such as NTN carrier aggregation and satellite MIMO, which are challenging technologies to implement. China Unicom identified NTN as a key aspect of 6G architecture design for ubiquitous coverage, advocating for unified NTN/TN control planes to enable seamless handovers. NTT Docomo and Telstra viewed NTN as essential for network resiliency.

Asian operators place greater emphasis on NTN compared to their European and American counterparts, who largely regard NTN as a niche supplemental service. Notably, US-based operators mentioned NTN only briefly.

Network Slicing. Operators widely agree that 5G network slicing was overly complex, rigid, and costly to implement. Many are advocating for a simpler design and implementation in 6G.

Security. Operators are addressing 6G security requirements by introducing significant enhancements to meet emerging

6G Deployment: Breaking from Tradition

In their vision on 6G, SK Telecom suggested that without a new killer service or explosive service/traffic demand, 6G would be deployed as a targeted technology in specific geographic areas and use cases (“Generation Mix”) as opposed to a nationwide deployment like its predecessors (“Generation Change”).




Figure 3 6G deployment scenarios as (a) “Generation change” and (b) “Generation mix.” [Source: SK Telecom]

threats and operational needs. First, a zero-trust architecture framework has been proposed to ensure security by design, shifting 6G defenses from static, network-based perimeters to a focus on users, assets, and resources. Second, quantum-safe end-to-end security has been advocated by several operators, including LG-U, AT&T, Verizon, and Deutsche Telekom, to safeguard against quantum-based threats and ensure the long-term security of 6G networks. Additionally, some operators, like AT&T, have called for AI to automatically manage 6G security, while Verizon has proposed a security framework to counter AI-based attacks. China Mobile proposed built-in agile security to enable the deployment of cryptographic primitives, algorithms, or protocols with minimal operational impact, along with fine-grained security isolation tailored to the specific needs of different businesses.

Integrated Sensing and Communication (ISAC). Operators view ISAC as a transformative technology for 6G, enabling new services and addressing emerging use cases. This stands in contrast to other applications commonly referenced in industry circles, which were almost entirely

absent from the visions operators advanced (e.g., no one mentioned remote surgery!). ISAC, however, stands out as an exception, with several operators—including AT&T, China Mobile, Deutsche Telekom, NTT Docomo, Oppo, KPN, Telenor, Verizon, and Vodafone—highlighting it in their 6G vision.

Operators see ISAC as an application with significant potential for monetization. As China Mobile explains: “With sensing, the cellular system acts as a data generator, as compared with simply the data pipe.” Consequently, there is a concerted effort to propose specific ideas related to ISAC. AT&T considers ISAC crucial for public safety and UAV detection.

Nevertheless, several challenges remain with ISAC, including the effectiveness of sensing in low-band spectrum, the physical layer

overhead, the impact on network KPIs and operations, energy efficiency, and potential regulatory updates.

Network Resiliency. Operators’ advocacy for network resiliency was largely integrated into other themes and not addressed as a standalone topic. A notable exception was among Asian operators, including NTT Docomo, Rakuten, and China Mobile, who advocated for “zero outage networks.” SK Telecom explicitly emphasized emergency data services during outages. NTN was highlighted by NTT Docomo and LG U+ as a means to enhance resiliency, while Rakuten focused on self-healing resilient RAN. Telus and Telefónica stressed the importance of end-to-end resilient design, and KT highlighted cloud-native network transformation as a key strategy for ensuring resiliency.

Key Takeaways

- **Simplification and Efficiency:** Operators advocate for a streamlined approach, avoiding unnecessary complexity and focusing on practical use cases with credible monetization potential.
- **Building on 5G Foundations:** Emphasis on evolving the 5G core, NR waveforms, modulation, and coding while minimizing hardware changes, especially in sub-6 GHz bands.
- **Spectrum Strategy:** Prioritizing upper-midband spectrum (6–8 GHz) for new deployments.
- **Security Enhancements:** Zero-trust architecture, quantum-safe protocols, AI-driven security management, and agile cryptography for resilient and adaptive defenses.
- **ISAC Integration:** Operators view ISAC as transformative, with real monetization potential, addressing public safety, sensing, and UAV detection.
- **NTN Integration:** Recognized as essential for coverage extension and emergency services, with calls for interoperability and unified control planes.
- **AI Adoption:** Asian operators lead with ambitious AI implementations, while European and American operators take a cautious, use-case-driven approach.
- **Energy Savings:** Operators advocate for significant energy efficiency gains, leveraging AI/ML, simpler protocols, and lean signaling.
- **Network Resiliency:** Advocacy for end-to-end resiliency, emergency data services, and cloud-native transformation. Asian operators put forward the concept of “zero outage networks.”

- **Cost Control:** Network sharing and focused, targeted deployments are strategies to reduce capex. This suggests that 6G may not be deployed across large footprints as was the case with previous technologies.
- **Memorable operator quotes:**
 - o **China Mobile:** "With sensing, the cellular system acts as a data generator, as compared with simply the data pipe."
 - o **Deutsche Telekom:** "No forced 10 years radio cycle. No forced 10x lower latency. No 10x higher peak data rate. No 10x more devices. No 10x more accurate positioning. No new IMT spectrum for 6G. No forced maintenance of interworking with (all) legacy technologies and services "
 - o **Orange:** "No need to increase requirements still not achieved by 5G deployments, e.g. peak rates, latency."
 - o **AT&T:** "6G performance improvements should focus on mean and cell edge user throughput instead of peak throughput."
 - o **T-Mobile:** FWA "Fallow Capacity" Model Works but at a High Price
 - o **Vodafone:** 6G must be so great that operators want to deploy it immediately in existing bands!
 - o **Telefonica:** "Without a clear path to offer new value in comm. services, supporting voice in 6G as defined today will not be sustainable by operators."
 - o **SK Telecom:** "5G lessons learned: Various visionary services expected, but no killer service! Gap between 5G vision recommendations and actual customer experience."

About Xona Partners

Xona Partners (Xona) is a boutique advisory services firm specializing in technology, media, and telecommunications (TMT). Established in 2012 by a team of seasoned technologists, startup founders, managing directors in global ventures, and investment advisors, Xona leverages its founders' cross-functional expertise to offer a unique, multidisciplinary approach to technology and investment advisory services. Our clientele includes private equity and venture funds, technology corporations, regulators, and public sector organizations. We assist our clients with pre-investment due diligence, post-investment lifecycle management, and strategic technology management, helping them identify new revenue streams and navigate the complex landscape of the TMT sector.

E-mail: advisors@xonapartners.com | Web: xonapartners.com

Fortifying Global Connectivity: Enhancing Resiliency and Strategic Investments in Subsea Fiber Cable Networks

Overview. Subsea cables are a critical global internet infrastructure and have become a focal point in geopolitical rivalries, elevating the importance of their security and resiliency to new heights. Consequently, network resiliency now plays a pivotal role in shaping cable architectures and deployment strategies. Although always an important consideration, resiliency has gained newfound prominence due to the growing involvement of hyperscalers and escalating geopolitical tensions. These factors have driven new deployments, with resiliency emerging as an important determinant of cable routes and topology.

This Insight Note addresses recent developments in subsea cable network security and resiliency, and their impact on network architecture and investments in this critical sector

Risks to Subsea Fiber Cable Networks. The growing concern over the security and resiliency of subsea cables underscores the reliance of the global economy on these critical systems, which carry around 99% of intercontinental data traffic. Events such as the 2022 sabotage of the Nord Stream gas pipeline have heightened awareness of the vulnerabilities facing subsea infrastructure. This concern has been amplified by recent cable cuts in the North Atlantic Ocean and Baltic Sea, the most recent occurring just weeks before the publication of this Insight Note. Nevertheless, it must be noted that cable cuts are not uncommon occurrences [see sidebar], but the backdrop of rising geopolitical tensions has intensified focus on the need for enhanced network security and resiliency.

Reliability of Subsea Fiber Cables.

Following the suspected Baltic seafloor cable sabotage, NATO deployed drone boats for round-the-clock surveillance of critical areas. While sabotage grabs headlines, most of the estimated 100-200 annual cable breaks get little attention. The most common cause, accounting for 80% of failures, is ships dragging anchors or fishing equipment. About 10% of failures are due to internal faults like abrasions or failed repeaters. Natural events such as volcanic activity, earthquakes, tsunamis, mudflows, and submarine landslides can also damage subsea cables. With approximately 65% of external faults occurring in shallow waters of under 100 meters, cables are typically buried at depths of 0.5 to 3 meters when deployed in shallow waters up to 1,500 meters deep.




Figure 1 Causes of subsea fiber cable failures in percent. [Source: Telegeography]

While cable cuts pose significant risks to subsea cable network infrastructure and are the most costly to repair when they occur in deep seas, they are not the only threats.

Cable landing stations (CLS) and the terrestrial manholes connecting termination stations to subsea cables are highly vulnerable to security threats. Often, it's easier to disrupt a subsea cable network by targeting its land-based infrastructure rather than the undersea segments. The FCC's recent rule review underscores the critical need to secure both terrestrial and subsea components of these networks. Therefore, CLS diversity is a critical topic that, while receiving attention during the design phase, is not a well-covered issue.

Subsea Network Resiliency Initiatives.

Different initiatives have emerged in 2024 to address resiliency in subsea fiber cables. In September 2024, the governments of the United States, EU, Australia, Canada, and other allied nations issued a [joint statement](#) emphasizing the importance of resiliency and security in subsea fiber cable networks. The statement outlined principles for designing undersea cable infrastructure with a focus on resilience, redundancy, and security. It also advocated for stronger collaboration between governments and industry on cable deployment, maintenance, and repair.

In November 2024, the [ITU](#) reinforced this priority by establishing an advisory body on submarine cable resilience in partnership with the International Cable Protection Committee (ICPC). This body aims to enhance network resiliency, reduce recovery times after disruptions, and improve redundancy in subsea cable systems.

Subsea network resiliency has also come under the scrutiny of regulators. In November 2024, the FCC issued a [Notice of Proposed Rulemaking \(NPRM\)](#), marking the first comprehensive review of subsea cable regulations since 2001. The proposed rules aim to enhance deployment efficiency while prioritizing the security, resilience, and protection of critical subsea cable infrastructure.

The European Union further emphasized the strategic importance of subsea cables when it issued recommendations on the security and resilience of submarine cable infrastructures in February 2024. The recommendations addressed improving coordination and strengthening governance to enhance cable security, increasing funding and promoting public-private partnerships, and developing preparedness plans to respond to incidents and minimize disruptions of subsea cable operations.

Ensuring Fiber Cable Network Resiliency.

Minimizing the effects of network failures on network availability and performance is at the core of resilient networks. From this standpoint, there are different approaches to improve network resiliency through network topology and operational processes.

Ensuring resiliency through topology implies connecting a node in the network through multiple paths. This is already the case with many nodes in the subsea cable network. However, economics dictate that certain paths are more economically viable than others, leaving more to be desired from a perspective of overall network resiliency.

Among the notable approaches to enhancing subsea cable network resiliency are route or path diversity and branching off a main trunk which helps to diversify cable landing stations.

Route or path diversity. While subsea cables are fundamentally point-to-point systems, they form an interconnected network that allows reaching a single point through multiple paths. Recent market developments have introduced further route diversification.

For example, the North Atlantic already hosts multiple cables connecting Europe and North America. Traffic between Europe and Latin America traditionally follows this route, although a secondary path via Africa exists. The inauguration of EllaLink in 2021

introduced a direct connection between Europe and Latin America, providing an alternative route that not only enhances resiliency but also reduces latency.




Figure 2 EllaLink provides a direct route between Europe and Latin America, bypassing the highly congested Trans-Atlantic route. This strategically positions Portugal and Brazil as key gateways to their respective regions.

Other examples of new routes include the Far North Fiber and Polar Connect projects, which connect Northern Europe with Japan. These cables are shorter than the traditional route to Asia through the Mediterranean and Red Seas, offering lower latency. They also bypass critical choke points such as the Suez Canal, the Bab al-Mandeb Strait and the Malacca Strait.

Implementing route diversity comes with its own set of challenges. Cables tend to follow similar routes due to technical, geological, financial, and regulatory reasons. While some routes are commercially justified based on the changing dynamics of data traffic storage and flow among the various regional operations of CSPs, other routes would require greater involvement from governments to realize.

A Role for Satellites?

The advent of low earth orbit (LEO) satellite constellations has prompted questions about their ability to compete with subsea fiber cables. Subsea fiber cables boast terabit-per-second capacity, vastly overshadowing the gigabit-per-second capacity of LEO satellites, which are limited by available frequency spectrum. Nevertheless, it is possible to route some critical traffic through satellites as a backup path, making the two technologies complementary solutions that enhance overall network reliability.

Another approach for satellites could involve using free-space optical technologies instead of RF spectrum for communication. These technologies utilize lasers to transmit data between the Earth and satellites. A few startups are pioneering these technologies, which could increase the capacity of satellites by an order of magnitude. However, these technologies face challenges since the performance of free-space optics is highly dependent on environmental conditions.



Figure 3 Far North Fiber cable connecting Norway to Japan.

Branching off the Main Trunk. Deploying branches off the main cable trunk serves multiple purposes. In addition to connecting various location, it provides diversity in cable landing stations.

SeaMeWe-6 between Marseille, France and Singapore is one such example where multiple branches connect various hubs, including the Middle East, Pakistan, Bangladesh, India and many other locations.

Branching off major routes makes economic sense for connecting areas with lower traffic, which otherwise would be financially unprofitable. It also serves to improve the overall financial feasibility of the project by aggregating traffic on the main trunk. As a result, we anticipate more cables will be constructed with multiple branches, particularly as new cable construction address emerging markets in Latin America, Africa and Southeast Asia.

Impact of Cloud Service Providers. Starting around 2010, cloud service providers (CSPs)—particularly Google, Meta, Microsoft, and Amazon—began constructing their own subsea cables to meet the increasing demand for bandwidth. Today, traffic from these CSPs accounts for approximately 70% of total subsea cable traffic. This development has had a significant impact on the subsea cable market in various ways.

CSPs deployed subsea cables along routes that connect their data centers, considering both performance and redundancy. As cloud infrastructure continues to evolve, additional routes are being deployed and planned for the future, incorporating redundancy for existing routes as well as route diversity. To ensure uninterrupted service, CSPs have overbuilt some routes, such as the North Atlantic connecting Europe and North America, and the North Pacific connecting the U.S. with Japan, both of which carry the highest amount of data traffic.

As a result of CSPs' extensive subsea cable construction, the price of traffic has dropped. This has crowded out private investments, as it has become challenging to achieve the expected 10-20% return on investment that these investors seek. The focus on network resiliency by CSPs has

further influenced cable routes and topology, prioritizing high-redundancy and diverse routes to enhance overall network reliability.

Private Investments in Subsea Fiber Cables.

Private investments in subsea fiber cables have been sporadic over the past decade. The dotcom bubble crash in the early 2000s dampened investments in that decade. The entry of cloud service providers (CSPs) in the 2010s crowded out private investors as the price of capacity on high-traffic routes dropped. Nevertheless, there are pockets of investment remaining.

Certain high-traffic routes, such as those in the North Atlantic, remain appealing to investors. Private investors also invest in regional cables, which often receive government subsidies, leading to public-private partnerships for financing and operating these cables. Additionally, niche markets present investment opportunities. For example, Tampnet, which operates fiber-optic cables servicing oil and gas platforms in the North Sea and the Gulf of Mexico, has received private equity investments over the years.

More recently, private equity investors have shown interest in cable deployment assets. For instance, KKR invested in Optic Marine Systems to expand its fleet of cable-laying vessels, a market that has seen little investment since the dotcom bubble crash.




Figure 4 Private investments in subsea fiber cables. [Source: Infralogic]

Key Takeaways

- **Geopolitical Rivalries:** Subsea fiber cables are critical to global internet infrastructure and have become focal points in geopolitical rivalries. This has elevated the importance of their security and resiliency to new heights.
- **Regulatory Emphasis:** Governments and regulatory bodies are increasingly emphasizing network resiliency. This includes initiatives to improve the security and reliability of subsea cables, encouraging investments in resilient infrastructure.
- **Architectural Impact:** Network resiliency now plays a pivotal role in shaping cable architectures and deployment strategies. Ensuring uninterrupted service influences the design of routes, aiming to enhance network robustness and reduce single points of failure.
- **Redundant Pathways:** Resiliency is enhanced by deploying multiple redundant paths and routes, ensuring there are alternative pathways for data traffic in case of disruptions. Establishing additional cable landing stations is critical, as these stations can be more vulnerable to vandalism compared to cutting the cables themselves.
- **Role of CSPs:** Cloud service providers (CSPs) like Google, Meta, Microsoft, and Amazon have driven the focus on network resiliency by constructing their own subsea cables. Their involvement has prioritized redundant and diverse routes to ensure continuous service. However, the drop in price of capacity has crowded out private investments in the subsea cable ecosystem.
- **Overbuilding Routes:** To ensure uninterrupted service, CSPs have overbuilt some high-traffic routes, such as the North Atlantic and North Pacific, enhancing network robustness and reducing the risk of single points of failure. However, investments by CSPs into subsea fiber has crowded-out private investments.
- **Public-Private Partnerships:** Collaborations between governments and private investors, including public-private partnerships, will be essential to support the development of resilient subsea cable networks. These partnerships will help lower the barriers to entry for the private sector into routes that would otherwise not be developed.
- **Specialized Investments:** Private equity investors are also focusing on niche markets and specialized assets, such as cables servicing oil and gas platforms.
- **Contingency planning:** We urge service providers to assess network resiliency from a holistic regional perspective. It is essential to develop comprehensive contingency plans to ensure continuity in the event of network outages. By addressing potential vulnerabilities and implementing robust strategies, service providers can better safeguard against disruptions and enhance the reliability of their services.

About Xona Partners

Xona Partners (Xona) is a boutique advisory services firm specializing in technology, media, and telecommunications (TMT). Established in 2012 by a team of seasoned technologists, startup founders, managing directors in global ventures, and investment advisors, Xona leverages its founders' cross-functional expertise to offer a unique, multidisciplinary approach to technology and investment advisory services. Our clientele includes private equity and venture funds, technology corporations, regulators, and public sector organizations. We assist our clients with pre-investment due diligence, post-investment lifecycle management, and strategic technology management, helping them identify new revenue streams and navigate the complex landscape of the TMT sector.

E-mail: advisors@xonapartners.com | Web: xonapartners.com

AI Bandwidth Bubble: Unpacking Mobile Traffic's Uncertain Path Forward

Overview. AI traffic already has a well-defined impact on long-haul and data center interconnect networks. However, as adoption accelerates, its effect on access networks, particularly mobile networks, remains far less clear. Because traffic directly shapes investment in network dimensioning, architecture, and spectrum, this uncertainty raises two central questions: what volume of traffic AI will generate, and what type of traffic it will produce.

The answers carry weight. They will determine how service providers and equipment vendors design networks to balance cost efficiency with performance requirements, while also influencing regulatory strategies for new spectrum. In short, AI-driven traffic could significantly affect the next wave of capital investment, with implications across silicon vendors, equipment manufacturers, regulators, and service providers alike.

That said, the impact of AI on mobile networks remains speculative, tied closely to the evolution of applications. Past investment cycles have shown the limits of prediction in such contexts. Against this backdrop, the purpose of this Insight Note is to examine how leading ecosystem players, spanning silicon vendors, equipment manufacturers, and service providers, are projecting these trends and the rationale behind their positions. Our goal is to establish a framework for analyzing potential AI traffic scenarios, enabling ecosystem participants to develop informed perspectives and plan strategically for the eventualities AI traffic may bring.

Declining Growth of Mobile Traffic

Forecasting mobile traffic trends is essential for planning network capacity, phasing capital expenditures, and shaping spectrum strategy. The 2000s marked a turning point: the rise of social media platforms (Facebook in 2004, YouTube in 2005), the advent of mobile Internet devices (iPhone in 2007), and the deployment of LTE in 2009 together fueled triple-digit annual growth in mobile data traffic.

Today, the landscape is different. Mobile penetration is near saturation, and video consumption has matured. Traffic continues to grow, but at a slower pace—global annual growth now averages around 20%, with country-level rates declining between 3–8% per year.

One cautionary note: fixed wireless access is expanding in many markets, particularly the U.S., and this application skews mobile traffic growth upward. Understanding these dynamics is critical for operators and regulators as they plan the next phase of investment and spectrum allocation.




Figure 1 Year-over-year mobile traffic growth rate in select markets. [Source: Xona Partners]

Mobile Traffic Trends and Investment Implications.

Mobile networks have undergone a full transformation from voice-centric to data-dominant systems over the past 20 years. These two decades were marked by large investments in network and spectrum assets, enabling the migration from 3G to 5G and the release of extensive spectrum holdings that generated billions for national regulators. Today, this trend is maturing as growth in mobile data slows and traffic approaches saturation (see box on declining growth in mobile traffic).

The deceleration in growth is directly linked to the type of traffic and consumer behavior. Video, both on-demand and live streaming, together with social media, account for the majority of mobile data. The number of hours a subscriber can devote to video and social media is finite. With mobile penetration now at saturation, meaning that anyone who can afford a service is already subscribed, traffic trends point toward a plateau in growth.

The implications are significant. Much of the capital budget of service providers is directed toward network upgrades and expansions, and data traffic is the justification for these investments. Without sustained growth in traffic, the business case for new investment weakens.

Table 1 Top application traffic categories in mobile and fixed access networks. [Source: Sandvine]

	Mobile Networks		Fixed Networks	
	Downlink	Uplink	Downlink	Uplink
Social Media	35%	23%	18%	9%
Video	31%	10%	39%	13%
File Sharing	7%	15%	9%	28%
Television	1%	0%	11%	2%
Gaming	7%	4%	7%	4%
Communications	4%	16%	2%	12%
Other	15%	32%	14%	32%
Traffic Ratio	9:1		9:1	

There was an expectation that demand for traffic would rise in parallel with 5G deployments. That demand has not materialized. Applications such as Virtual Reality, Augmented Reality, and extended

Saving Networks from Video Dominance

The mobile internet is dominated by video traffic, which accounts for 60-75% of total data usage. As video consumption surged, social media and streaming platforms focused on bandwidth efficiency, aiming to deliver the best quality of experience (QoE) while conserving capacity.

Today, platforms transmit video using only 30-50% of the bandwidth required a decade ago. Netflix illustrates this evolution with per-title encoding, tailoring bitrates to content complexity and cutting them by 10-20%. The company advanced to per-chunk and per-shot encoding, applying different encoders to sets of frames. AI-driven dynamic optimization further refined encoder selection, achieving another 20% reduction in bitrate.

Codec innovations also saved bandwidth. Providers used H.264/AVC in the early 2010s, shifted to H.265/HEVC and VP9 in the mid-2010s, and in 2020 Netflix began streaming AV1, which delivers 30-40% bitrate savings compared with H.264. These advances improved QoE while reducing bandwidth requirements.




Figure 2 Type of traffic carrier by mobile networks. [Source: Ericsson]

reality more broadly (AR/VR/XR), along with traffic from vertical markets including cloud gaming and autonomous vehicles, have not achieved the anticipated commercial traction. Their contribution to traffic remains marginal across both mobile and fixed networks.

AI Traffic: Training vs. Inference. The rise of generative AI (GenAI), particularly with the launch of ChatGPT in November 2022 and subsequent investments in AI infrastructure, has raised questions about the impact of AI on connectivity networks, including metro, long haul, and access networks. Market evidence shows that while AI is fundamentally reshaping long haul networks, including submarine systems, its impact on downstream networks is less certain and more ambiguous. This is especially true for access networks, including mobile. To explain this distinction, it is useful to differentiate between AI training and inference workloads. AI training runs in centralized data centers that process massive volumes of data in contexts that are not sensitive to latency. In contrast, AI inference workloads process smaller volumes of data in distributed data centers where latency and proximity to the user are critical, depending on the requirements of the application. Most attention and investment in AI infrastructure has focused on training, including high speed data center interconnects ranging from 100 to 400 Gbps, with targets of 800 Gbps and beyond in the future.

AI inference, by comparison, is lagging behind training, reflecting the slower uptake of AI applications. Inference traffic will appear in metro and access networks, which means the impact of AI on mobile traffic is directly tied to inference workloads and the applications that drive them. This highlights one of the core challenges in projecting traffic impact, since the applications that will ultimately succeed are not yet known.

Most current GenAI applications, such as ChatGPT sessions, graphics generation, and

AI and the Future of Mobile Demand

Nokia and Ericsson provide forecasts for mobile AI traffic based on specific contexts and assumptions. The GSMA offers a broader view, estimating that mobile data traffic will increase threefold by 2030 compared with 2024 levels, with AI potentially contributing 20-80% of that growth. Omdia projects 2031 as the crossover point when global AI network traffic will exceed conventional traffic, and expects net new AI mobile traffic to grow at a compound annual rate of 130% between 2025 and 2033.

Table 2 AI impact on mobile traffic - vendor estimates.

Forecast	Quantitative Forecast
Nokia 2025	30% is AI's share of wide area network traffic by 2034 (all types)
Ericsson 2025	100% uplink traffic growth driven ~20% AI/AR glasses penetration.




Figure 3 Total network traffic, AI and non-AI. [Source: Omdia, 2025]

A close analysis of forecasts suggests that AI may not overwhelm access networks with new traffic but instead reallocate existing patterns. Access network operators should approach traffic forecasting with caution, considering substitution effects and evolving user behaviors.

video creation, are not demanding in terms of latency or even in bandwidth, so they do not need to run close to the user. Time

sensitive and interactive applications, however, will require low latency and will be pushed closer to the user, including into the access network and even onto user devices. This has potential implications for mobile network architecture and for silicon requirements at the device level.

The challenge therefore centers on three factors: the type of application and its performance parameters, the location of the corresponding AI workload, and the rate of user adoption. As a result, projections of AI traffic impact on access networks, both in timing and magnitude, remain speculative.

AI Applications and the Uplink Shift. The mobile ecosystem foresees several AI applications as potential drivers of traffic growth. These applications mirror those once expected to gain traction with 5G but failed to do so. Industry players now argue that integrating AI will enhance their utility and make them more compelling to users:

- AI chatbots
- AI-assisted cloud applications, such as fitness or shopping
- AI-powered gaming
- AI-powered XR applications

- Autonomous vehicles, humanoid droids, and drones

Equipment vendors including Ericsson, Nokia, and Huawei, along with market research firms such as Omdia and GSMA Intelligence, state that these and other AI applications will drive a surge in uplink traffic (see AI and the future of mobile demand box). They assume that mobile devices will act more as sources of data, in contrast to today's downlink heavy traffic pattern. Consequently, mobile traffic would reflect a more balanced pattern.

Vendors also have a vested interest in raising concerns about explosive uplink traffic because that narrative supports sales of 5G Advanced equipment, which is designed to improve uplink performance. Forecasts do not suggest that uplink traffic will dominate overall usage, but rather that its share of total traffic will rise. In our view, networks are not typically constrained on the uplink and should be able to absorb the initial growth, particularly since most sub-2 GHz bands operate in frequency division duplex mode.

The Nature of AI-Driven Traffic. AI agents could generate more persistent and “chatty”

Table 3 AI impact of mobile networks.

Dimension	Impact description	Operational implication
Uplink traffic	Higher uplink share as devices act more as data sources	Strengthen uplink capacity planning, revisit power control and UL MIMO strategies
Signaling load	AI agents generate persistent, chatty sessions with frequent small messages, setup and teardown, keep alive, and mobility events	Scale signaling capacity, refine mobility and session policies, optimize network architecture
Latency sensitivity	Interactive AI services demand lower and tighter latency bounds	Deploy edge compute, prioritize low latency paths, enhance timing and synchronization
Mobility and session churn	Frequent micro interactions increase handovers and session churn	Tighten handover parameters, optimize session lifecycles
Caching efficiency reduction	AI traffic may not cache efficiently due to personal and dynamic characteristics	Increase origin capacity, refine cache policies

traffic. This activity would affect not only information load but also signaling flows in both directions. Frequent small messages, session setup and teardown, device keep-alive, and mobility events increase signaling and place stress on multiple parts of the network, including radio access and core elements. This effect differs from pure throughput pressure and requires targeted solutions, potentially including a re-architecture of the network.

AI traffic may also reduce caching efficiency. There is a view that its personal and dynamic nature makes it difficult to cache effectively. As a result, Content Delivery Networks (CDNs) could operate less efficiently, which would increase transport volumes.

These traffic characteristics support the case for evolving mobile network architecture to incorporate edge computing. Doing so would help maintain latency requirements, reduce backhaul load, and better manage AI-related data and signaling traffic.

Key Takeaways

- **Uncertainty in AI's Mobile Impact:** While AI clearly drives traffic in long-haul and data center networks, its effects on mobile access networks remain speculative, hinging on unknown application types and adoption rates, and traffic volumes, which could reshape network investments, architecture, and spectrum strategies.
- **Slowing Mobile Data Growth.** Mobile traffic is maturing and plateauing due to saturation in penetration and finite consumer time for dominant video/social media usage; expected 5G-driven surges from AR/VR, cloud gaming, and verticals like autonomous vehicles have not materialized, weakening the business case for ongoing capital expenditures.
- **Training vs. Inference Distinction.** AI training occurs in centralized data centers with high-bandwidth, latency-tolerant needs, while inference, which is key for mobile, requires distributed, low-latency processing closer to users; current GenAI apps (e.g., ChatGPT) are bandwidth-light and latency insensitive, but future interactive ones could push workloads to edges and devices.
- **Forecast Reliance on Revived Applications:** Most industry forecasts for AI-driven traffic growth are based on AI-enhanced versions of applications like AR/XR, cloud gaming, and autonomous vehicles that were anticipated to surge with 5G but failed to gain meaningful commercial traction; this revival narrative, promoted by vendors and analysts, highlights the speculative risks in projections.
- **Uplink Traffic Shift.** AI applications (e.g., chatbots, AI-enhanced gaming/XR, autonomous systems) are projected to increase uplink data shares by making devices more data-generative, balancing downlink-heavy patterns; vendors like Ericsson/Nokia promote 5G Advanced upgrades, though networks can likely absorb initial growth.
- **Unique AI Traffic Characteristics.** AI may produce "chatty" patterns with frequent small messages, boosting signaling loads (e.g., sessions, mobility events) over pure throughput; it could also erode CDN caching efficiency due to personalized/dynamic content, necessitating edge computing to manage latency, backhaul, and overall network stress.
- **AI may reallocate rather than flood traffic.** A close analysis of forecasts suggests that AI may not overwhelm access networks with new traffic but instead shift existing traffic patterns. Operators should treat forecasts cautiously and model substitution effects and evolving user behavior.
- **Network architecture will need targeted changes.** The impact of AI differs from pure throughput stress and may require solutions such as signaling optimization, control plane scaling, and selective re-architecture rather than only adding raw capacity.

About Xona Partners

Xona Partners (Xona) is a boutique advisory services firm specialized in technology, media and telecommunications. Xona was founded in 2012 by a team of seasoned technologists and startup founders, managing directors in global ventures, and investment advisors. Drawing on its founders' cross-functional expertise, Xona offers a unique multidisciplinary integrative technology and investment advisory service to private equity and venture funds, technology corporations, as well as regulators and public sector organizations. We help our clients in pre-investment due diligence, post investment lifecycle management, and strategic technology management to develop new sources of revenue.

E-mail: advisors@xonapartners.com

Web: xonapartners.com