Implications of an increasingly local internet

Creating value as traffic flows adapt to AI demand

GLF Board focus

How AI-driven traffic will reshape traffic patterns and the implications for international carriers

Introduction

Demand from AI, cloud, content, and data-intensive enterprise workloads is driving a growing share of digital processing closer to where data is generated, consumed, or regulated. In this paper, "the internet becoming more local" refers not only to public internet traffic, but to the broader localisation of digital workload flows across public internet, private interconnect, cloud on-ramps, data-centre-to-data-centre connectivity, and enterprise network environments. This shift is creating a multi-tier digital infrastructure architecture: large-scale AI training, bulk storage, and global platforms remain concentrated in a limited number of hyperscale locations, while inference, data processing, content delivery, and regulated enterprise workloads increasingly distribute into metro, regional, and on-premise environments.

Multiple forces are converging to drive this shift. AI is an important accelerant, particularly through latency-sensitive inference and data-intensive training, but localisation is not an AI-only phenomenon. It is also being driven by cloud architecture, content distribution, enterprise resilience requirements, data sovereignty, power availability, grid constraints, permitting timelines, and the economics of moving large volumes of data over long distances. Together, these factors are changing where digital workloads are processed and where connectivity value is captured.

Industry indicators support this rebalancing of workload geography. Analysts estimate that by the mid-2020s, most enterprise-generated data will be created and processed outside traditional centralised data centres and public cloud regions, reflecting the growth of edge and distributed architectures. Hyperscalers are extending their platforms into dozens of metropolitan areas through local zones and similar constructs, explicitly targeting low-latency and data-residency use cases. At the same time, global data centre electricity consumption is forecast to roughly double by 2030, highlighting that localisation is occurring alongside continued overall growth, not instead of it.

Recent empirical studies show that distributing AI inference closer to data sources can reduce end-to-end latency by up to 90[1]% and cut bandwidth usage for upstream traffic up to 60%[2], underlining that localisation is not only desirable but increasingly necessary for real-time services. Typical cloud-only architectures deliver 100ms+ round-trip latency, whereas well-designed edge deployments regularly achieve 5–20ms[1], enabling industrial automation, immersive applications, and time critical decisioning that are impossible at higher latency levels.

For international carriers, three conclusions are particularly relevant:

1. **The internet is becoming more local in structure, not smaller in scale.** Global traffic volumes and compute demand continue to expand, driven by AI, cloud adoption, and digitalisation across sectors. What is changing is the shape of that traffic: more east–west flows within metropolitan regions, greater regionalisation of processing, and increased

[1] IJITRS, October 2025 – (link) a benchmark study comparing edge vs cloud AI in a distributed test bed found average latency of 4.2ms at the edge versus 31.7ms for cloud-centric processing, an ~86.7% reduction in round-trip latency when inference is performed closer to the data source

[2] IAEME, January - February 2025 – (link) An engineering-focused review of cloud–edge real-time systems reports that local edge filtering and processing can reduce network bandwidth usage by ~30–60% in industrial IoT deployments by minimizing unnecessary upstream data transmission to the cloud

diversity in where value is created. Hyperscale hubs remain essential, but they are increasingly complemented by a dense fabric of metro and regional nodes.

2. **Control points are moving closer to users and enterprises.** As workloads distribute, relevance shifts toward assets that sit at the intersection of connectivity and compute: dense metro fibre, interconnection-rich data centres, cloud on-ramps, and low-latency access to AI inference. For individual carriers, this means that proximity, performance assurance, and ecosystem integration matter more than purely raw asset footprint or network reach. Long-haul connectivity remains critical, but its value becomes more corridor-specific, tied closely to where large-scale compute clusters and replication flows are located. AI application developers, their customers, and infrastructure partners will have significant influence over the shape of the network.

3. **Fragmentation risk is rising, making industry collaboration more critical.** Localisation driven by sovereignty, regulation, and operational constraints increases the risk of a more fragmented global internet, characterised by inconsistent rules, bespoke interconnections, and higher operational complexity. More than 100 jurisdictions now apply some form of data localisation or sovereignty requirement, including stringent regimes in the EU, India, China, parts of the Middle East, and other major markets. Without coordination, the industry risks replacing a globally interoperable system with a patchwork of local optimisations. For the international connectivity ecosystem, this elevates the importance of collaboration on standards, APIs, assurance models, and shared approaches to compliance and auditability.

The strategic “so what” for GLF members is therefore twofold. At the individual carrier and operator level, consideration might be given to rebalance investment toward metro density, interconnection, and performance-led services, while embedding sovereignty and power considerations into network design. At the industry level, organisations including GLF have a critical role to play in preserving global interoperability. As the internet’s centre of gravity comes more local, industry frameworks for service automation, performance assurance, and cross-border compliance are essential to ensure that localisation enhances resilience and innovation rather than fragmenting the global digital economy.

An evermore local internet

Defining “the internet coming local”

The phrase “the internet is coming local” refers to the structural shift by which compute and storage are placed within the network close to end users. Instead of relying solely on a small number of hyperscale facilities requiring long-haul connectivity, workloads are increasingly distributed across three categories of location:

- **Hyperscale hubs** – large, power-dense campuses supporting AI training, bulk storage, and global platforms.
- **Metro and regional data centres** – hosting inference, content distribution, enterprise applications, and regulated workloads.
- **On-premise and near-edge locations** – factories, hospitals, campuses, and network edges where latency, resilience, or sovereignty requirements dominate.

This evolution is driven by latency sensitivity, regulatory constraints, cost and energy considerations, and the rapid growth of AI inference workloads. Importantly, **localisation complements rather than replaces** global infrastructure.

In this paper, “internet” is therefore used as shorthand for the broader digital connectivity fabric. Some localised flows will traverse the public internet, but many of the highest-value flows will use private interconnect, cloud connectivity, data-centre fabrics, enterprise networks, or carrier-managed services. This distinction matters commercially: public-internet traffic may be monetised differently from private connectivity, cloud on-ramps, wavelengths, dark fibre, or assured enterprise services. The strategic issue for carriers is not simply whether traffic remains “on the internet,” but where digital value is exchanged, who controls the interconnection point, and which party can assure performance, compliance, and resilience.

Drivers of a more local internet

The shift toward a more localised digital infrastructure architecture is not the result of a single technology change. AI is accelerating the trend, but it sits alongside cloud distribution, content delivery, enterprise resilience, sovereignty, power, and network-economics drivers that collectively reshape how and where digital workloads are deployed.

Latency and real-time performance requirements are the most immediate drivers. AI inference, real-time analytics, immersive media, industrial automation, and mission-critical enterprise applications require deterministic, low-latency performance that is difficult to guarantee when workloads are served exclusively from distant hyperscale regions. Locating compute and storage closer to users, machines, and data sources reduces round-trip latency, improves service quality, and lowers reliance on long-haul transport for time-sensitive traffic.

Data sovereignty and regulatory constraints are a second, increasingly structural force. Governments and regulators are expanding data protection, residency, and sovereignty frameworks across jurisdictions, particularly for personal, financial, healthcare, and public-sector data. These requirements often mandate that data be processed and stored within

national or regional boundaries, driving demand for local and in-country compute, storage, and connectivity even when global cloud platforms are used.

AI economics and workload characteristics further reinforce localisation. While large-scale model training remains concentrated in power-rich hyperscale campuses, inference workloads scale horizontally and are highly sensitive to latency, network cost, and energy efficiency. Distributing inference closer to demand reduces backhaul traffic, lowers egress costs, and enables more responsive AI-driven services.

Infrastructure and sustainability constraints also play a critical role. Power availability, grid capacity, permitting timelines, and community acceptance are increasingly binding constraints on large, centralised data centre developments. Distributing workloads across metro and regional facilities can mitigate these constraints while improving resilience and energy efficiency.

Together, these indicators show that growth in edge devices, AI inference, and colocation capacity is outpacing that of traditional centralised infrastructure, while data centre electricity use is becoming a binding constraint in several mature markets.

Evidence that workloads are coming more local

This shift is now visible in changes in infrastructure build patterns, traffic flows, and deployment models. Taken together, these trends point to a structural redistribution of where data is processed and exchanged across the network. The table below sets out some relevant data indicators that suggest the on-going trend towards a more distributed local internet.

Exhibit 1: relevant indicators demonstrating ongoing internet localisation

Metric	2025	2030	25-30 CAGR	Source
Connected IoT devices (bn) *(proxy for edge data generation)*	21.1	39.0	13.2%	IoT Analytics
Global DC electricity consumption (TWh) *(proxy for infrastructure constraint)*	477	945	14.6%	IEA
AI Inference market (USD bn) *(proxy for local AI deployment)*	106.1	255.0	19.2%	MarketsandMarkets
Edge AI market (USD bn) *(proxy for inference market near edge)*	24.9	66.5	21.7%	Grand View Research
CDN market (USD bn) *(proxy for localisation via caching)*	26.5	45.1	11.3%	Mordor Intelligence
Network API market (USD bn) *(proxy for network interoperability)*	2.0	6.1	25.7%	MarketsandMarkets
DC colocation market (USD bn) *(proxy for neutral interconnect ecosystem)*	96.2	204.4	16.3%	MarketsandMarkets

Five indicators are particularly relevant for international carriers:

1. Enterprise data is increasingly processed outside centralised cloud regions

Industry analysts consistently highlight that enterprise data processing is moving away from a small number of centralised cloud and hyperscale data centres. Gartner estimates that by the mid-2020s, most enterprise-generated data will be created and processed outside traditional data centres and public cloud regions, reflecting the growth of edge, on-premise, and distributed architectures. This shift is driven by latency requirements, resilience considerations, and regulatory constraints, particularly in sectors such as manufacturing, healthcare, financial services, and the public sector.

From a network perspective, this trend directly increases demand for metro and regional connectivity, as data flows become more localised and east–west in nature. Rather than data traversing long-haul networks to reach distant cloud regions, processing increasingly occurs within metropolitan or national boundaries, reinforcing the importance of dense fibre networks and proximity to compute.

2. Hyperscaler expansion into Metro and Edge locations

Major cloud providers have extended their infrastructure beyond core regions into dozens of metropolitan areas through constructs such as local zones, edge regions, and on-premise extensions. As of the mid-2020s, AWS Local Zones alone are deployed across more than 30 metro areas, explicitly targeting low-latency workloads, data residency requirements, and hybrid enterprise architectures.

Whilst this expansion complements, rather than replaces, hyperscale regions, it materially changes traffic patterns by anchoring cloud access and application execution closer to users. For carriers, this reinforces the strategic importance of cloud on-ramps, metro interconnection, and last-mile fibre access as primary value drivers.

3. The emergence of neoclouds

The emergence of neoclouds reflects growing demand for AI-native compute that can be deployed closer to data, users, and regulatory boundaries. Focused on GPU and accelerator access, providers such as CoreWeave and Lambda are building distributed footprints across regional data-centre hubs rather than concentrating capacity in a small number of hyperscale regions. This model supports the localisation of AI compute driven by latency, data sovereignty, and power constraints.

For international carriers, neoclouds are strategically important customers: their distributed architectures increase demand for high-capacity inter-DC connectivity, compliant cross-border transport, and federated network services linking regional AI nodes into a coherent global fabric.

4. AI workload bifurcation: centralised training, distributed inference

AI is a critical accelerant of localisation. While large-scale model training remains concentrated in a limited number of power-rich hyperscale campuses, AI inference workloads are increasingly distributed to reduce latency, network cost, and energy consumption. Inference supports real-time decision-making in applications ranging from industrial automation and autonomous

systems to personalised digital services, all of which benefit from compute located closer to users or machines.

This bifurcation is reflected in market dynamics. The edge and distributed computing market is estimated to reach approximately US$168 billion in the mid-2020s, growing to around US$249 billion by 2030, representing a sustained growth rate of roughly 8% CAGR[3]. These figures underscore that distributed processing is becoming an integral part of the compute landscape, not a niche overlay.

5. Data growth, IoT, and infrastructure constraints reinforce local processing

Underlying these shifts is continued growth in data generation at the network edge. Global IoT connections are forecast to reach approximately 21 billion by 2025, rising to around 39 billion by 2030[4], significantly increasing the volume of data produced outside traditional data centre environments. Processing this data locally reduces backhaul traffic, improves responsiveness, and mitigates network congestion.

At the same time, infrastructure constraints reinforce localisation. Global data centre electricity consumption is projected to rise from roughly 460 TWh in the early 2020s to around 945 TWh by 2030[5], while in the United States alone, peak data centre power demand could reach up to 106 GW by the mid-2030s under aggressive growth scenarios[6]. These constraints make it increasingly challenging to concentrate all workloads in a small number of mega-campuses, encouraging more distributed deployment models.

[3] MarketsAndMarkets – Edge Computing Market
[4] IOT Analytics – State of IOT 2025
[5] IEA – Energy Demand from AI
[6] Bloomberg – AI and the Power Grid

Implications for Fibre Operators

What will change?

The localisation of internet workloads materially alters the demand profile, economics, and operating environment for fibre network owners. As compute and storage move closer to end users, enterprises, and machines, end-to-end fibre quality-of-service requires not only long-haul transport infrastructure, but metro-scale digital ecosystems that support AI inference, cloud access, and data-intensive applications.

First, **demand increases significantly within denser metropolitan fibre infrastructure**. Edge and distributed computing models rely on high-capacity, low-latency connectivity between end users, access networks, data centres, and cloud on-ramps. This drives growing requirements for middle-mile and metro fibre, particularly to interconnect clusters of data centres and edge sites. Market indicators support this shift. The global optical fibre connectivity market is forecast to grow at a compound annual rate of around 9% through the early 2030s[7], reflecting sustained demand for high-speed, low-latency connectivity.

Second, **traffic patterns become more localised**. Rather than traversing national or international backbones to reach distant cloud regions, an increasing share of traffic circulates within metropolitan areas between enterprises, local data centres, and edge nodes. For fibre owners, this changes network utilisation profiles and places greater emphasis on route diversity, resiliency, and proximity to compute assets, rather than sheer long-haul reach. Hybrid architectures, combining centralised cloud regions with metro and edge locations, further increases the importance of resilient local fibre connectivity to support workload failover and continuity.

Third, **new revenue opportunities emerge alongside increased infrastructure intensity**. Fibre owners are well positioned to monetise this shift through leasing dark fibre or providing high-performance lit services to hyperscalers, data centre operators, and enterprises requiring deterministic connectivity for AI inference and real-time applications. Public-sector programmes and regulated industries provide additional demand catalysts. In the United States, initiatives such as the Broadband Equity, Access, and Deployment (BEAD) programme are expanding opportunities for fibre investment, while sectors such as financial services, healthcare, and government increasingly require compliant, locality-aware connectivity solutions.

However, localisation also **raises structural constraints and operational complexity**. Edge proliferation increases requirements for fibre densification at a time when supply chains are under pressure. Industry reports indicate shortages in optical networking equipment, with some vendors reporting extended lead times. At the same time, power availability has become a binding constraint, potentially impacting expansions linked to edge compute sites. This, in turn, affects the viability and timing of associated fibre deployment.

[7] GMInsights – Optical Fiber Connectivity Market

What could this mean for Fibre Operators' strategies?

For fibre operators, the increasing localisation of internet workloads requires a strategic reorientation that goes beyond incremental network expansion. The shift toward metro-centric, latency-sensitive, and regulation-aware architectures changes where capital should be deployed, how services are defined, and how value is captured across the ecosystem.

First, **capital allocation might rebalance toward metro density and interconnection adjacency**. While long-haul connectivity remains essential, its growth profile is increasingly corridor-specific and tied to the location of hyperscale training clusters. By contrast, sustained demand growth is emerging in metropolitan rings, middle-mile routes, and direct connectivity into data centre campuses and cloud on-ramps. Fibre operators that prioritise dense, resilient metro footprints – rather than solely maximising geographic reach will be better positioned to capture the traffic and revenues associated with AI inference, hybrid cloud, and distributed enterprise workloads. For international carriers this means that where they lack access networks offnet agreements become critical to ensure timely and commercially viable network access for their customers.

Second, **offnet procurement models will need to evolve**. Traditional off-net procurement models, built around static capacity commitments, bilateral negotiations, and long provisioning lead times, are increasingly misaligned with the requirements of AI-driven traffic patterns. As AI inference and distributed training workloads scale unevenly across regions and shift dynamically between data centres and edge locations, customers will increasingly expect connectivity that can be activated, resized, and performance-assured in near real time. Meeting these expectations requires international carriers to fundamentally rethink how off-net capacity is sourced and managed. This includes moving toward federated procurement frameworks, underpinned by standardised APIs, that allow carriers to discover routes, provision capacity, and enforce SLAs seamlessly across partner networks. It also requires greater transparency on performance, automated assurance, and integrated settlement mechanisms to support on-demand services. Without these changes, carriers risk becoming bottlenecks in AI connectivity chains rather than enablers of agile, distributed AI infrastructure.

Third, **product strategy can evolve from capacity-led offerings to performance- and outcome-based services**. As workloads move closer to end users, customers place greater value on deterministic latency, availability, and resilience than on headline bandwidth. Fibre operators can therefore develop services that are explicitly defined by end-to-end performance parameters across a connectivity fabric, such as latency bounds, jitter, diversity, and failover characteristics. This shift supports premium pricing models and differentiates fibre assets in an environment where raw capacity is increasingly undifferentiated.

Fourth, **sovereignty and compliance must be embedded by design**, not treated as bespoke add-ons. Data localisation requirements are becoming a persistent feature of the operating environment, particularly for government, financial services, healthcare, and critical infrastructure customers. Fibre operators need network architectures, operational processes, and assurance mechanisms that can demonstrate jurisdictional compliance, auditable routing, and controlled interconnection without undermining global service consistency. This capability will increasingly influence vendor selection and long-term partnership decisions.

Fifth, **power and local infrastructure partnerships become a strategic lever, not only a constraint-management exercise**. No single fibre operator can independently solve for power availability, permitting, end-to-end international connectivity, and regulatory engagement across all markets. As power becomes one of the binding constraints on localised digital infrastructure, carriers can create advantage by partnering earlier and more directly with utilities, municipalities, data-centre operators, hyperscalers, and local authorities. This could include joint planning of fibre and power availability, co-investment in grid-adjacent digital infrastructure zones, use of municipal rights-of-way to accelerate deployment, and sustainability-linked partnerships that align connectivity expansion with local energy and community objectives. In this model, carriers move beyond reactive network build-out and become active participants in local digital infrastructure planning. In parallel, industry collaboration through common APIs, service automation frameworks, and assurance standards remains critical to avoid fragmentation as networks and workloads distribute

Finally, **operational excellence and sustainability emerge as sources of competitive advantage**. Managing a larger number of distributed fibre assets increases operational complexity and cyber risk, elevating the importance of automation, monitoring, and security by design. At the same time, environmental and community considerations increasingly influence project viability. Fibre operators that proactively engage stakeholders, invest in energy-efficient technologies, and align expansion plans with local sustainability objectives will reduce delivery risk and improve long-term returns.

In summary, the strategic imperative for fibre operators is clear: success in a more localised internet will depend less on scale alone and more on density, performance, compliance, and collaboration. Those that adapt their strategies accordingly will be positioned to create long-term value and remain an indispensable ecosystem participant as the architecture of the internet continues to evolve.

Implications for International Connectivity

How could this impact international connectivity?

First, **traffic growth becomes more uneven and increasingly corridor-specific**, altering the economics of international connectivity. Aggregate international traffic volumes continue to rise, driven by large-scale AI training, the expansion of global digital platforms, and ongoing requirements for data replication and synchronisation across regions. However, as workloads localise, a growing proportion of application traffic will be terminated and processed within national or regional boundaries rather than traveling across global networks end-to-end.

This shift concentrates incremental growth on a smaller number of strategically critical international corridors that connect major hyperscale, AI training, and data gravity hubs relevant to AI inference use-cases. Traffic growth is therefore less evenly distributed and more sensitive to changes in compute siting, regulatory policy, and enterprise workload placement decisions. For international carriers, this increases variability in capacity utilisation and heightens the risk of both congestion on key routes and under-utilisation elsewhere.

As a result, returns on capital are increasingly determined by the precision with which subsea systems, terrestrial backbones, and metro interconnection assets are aligned to these evolving traffic corridors. Static network planning assumptions based on broad regional growth are no longer sufficient. Instead, investment outcomes depend on a granular understanding of how traffic is generated, where it is processed, and how it will flow between local, regional, and global layers of the network.

In this environment, deep visibility into traffic flows on specific routes and corridors becomes a non-negotiable strategic capability. Carriers must develop more sophisticated forecasting, analytics, and partnership models to inform build, buy, and collaborate decisions. Without this, the industry risks misallocating capital, reinforcing volatility, and weakening its ability to support the next phase of global digital growth as the internet becomes more distributed in structure.

Second, **off-net procurement models will need to evolve toward instant, always-on, and performance-assured service models.** Traditional off-net procurement models, built around static capacity commitments, bilateral negotiations, and long provisioning lead times, are increasingly misaligned with the requirements of localised, AI-enabled, and cloud-integrated traffic patterns. In this context, value accrues disproportionately at locations where access networks, metro and long-haul fibre, data centres, and cloud platforms converge. These interconnection-rich environments increasingly determine performance outcomes, cost structures, and customer experience, rather than the characteristics of any individual network segment in isolation.

Dense metropolitan interconnection ecosystems therefore play a growing role in mediating access between local processing environments and global networks. They function as aggregation points where traffic is exchanged, optimised, and redirected across providers, geographies, and compute layers. As workloads localise, these hubs become essential not only for latency-sensitive applications such as AI inference and real-time analytics, but also for resilience, enabling workload mobility and failover across local, regional, and global resources.

For international carriers, proximity to and integration with these ecosystems increasingly influences relevance to enterprise and cloud customers.

This shift elevates the strategic importance of interconnection policy and commercial design. Three ecosystem models are likely to compete:

- **Carrier-led interconnection:** carriers use their fibre reach, operational relationships, and regulated-market credibility to manage access between local processing environments and broader regional or global networks. This model can support performance assurance, route diversity, and compliance-aware connectivity, but carriers must invest in automation, neutral access models, and faster provisioning to avoid becoming a bottleneck.

- **Data centre-led interconnection:** neutral data centre and colocation providers act as the aggregation layer where carriers, clouds, enterprises, and content platforms meet. This model benefits from ecosystem density, neutrality, and customer choice, but it can shift bargaining power away from network operators if carriers are reduced to access providers into someone else's platform.

- **Cloud-led or vertically integrated interconnection:** hyperscalers and large platforms internalise more traffic within proprietary backbones, edge zones, and private interconnect fabrics. This model can deliver rapid innovation, integrated customer experience, and strong automation, but it risks reducing openness, limiting multi-cloud flexibility, and fragmenting global interoperability around closed ecosystems.

The strategic challenge for international carriers is not to resist all forms of vertical integration, but to ensure that open, neutral, and interoperable interconnection models remain commercially viable. GLF can support this by promoting common demarcation models, service assurance standards, API-based ordering and provisioning, and compliance frameworks that allow local interconnection to scale without locking customers into closed ecosystems.

Third, **fragmentation risk increases materially as sovereignty and localisation requirements proliferate across jurisdictions**. Governments are introducing increasingly detailed rules governing data residency, routing, security controls, and lawful access, often with limited alignment across borders. While each requirement may be rational in isolation, their cumulative effect is to introduce friction into cross-border connectivity and to complicate the delivery of seamless, end-to-end services for multinational enterprises and global platforms.

As localisation increases, the industry risks drifting toward a patchwork of bespoke, jurisdiction-specific network designs and operating models. In this world, international carriers would be required to adapt routing, interconnection, monitoring, and assurance processes on a country-by-country basis, increasing operational complexity and cost. Over time, this fragmentation would erode economies of scale and slow service innovation weakening the efficiency and resilience of the global connectivity ecosystem.

Importantly, this is not a cyclical regulatory challenge that can be managed through incremental adaptation. It is structural and cumulative: once local requirements are embedded in network

architectures and operating models, they are difficult and costly to unwind. Without coordination, successive layers of regulation risk hard-coding fragmentation into the fabric of the internet itself, undermining the concept of global ubiquitous interoperability.

Fragmentation also makes partnership formation harder. As local rules, security expectations, routing constraints, certification requirements, and commercial practices diverge, carriers face higher onboarding costs, longer due-diligence cycles, and greater uncertainty when stitching together multi-operator services. This can slow the formation of the very partnerships required to serve localised workloads at scale. The risk is therefore not only that the internet becomes technically fragmented, but that the commercial fabric of international connectivity becomes harder to assemble. GLF should explicitly debate how far common trust, identity, compliance, and service-assurance mechanisms can reduce this friction while preserving each operator's commercial independence.

For international carriers, this elevates fragmentation from a compliance issue to a system-level strategic risk. Addressing it requires more than bilateral solutions; it demands collective approaches to interoperability, compliance frameworks, and assurance mechanisms that allow local requirements to be met without breaking global connectivity. This reinforces the need for coordinated industry action to ensure that sovereignty and localisation enhance trust and resilience rather than accelerating structural fragmentation of the internet.

What does this mean for international collaboration?

In this environment, **industry collaboration becomes a strategic necessity rather than a discretionary activity**. No single carrier can preserve end-to-end service consistency in an increasingly localised and regulated ecosystem.

First, carriers must collaborate on **common technical and operational standards**. As networks and workloads distribute, interoperability increasingly depends on shared APIs, service orchestration frameworks, and performance assurance models. Alignment in these areas reduces integration friction, supports automation, and enables scalable, cross-border services despite growing local complexity.

Second, collaboration must extend to **shared approaches to compliance and assurance**. Rather than each operator developing bespoke solutions for sovereignty and auditability, the industry would benefit from common frameworks that define how compliance is demonstrated, monitored, and validated across interconnected networks. This is particularly important for multinational enterprises and public-sector customers that require consistent treatment across jurisdictions.

Third, a further prerequisite is a **trusted digital identity model for operators and ecosystem participants**. As services become more federated and compliance becomes increasingly mutualised across partners, customers and counterparties will need confidence that each operator is who it claims to be, is qualified to deliver the relevant infrastructure or service, and can be trusted within a multi-party service chain. GLF should therefore consider supporting a decentralised digital identity and certification model for international connectivity, aligned where

possible with existing industry work such as Mplify's decentralised digital identity initiatives. Such a model could allow operator identity, service qualification, compliance status, and certification credentials to be verified by trusted third parties without requiring each bilateral relationship to recreate the same assurance process.

Fourth, **cross-ecosystem collaboration becomes essential**. Many of the constraints shaping localisation—power availability, permitting, sustainability, and community acceptance—sit outside the traditional control of connectivity providers. Effective responses therefore require closer coordination between carriers, data centre operators, hyperscalers, utilities, regulators, and local authorities to align infrastructure planning, investment cycles, and delivery timelines.

To make this actionable, GLF should focus on a limited set of practical interoperability building blocks rather than attempting to define a broad new architecture from scratch. The near-term roadmap should include:

1. **Service automation and API alignment:** identify which existing carrier API frameworks can be used or extended for localised, cross-border services, including product discovery, quotation, ordering, provisioning, assurance, and settlement.

2. **Performance and assurance baselines:** define a minimum set of service attributes for latency, jitter, availability, diversity, restoration, and telemetry so that localised services can be stitched into end-to-end customer propositions without bespoke integration for every market.

3. **Trust and compliance baselines:** define a common minimum security and compliance baseline that can be recognised across participating operators, with regional overlays for jurisdiction-specific requirements.

4. **Digital identity and certification:** establish a decentralised operator identity and qualification model so that participants in multi-operator service chains can be trusted, authenticated, and certified by third parties.

5. **Priority use cases:** test these mechanisms on a small number of high-value use cases, such as compliant cloud on-ramp extension, local inference connectivity, inter-data-centre metro connectivity, and regulated enterprise cross-border connectivity.

As the internet becomes more local, the challenge for international carriers is not simply to carry more traffic, but to preserve global interoperability while value shifts closer to users, enterprises, and regulated data. This is a moment for the industry to shape the next architecture of connectivity: open rather than fragmented, automated rather than bespoke, trusted rather than opaque, and resilient by design. GLF can play a central role by aligning carriers and ecosystem partners around the standards, identity, assurance, and collaboration models needed to keep a localising internet globally connected.