

Paper 1 of 6



# The 10G Wall: Why Access Networks Are Running Out of Road

The access network faces a structural capacity problem that no amount of traffic engineering resolves. This paper examines why — and what the architecture of a viable upgrade path must look like.

## The Challenge in Brief

This paper focuses on aggregation and access backhaul — not core transport or DCI. That distinction matters, because the upgrade economics, equipment constraints, and traffic profiles are fundamentally different at the network edge from those in the core.

Access networks were dimensioned for a different era. The SFF-slot equipment running 10G DWDM optics across aggregation rings, mobile backhaul, and enterprise access points was designed for traffic profiles that no longer reflect operational reality. Three forces have broken that model: the densification of 5G backhaul, the emergence of AI inference workloads at the edge, and the steady redistribution of compute away from centralized hyperscale data centers toward geographically distributed edge sites.

None of this was unpredictable. What was underestimated — systematically — was the speed of the transition and the depth of the installed base that would need to absorb it. The result is a structural capacity problem that traffic engineering cannot paper over. It requires an architectural response.



## The Demand Is Structural, Not Cyclical

The conventional framing of access network capacity exhaustion treats it as a forecasting problem — operators underestimated traffic growth and need to catch up. That framing is wrong, and it leads to the wrong solutions.

The demand surge hitting access networks today is structural: it reflects a fundamental shift in where compute lives, not simply how much data end users are consuming. STL Partners projects the total edge computing addressable market will grow from \$9 billion in 2020 to \$445 billion by 2030, and forecasts that the number of edge data center sites will triple between 2023 and 2025. These are not workloads being streamed from hyperscale data centers to end users. They are workloads being redistributed — from the center of the network toward its periphery — and the backhaul and aggregation infrastructure sitting between those edge sites and the broader network was never designed to carry them.

"The total edge computing addressable market will grow from \$9 billion in 2020 to \$445 billion in 2030. The number of edge data center sites will triple, from 491 in 2023 to 1,584 in 2025."

— STL Partners

Cisco's data gives the scale of what that redistribution means in practice: 85 zettabytes of useful raw data were created in 2021, but only 21 zettabytes were stored and processed in data centers. The gap between data generated and data processed is the demand signal that access networks must carry. It is not a traffic spike. It is a permanent step change in the volume and geographic distribution of workloads — and it is accelerating.




Figure 1: Near-peak bandwidth demand vs. 10G link capacity by use case. Sources: Ericsson Microwave Outlook 2022 (5G — 5–20 Gbps per urban site); SynaXG / NVIDIA MWC 2026 (Edge AI — 36 Gbps measured on AI-RAN platform); CableLabs DOCSIS 4.0 specification (MSO — 10 Gbps/node × 8 nodes per hub). Near-peak planning conditions in dense deployments. Arycs Technologies analysis.

### Three Specific Drivers — and Why Each One Breaks 10G Differently

It is worth being precise about what is actually driving the demand, because each driver stresses access infrastructure in a distinct way.

**5G backhaul densification.** 5G mmWave deployment creates highly localized, asymmetric demand spikes that legacy backhaul was not dimensioned for. A single sector in a dense urban deployment can push sustained throughput that saturates a 10G uplink during peak periods. The problem is not average load — it is the ratio between peak and average, which is far higher in 5G than in 4G. Operators planning deterministic capacity for 5G backhaul cannot plan to 10G line rate; when latency, jitter, and queue depth are constrained, operators must stay well below line rate, compressing usable headroom further.

**Edge AI inference.** The redistribution of AI workloads to the edge is not a future scenario. Industries from manufacturing to healthcare are deploying real-time inference at edge sites today — predictive maintenance, patient monitoring, video analytics, anomaly detection. These workloads require low-latency, deterministic connectivity between edge compute nodes and their data sources. The connectivity requirement is not simply bandwidth; it is the combination of bandwidth, latency, and reliability that 10G links at high utilization cannot reliably deliver.

**Edge data center interconnect.** As edge data center sites multiply — and STL Partners' tripling forecast implies a very rapid buildout — the interconnects between them become a primary constraint. Direct detect technology was adequate for short-reach intra-campus links. It is not adequate for the 50km-plus distances between edge sites in regional and metro topologies, where external amplifiers and dispersion compensators add complexity and cost that operators are actively trying to eliminate.

*"400ZR shipments will dominate in edge applications. The access market needs a simple, pluggable, low-cost upgrade to the 10G DWDM optics it has been using for years."*

— **Cignal AI, Transport Applications Report Q4 2021**

### The Installed Base Is the Constraint — and Forklift Logic Makes It Worse

The depth of the 10G SFF installed base is consistently underestimated in upgrade planning. Aggregation nodes, backhaul rings, and enterprise access equipment built around SFF slots have operational lives that regularly exceed original depreciation assumptions. This equipment is not about to be replaced on a schedule — it will be replaced when capacity forces the issue or when a planned refresh cycle coincides with the need.

The problem with the forklift response — full chassis and line card replacement — is not that it fails to solve the capacity problem. It is that it solves the wrong version of the problem at the wrong time. A multi-year equipment refresh program, requiring coordinated maintenance windows, new management systems, operator retraining, and hard service cutovers, consumes operational capacity that most access network teams simply do not have in reserve. And it introduces service risk at precisely the moment when the network is already under capacity stress.

The more damaging pattern is the interim workaround: operators add parallel 10G links to buy headroom, doubling fiber lease costs and management complexity, deferring the real upgrade decision while the cost of not upgrading compounds. By the time the forklift program is finally approved and resourced, the cost of the workaround period has often exceeded what a more surgical upgrade would have cost at the outset.

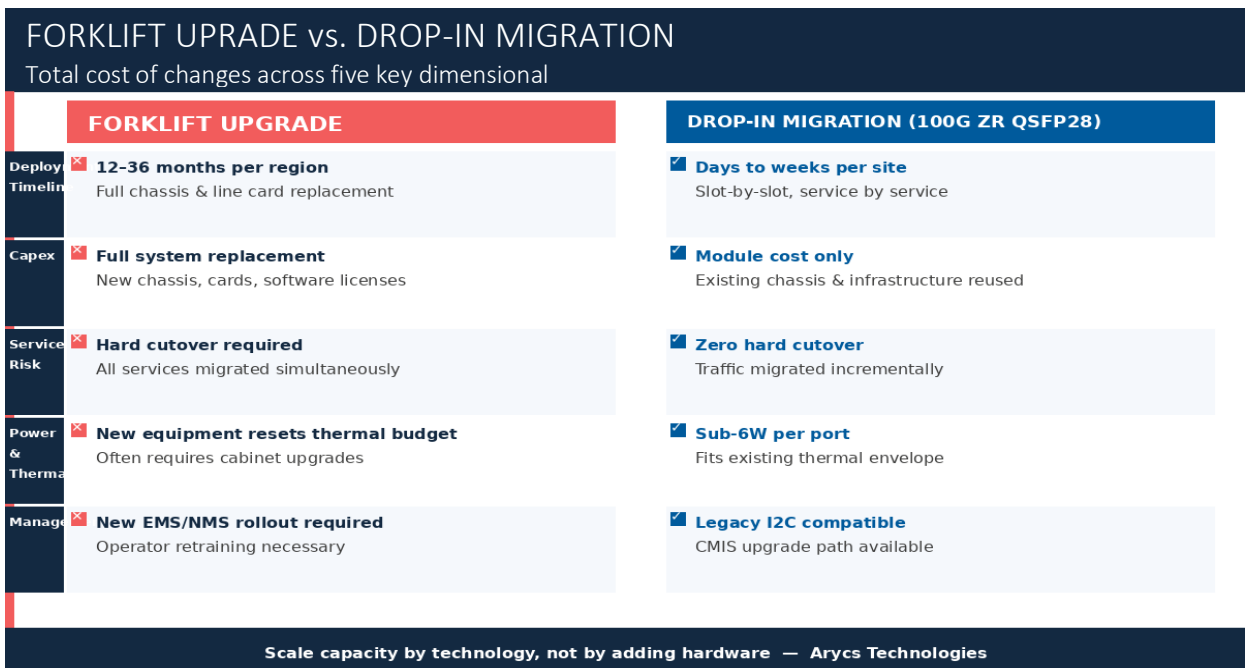


Figure 2: Forklift upgrade vs. incremental drop-in migration — a comparison across five operational dimensions.

## What a Viable Upgrade Architecture Must Satisfy

The structural nature of this demand shift means that point solutions — adding capacity here, patching a link there — do not address the underlying problem. An upgrade architecture for the access domain needs to satisfy several conditions simultaneously.

It must be compatible with the existing installed base: new optics should drop into existing equipment slots rather than requiring chassis or line card replacement. It must operate within the power and thermal constraints of access environments, which are fundamentally different from data center environments. It must support both legacy management interfaces — to function in equipment that predates CMIS — and modern telemetry interfaces, to support the AI-managed, zero-touch network architectures that operators are building toward. And it must be based on open standards with genuine multi-vendor interoperability, because operators who traded one form of vendor lock-in (legacy DWDM proprietary gear) for another (single-vendor pluggable ecosystems) have not actually solved their strategic procurement problem.

These are not new requirements. They are the requirements that have consistently not been met by the options available in the access domain — until recently. The combination of silicon photonics integration, access-optimized DSP design, and the maturation of coherent standards for sub-100G applications has created a solution class that did not exist five years ago. The next paper in this series examines the technology shifts that made it possible.

[#NetworkArchitecture](#) [#AccessNetworks](#) [#CoherentOptics](#) [#LegacyMigration](#) [#AIScaling](#)

---

## About Arycs Technologies

Arycs delivers power-efficient, coherent-class optical connectivity based on silicon photonics, coherent DSP, and advanced optical architectures. Our solutions provide industry-leading bandwidth per watt, deterministic performance, and flexible network evolution for AI, cloud, telecom, and edge infrastructure. Designed for real-world deployment, Arycs Technologies enables networks to scale with growing AI demand without disruptive redesign or hardware replacement.

[arycs-tech.com](http://arycs-tech.com) | [LinkedIn: Arycs Technologies](#)