




## Executive Briefing

# WHY LAYER 1 NEEDS NAAS: A GUIDE FOR ENTERPRISES

Network-as-a-Service (NaaS) offers enterprises more than flexibility – it can unlock material, measurable financial gains. This report explores the quantitative benefits of Layer 1 NaaS across three real-world scenarios.



# Executive Summary

Layer 1 connectivity provides the physical backbone for enterprises where network performance is a critical enabler for business outcomes. These enterprises are typically very large organisations – from global banks to cloud providers and large media platforms - who require ultra-secure, deterministic, high-capacity infrastructure to ensure business continuity, regulatory compliance and deliver competitive differentiation.

Layer 1 connectivity has traditionally been provided to enterprises as a static long-term contract. As workloads become more dynamic and resource-intensive – driven by AI adoption, decentralised infrastructure and cloud-native operating models – static provisioning exposes fundamental inefficiencies. Enterprises are forced to overbuy capacity for peak traffic resulting in costly underutilised capacity, while fixed long-term contracts and lengthy provisioning cycles hinder innovation and agility.

Network-as-a-Service (NaaS) in L1 networks is emerging as a response to this challenge. By delivering connectivity in a more programmable model, NaaS provides enterprises the same flexibility, visibility and control from their networks as they experience with their cloud environments. The qualitative benefits are well understood: more agile scaling, real-time observability, and direct configuration control. However, a critical component of the enterprise business case is missing: a clear and compelling evidence linking these qualitative benefits to measurable business outcomes.

This report seeks to address that gap, exploring the benefits that NaaS at Layer 1 can deliver to enterprises, and how these benefits translate into material, financial impact. We illustrate the quantitative value of NaaS through three enterprise scenarios grounded in real-world enterprise experiences:

- **Scenario 1: Scaling connectivity for dynamic network demand.** NaaS enables enterprises to provision high-capacity links only when needed, reducing the need for overprovisioning and underutilised capacity. For a global CDN provider supporting high-traffic events STL estimates that NaaS can reduce annual L1 connectivity costs by 53%, delivering \$130,000 in cost savings.
- **Scenario 2: Accelerating innovation without lock-term lock-in.** Shorter provisioning cycles and on-demand scalability reduces costly commitments from enterprises, accelerating innovation cycles and time to market. In practice, a large retail bank seeking to pilot a GenAI application could reduce time-to-deployment by eight months with NaaS, bringing forward productivity gains and revenue gains worth an estimated \$830M within the first 12 months.
- **Scenario 3: Building network resilience for distributed global operations.** Through end-to-end visibility and automated recovery, NaaS can reduce the frequency and duration of unplanned network outages. A global financial services firm operating a major regional hub in a different market could reduce network-outage related productivity losses by 65%, saving an estimated \$4.2M annually.

# Table of Contents

Executive Summary ..... 2

Unlocking the Layer 1 NaaS opportunity ..... 5

    Layer 1 connectivity: The foundation for mission critical enterprises..... 5

    Enter Network-as-a-Service (NaaS), a response to enterprise demand for greater flexibility, visibility and control..... 6

    Layer 1 NaaS: Becoming a reality..... 7

Quantifying the value proposition of Layer 1 NaaS ..... 9

    Scenario 1: Scaling connectivity for dynamic network demand ..... 9

    Scenario 2: Accelerating innovation without long-term lock-in ..... 10

    Scenario 3: Building network resilience for distributed global operations..... 12

Conclusion..... 14

    Recommendations for enterprises ..... 14

    Recommendations for service providers..... 14

Appendix ..... 15

# Table of Figures

Figure 1: Layer 1 connectivity target enterprise profile..... 5

Figure 2: STL Partners’ definition of NaaS..... 6

Figure 3: Estimated annual cost savings from usage-based bandwidth scaling ..... 10

Figure 4: Timeline from idea to launch for GenAI application ..... 11

Figure 4: Timeline from idea to launch for GenAI application ..... 11

Figure 5: Estimated financial impact on a large bank launching a GenAI customer care agent in first 12 months ..... 12

Figure 6: Estimated annual cost savings in productivity loss from network outages ..... 13

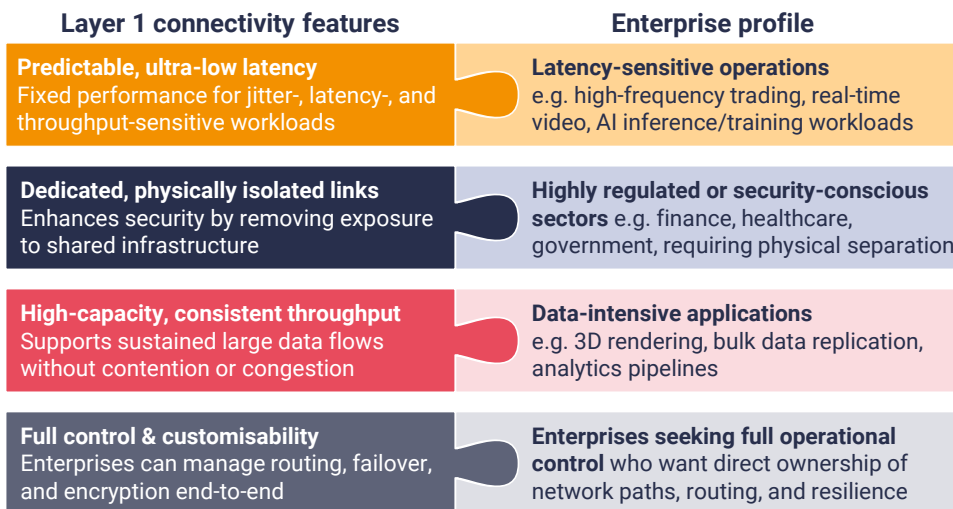
# Unlocking the Layer 1 NaaS opportunity

## Layer 1 connectivity: The foundation for mission critical enterprises

For enterprises with performance-critical, data-intensive and compliance-driven operations, Layer 1 connectivity provides the physical backbone while allowing the enterprise to retain control of network design, topology, security and protocols. These dedicated optical links are ultra-low latency, high bandwidth and highly secure making them ideal for enterprises who require full control over performance, security and data sovereignty.

Whilst Layer 1 underpins much of the global internet and cloud backbone, it is also used by select enterprise segments that depend on deterministic, high-performance connectivity. These are typically very large organisations with skilled in-house teams for whom network performance is a key competitive differentiator. Common adopters include large media streaming platforms (e.g. Netflix), cloud and managed service providers, global banks and digital-native manufacturers (e.g. Tesla) – all of whom require secure, low-latency, high-throughput infrastructure to support their business-critical workloads. Typical enterprises that choose to consume connectivity through Layer 1 network services are outlined in Figure 1.

**Figure 1: Layer 1 connectivity target enterprise profile**



Source: STL Partners

Enterprises relying on Layer 1 connectivity are often pioneers in digital experiences and so are often the first to feel the friction between legacy infrastructure models and evolving digital demands. Several key trends that these enterprises are facing include:

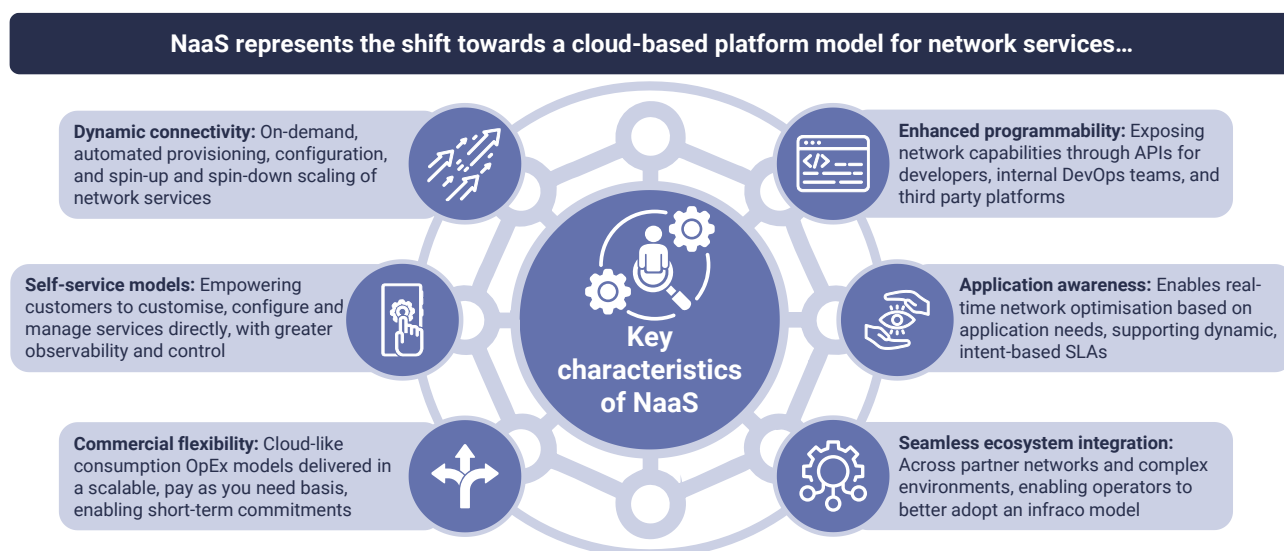
- AI and data-intensive workloads becoming more dynamic and distributed, with training, inference, and storage increasingly split across multiple locations. This shift creates demand for bursty, high-throughput links and on-demand scalability that traditional networking models struggle to support.

- Increasingly decentralised infrastructure – as enterprises scale across data centres, cloud regions and edge sites – creating a need for low-latency, high-resilience interconnectivity across geographies.
- Growing expectations for cloud-like orchestration mean enterprises want to manage their physical network infrastructure with the same automation, visibility, and control they experience in the cloud. Traditional static provisioning and long-term fibre contracts no longer align with these demands.

## Enter Network-as-a-Service (NaaS), a response to enterprise demand for greater flexibility, visibility and control

In response to changing enterprise expectations, service providers have accelerated investment in Network as a Service (NaaS) which marks a shift away from static, legacy models. STL Partners defines NaaS as the shift toward a cloud-based platform model for network services – see Figure 2.

**Figure 2: NaaS provides a reimagined business model for enterprise customers**



Source: STL Partners

Together, these elements illustrate a reimagined business model for network services, enabling enterprises to consume connectivity as a more dynamic, flexible and programmable service. In practice, this delivers greater:

- **Flexibility:** Enterprises can adjust network services dynamically – scaling bandwidth or spinning up new connections in line with shifting workload demands. This reduces overprovisioning, shortens time-to-deployment, and supports faster infrastructure expansion.
- **Visibility:** Real-time observability gives enterprises clearer insight into network performance, availability, and cost – enabling faster root-cause diagnosis, better capacity planning, and stronger SLA enforcement.

- **Control:** With APIs and self-service portals, enterprise teams can directly configure and manage their own network paths, integrating connectivity decisions into automated workflows (e.g. DevOps pipelines, AI model orchestration), and reducing dependency on third-party ticketing or manual intervention.

## Layer 1 NaaS: Becoming a reality

However, not all NaaS is equal: while many providers position their offerings under the NaaS label, they deliver only partial overlay NaaS functionality without full programmability across all network layers. Most commercial offerings today are focused on Layer 2 and 3 services – such as virtual Ethernet, SD-WAN, cloud interconnect, and managed routing. These are well suited to application and IT layer needs but fall short of the performance and infrastructure control required by Layer 1 use cases.

Historically, this means enterprises that depend on Layer 1 continue to face traditional connectivity limitations:

- Static provisioning forces overbuying for peak demand, driving up costs and leaving capacity idle for most of the year.
- Limited programmability restricts operational agility, despite high expectations for control.

This is in part driven by the technical complexity in automating the physical layer. Those providing Layer 1 networks encounter several challenges:

1. **Automating a traditionally manual process:** Provisioning fibre circuits has long relied on manual workflows and bespoke configurations. Whilst automating at Layer 2 and 3 is abstracted from certain physical constraints, deploying NaaS at Layer 1 requires rethinking how physical connections are managed and configured.
2. **Integrating fragmented systems and vendors:** OSS/BSS platforms, fibre infrastructure, and equipment from multiple OEMs must be unified into one orchestration layer. The heterogeneity of hardware and software at Layer 1 makes seamless integration a major challenge.
3. **Adhering to regulatory compliance:** Regulations on fibre infrastructure, data flows, and service delivery vary across markets. Encoding compliance into an automated, on-demand model adds significant complexity unique to Layer 1.

However, some providers are actively working to close this gap. For example, Lightstorm, an Indian interconnect specialist, has deployed robotic arms to automate optical connections in data centres, enabling true self-service provisioning at Layer 1.<sup>1</sup>

This opens up new possibilities for enterprises that purchase Layer 1 network services – particularly those for whom performance, security and infrastructure control are non-negotiable. In the following

---

<sup>1</sup> For more detail, read our case study "[Productising NaaS: A case study from India's most disruptive player](#)"

section, we explore how NaaS consumption models can equate to tangible financial outcomes for enterprises.

# Quantifying the value proposition of Layer 1 NaaS

While enterprises recognise the qualitative benefits of NaaS – flexibility, agility, visibility – the financial case remains underexplored.

In this section, we outline three enterprise scenarios that illustrate the financial impact of adopting Layer 1 NaaS. Each scenario is grounded in a real-world enterprise need, detailing the limitations of traditional connectivity models and the material benefits enabled by NaaS programmability and flexibility.<sup>2</sup>

## Scenario 1: Scaling connectivity for dynamic network demand

Network traffic is becoming increasingly volatile; patterns can be shaped by AI-driven applications, growing digital content consumption, and demand surges tied to specific events. As bursty traffic becomes more common, traditional connectivity models that can only handle this via overprovisioning becomes an expensive and inefficient way to manage capacity. NaaS offers more granular control and flexibility to enterprises over how and when they scale, without being locked into fixed 12–24-month contracts.

**Enterprise situation:** A global CDN provider supports several high-traffic events annually – including live OTT sports streaming, monthly OS rollouts for mobile OEMs, and enterprise software distributions. These events trigger periodic bandwidth spikes across their metro and long-haul backbone links, requiring high-throughput capacity for short, intensive periods.

**Enterprise challenge with legacy connectivity:**

- Static legacy models require purchasing peak bandwidth year-round, resulting in underutilised capacity and sunk cost.
- Alternatively, underprovisioning risks degraded performance and reputational damage from downstream partners (such as OTT content providers).
- Fixed contract terms and billing cycles do not align with the seasonal or event-driven nature of traffic peaks, limiting operational flexibility.

**With NaaS:**

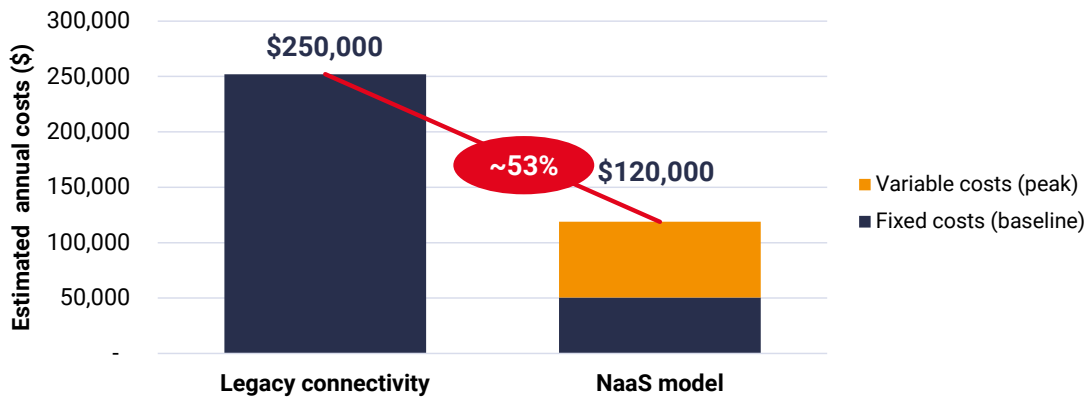
- Enterprises can instantly provision high-capacity links for peak events and deactivate them afterwards.
- Capacity can be scaled across multiple links and geographies, billed on an hourly basis.
- Avoids long-term overprovisioning and aligns spend with usage during high-traffic events.

<sup>2</sup> For more detail on our methodology behind the following enterprise scenarios, see Appendix.

**Cost impact for enterprise:** Reduces total network spend by eliminating the need to overprovision. Enterprises pay only for the capacity they use, when they need it – improving cost efficiency and agility.

**Figure 3: Estimated annual cost savings from usage-based bandwidth scaling**

Estimated annual backbone connectivity costs for large global CDN provider supporting event-driven traffic within a single country market (e.g. India). With NaaS, a CDN can reduce costs by 53% by scaling bandwidth on a usage-based model.



Source: STL Partners analysis

**Secondary cost impact:** Under-provisioning for high-traffic events risks degraded performance, with financial consequences for downstream CDN customers like OTT content providers. Even small drops in Quality of Experience (QoE) during major events (e.g. the IPL final or Super Bowl) can cause viewers to stop watching. We estimate that 5% session drop-off due to buffering could result in nearly 100 million lost ad impressions – equating to c. \$6 million in lost ad revenue per event.

Additionally, degraded network performance (e.g. jitter) can prevent an OTT’s ability to deliver context-specific ads – such as those personalised by viewer location or behaviour. To avoid disrupting the stream, platforms may default to simpler or static ad formats, which carry lower cost-per-mille (CPM) rates. This reduces the overall revenue they can command from advertisers, compounding the financial impact of performance issues.

## Scenario 2: Accelerating innovation without long-term lock-in

Most enterprises want to innovate and experiment with new technologies like GenAI, but legacy connectivity models can make it difficult for enterprises that use Layer 1 network solutions to trial and scale these initiatives without significant upfront investment. Long provisioning cycles and rigid contracts increase the barrier to innovation, particularly when connectivity costs can’t be right sized

to match trial periods. NaaS offers enterprises greater agility and faster time-to-innovation to unlock competitive advantage and accelerate productivity gains.

**Enterprise situation:** A large bank is piloting a GenAI-powered customer chatbot to improve financial advice, customer support, and real-time engagement. To train and deploy the model, the bank must connect its secure DC-DR (data centre-data recovery) environment to a neocloud offering GPU-based compute optimised for GenAI workloads. Its primary cloud provider cannot meet this need due to limited GPU availability in compliant geographies, as strict data regulatory compliance prevents financial data from leaving local infrastructure - making public cloud options unviable without a direct, secure, link.

**Enterprise challenge with legacy connectivity:**

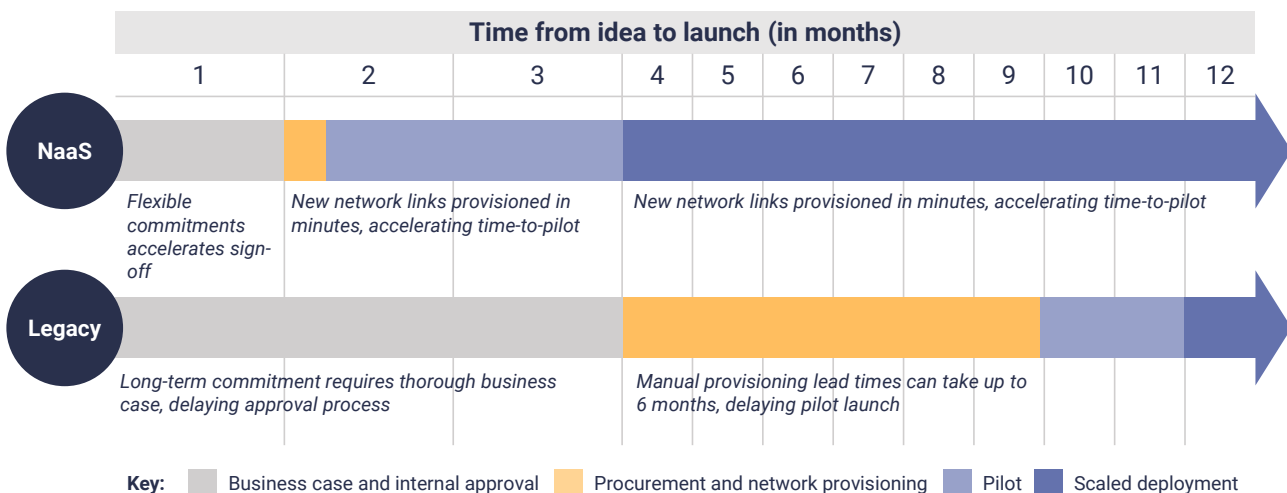
- High upfront cost of 12–24-month contracts deters experimentation.
- Long provisioning timelines delay go-live, creating internal friction and business case hesitancy.
- Inability to spin up secure, high-performance interconnects quickly means pilots may be shelved entirely.

**With NaaS:**

- Secure, high-throughput interconnects can be provisioned on demand for short-term trial periods (e.g. 4–8 weeks).
- Bandwidth can be scaled up or down flexibly, with no long-term commitment – making experimentation lower risk and more accessible.
- GenAI capabilities can be trialed, validated, and launched more quickly – accelerating access to productivity gains and customer-facing impact.

**Cost impact for enterprise:** Capture upside of earlier go-to-market for AI innovation, with reduced risk and experimentation cost. Avoid deferred productivity gains (e.g. in customer support) by accelerating time to launch.

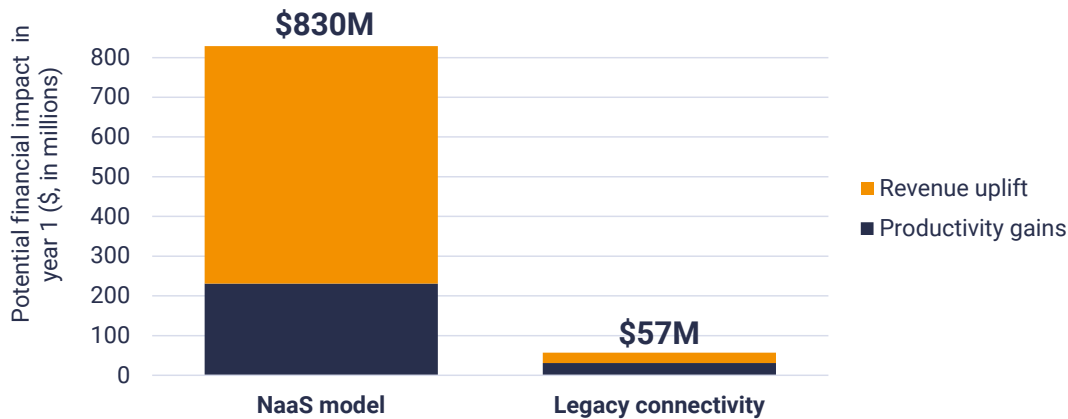
**Figure 4: Timeline from idea to launch for GenAI application**



Source: STL Partners

**Figure 6: Estimated financial impact on a large bank launching a GenAI customer care agent in first 12 months**

With NaaS, a large bank can capture an estimated \$830M in GenAI-driven gains in year 1, by accelerating pilot and scaled deployment of customer care agent by 8 months compared to legacy connectivity.



Source: STL Partners analysis

## Scenario 3: Building network resilience for distributed global operations

As enterprises expand globally and distribute their operations across multiple regional hubs (e.g. to access talent pools in other markets or to offshore certain teams), their reliance on seamless, resilient interconnectivity intensifies. Downtime in one region can now disrupt operations, service delivery, or financial workflows around the world. Legacy connectivity models offer limited observability into network performance limiting enterprises to mitigate faults swiftly or proactively.

**Enterprise situation:** Following industry trends, a global financial services firm operates a major Global Capability Centre (GCC) in India with over 50,000 employees supporting certain functions including product development, finance, and customer delivery. These employees rely heavily on continuous access to business-critical systems and operational data hosted in the firm's core data centres across the US, EMEA, and APAC. As such, resilient connectivity between the GCC and these international hubs is essential to maintain business continuity and avoid disruptions to daily operations.

**Enterprise challenge with legacy connectivity:**

- Traditional long-haul links provisioned via manual processes offer limited visibility.
- Without programmability or real-time observability, network issues may go undetected or unresolved, creating hours of unplanned outages across GCC operations.
- A recent fault between the GCC and core DC resulted in hours of lost productivity across thousands of employees.

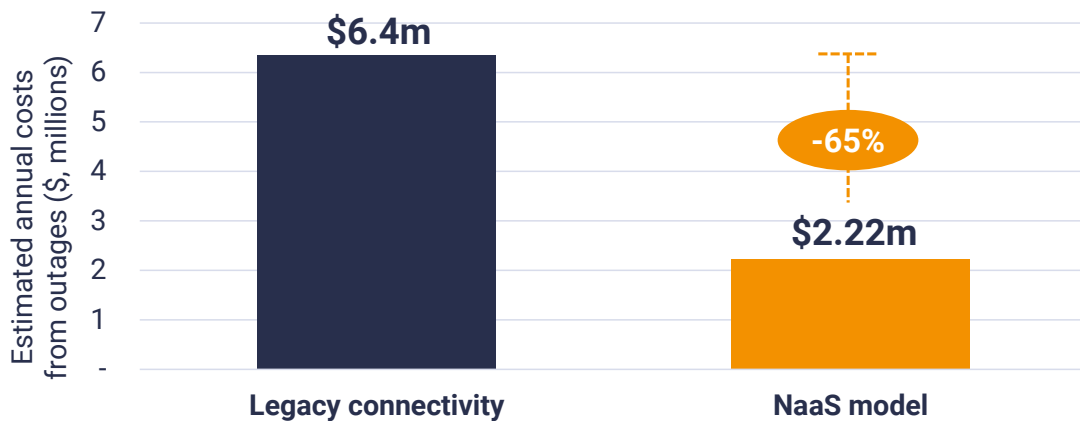
**With NaaS:**

- Enterprises can end-to-end visibility across their entire international inter-DC network enabling clearer view of network health and performance.
- Faults can be detected quickly, triggering automated responses or dynamic rerouting to reduce downtime through proactive fault detection and automated recovery.
- IT teams are empowered with real-time observability and self-service tools, reducing reliance on ticket-based workflows and enabling faster, more efficient resolution of issues.

**Cost impact for enterprise:** Reduce the frequency and duration of unplanned network outages. Minimise productivity loss from large-scale disruption to mission-critical global operations.

**Figure 7: Estimated annual cost savings in productivity loss from network outages**

For a large financial services firm with a major Global Capability Centre (GCC) in India operating 50,000 employees, NaaS could reduce annual outage-related productivity losses by 65%.



Source: STL Partners analysis

# Conclusion

## Recommendations for enterprises

- **Anticipate how AI will reshape Layer 1 network requirements.** As enterprises adopt AI at scale – from distributed training to real-time inference – they will face increasingly bursty and intensive network demands. These traffic patterns can be short-lived but bandwidth-heavy, making traditional fixed-capacity network models costly and inefficient. Enterprises should assess whether their Layer 1 infrastructure can support this new level of dynamism – particularly in scenarios where deterministic performance, security, and compliance are non-negotiable.
- **Embed commercial flexibility into ROI frameworks.** When evaluating the value of Layer 1 NaaS, enterprises should go beyond traditional cost comparisons and factor in benefits such as reduced time-to-market, improved resource utilisation, and the ability to flex capacity on demand. ROI calculations must account for avoided costs (e.g. overprovisioning, downtime) as well as opportunity gains (e.g. faster innovation cycles, earlier productivity capture).
- **Reframe connectivity as a lever for growth and resilience.** Enterprises need to shift from viewing network connectivity as a static utility to treating it as a dynamic enabler of digital performance, innovation, and resilience. In increasingly distributed and cloud-native environments, connectivity is not just an IT line item – it underpins how fast, securely, and efficiently the business can operate and scale.

## Recommendations for service providers

- **Position NaaS through enterprise outcomes, not just technical features.** To drive enterprise adoption of NaaS (Layer 1 and beyond), service providers must move beyond marketing technical specifications and instead position offerings around the business outcomes they enable – whether that's reduced downtime, accelerated AI deployments, or more efficient infrastructure scaling. This requires adapting sales messaging and solution framing to align with enterprise KPIs and strategic objectives.
- **Strengthen the financial narrative for NaaS through real-world enterprise case studies.** Many enterprise decision-makers remain sceptical of NaaS due to vague or unsubstantiated financial claims. Providers should invest in building and sharing evidence-based case studies that clearly demonstrate cost savings, operational improvements, and avoided risks – especially in enterprise segments where Layer 1 connectivity is most relevant.
- **Evolve infrastructure and enable NaaS deeper into the network stack.** True NaaS value cannot be unlocked through “NaaS-like” features at Layer 3 alone. Providers must develop the underlying capabilities – such as automation, telemetry, and self-service provisioning – to deliver NaaS programmability at the physical layer (Layer 0/1). While the technology may be complex, the business case for addressing this currently underserved enterprise segment is growing stronger.

# Appendix

## Enterprise NaaS cost benefit estimation

STL has estimated the financial impact of adopting Layer 1 NaaS through a set of illustrative enterprise scenarios. These scenarios have been developed with Lightstorm product teams, drawing on insights from real-world customer adoption. Pricing assumptions of Layer 1 NaaS services are based on industry benchmarks of a service provider based in India.

### Scenario 1: Scaling connectivity for dynamic network demand

Figure 3 presents an estimate of the annual L1 connectivity costs for a large global CDN provider within a large single national market (e.g. India). During high-traffic periods — such as live sports streaming or software rollouts — the provider must rapidly scale capacity across ~15 backbone links from a typical 10 Gbps baseline to up to 50 Gbps. In a legacy model, the provider would need to overprovision capacity year-round, resulting in significant underutilised spend. By contrast, a NaaS model could enable the CDN to scale capacity dynamically, only for the duration of high-traffic events. STL estimates that this could reduce costs by 53% - equating to around \$130,000 savings annually.

### Scenario 2: Accelerating innovation without long-term lock-in

Figure 6 estimates a large North American retail bank piloting and deploying a GenAI-powered customer care application. Calculations are based on the bank having total annual revenues of \$20Bn, with a customer care team of 20,000 employees (with an average salary of \$50,000). Industry benchmarks suggest that GenAI applications in financial services industries can deliver 30% productivity uplift for customer agent teams,<sup>3</sup> and a 7% broader revenue uplift from enhanced customer engagement.<sup>4</sup>

A legacy connectivity model could delay GenAI application rollout by 8 months, deferring the capture of these benefits (see timeline in Figure 4). Assuming benefits are phased over time, we estimate NaaS can deliver \$830M in productivity and revenue gains in the first year, compared to \$57M under a slower legacy rollout.

### Scenario 3: Building network resilience for distributed global operations

Figure 7 estimates the financial impact of network outages on a team of employees working offshore for a global financial services firm. The offshore team has 55,000 employees, 60% of whom are highly dependent on digital systems, data, and collaboration tools, with an average loaded cost of \$40,000 per employee per year.<sup>5</sup>

---

<sup>3</sup> McKinsey, [The economic potential of Generative AI: The next productivity frontier](#), June 2023.

<sup>4</sup> Nvidia, [State of AI in Financial Services: 2025 Trends](#), June 2025.

<sup>5</sup> Average loaded cost defined as total cost incurred by the employer including salary but other costs including, training, development, overheads, admin, benefits - estimated to be 1.6x salary. Average annual salary for GCC employee estimated to be \$25,000.

In a legacy model, the offshore team experiences 10 outages per year, each lasting approximately one hour. With NaaS, we estimate a 30% reduction in the number of outages and a 50% reduction in duration per outage, driven by improved observability, automated recovery, and real-time control. This totals 65% reduction in outage-related productivity losses – saving around \$4.2M annually.

# PARTNERS



Research



Consulting



Events