Near-Term 25G PON Use Cases

By Benoît Felten, Chief Research Officer

January 2022

Executive Summary

The standard for 25G PON is being developed through a multi-source agreement (MSA) and ecosystem that comprises over 40 partners from around the world. As service providers around the world are considering or implementing XGS-PON to allow them to deliver multi-gigabit service plans, the need for 25G might seem a long time in the future. This report examines the near-term use cases for 25G technology.

25G PON has a number of characteristics that make it a compelling solution to address some current needs of service providers:

- It is simultaneously compatible with existing GPON and XGS-PON solutions, which means that all three can co-exist in a single network implementation serving different end points.
- It does not require dedicated OLTs or different splitters from GPON or XGS-PON.
- It can deliver symmetrical throughput and (with the right chipsets) ultra-low latency.

In the mobile space, this makes 25G PON a relevant solution for both cell site densification (reusing an existing fiber access network for mobile transport to the new micro-cells) and mobile transport (disaggregated RAN). In time, the perspective of a single fixed access network to serve all needs (residential, B2B and mobile) becomes a possibility.

In the B2B space, 25G PON could allow service providers to recreate differentiation between residential and dedicated B2B offers, thus increasing ARPU potentials and margins. It will also enable Industry 4.0 solutions that often need high throughput, not only onsite, but in between sites or to the cloud. On campus-type sites, it makes the proposition of passive optical LAN even more attractive.

In the residential space, 25G PON could not only offer blazing speeds for consumers, but allow next-generation applications based on augmented reality, virtual reality or holographic communications to emerge and thrive.

But what does the business model look like for near-term implementation?

In the residential market, it's hard to make a case unless customers are willing to pay considerable premiums for super high speeds. The cost of 25G ONTs is simply too high right now for a business case to be viable, although costs are expected to drop significantly by 2025.

In the business market, on the other hand, targeting businesses tactically could allow for selected implementations of 25G PON, generating extra ARPU and attracting new customers in the B2B and wholesale space. Our high-level modeling shows that a limited implementation as early as 2022 could pay back the investment in about 12 months.

Because of the very high throughput that 25G PON allows, a scenario where a new entrant would deploy a 25G network from day one was also examined. This would not only give said new entrant an advantage in the market in delivering better services, it would allow said entrant to deploy with a leaner passive infrastructure where splitting could be done at 1:128 instead of 1:32, which is the norm currently. In urban contexts, this could mean significant CAPEX savings. However, the cost of the ONTs in the short term would likely undermine those benefits, which means this case may be viable in a few years, but is likely not now.

All in all, it's clear that 25G PON offers immediate opportunities in the business market. Since it can be deployed alongside GPON and XGS-PON on the same active access network, it's a move that service providers should seriously consider, with nothing to lose and potentially significant gains.

Table of Contents

I. Introduction

As many operators around the world are starting to replace their legacy GPON technology with 10G PON (predominantly XGS-PON), the time has come to think about what the next step might be. There are a number of reasons why this kind of prospective thinking matters at this junction:

- depending on the markets, there may already be needs for capacity beyond 10 Gbps for external and internal use cases;
- wireline and wireless are increasingly interdependent, and being ready for 5G deployment at scale is increasingly important;
- because technology leadership is a proven way to gain market share and maximize margins.

Giving service providers some choice in this matter is important, and that next step, depending on the needs and maturity of each operator, might be a 25G solution, a 50G solution, or both. In October 2019, 10 leading operators and communication equipment vendors[1] initially formed a multi-source agreement (MSA) to develop a 25G standard and eco-system. The 25GS-PON MSA now includes over 40 members from around the globe.

The purpose of this paper is not to describe the 25G solution in-depth, but rather examine near-term business cases using 25G. In May 2021, Belgian operator Proximus announced its adoption of 25G solutions as part of its national FTTH rollout, so the wider implementation of 25G PON solutions is just around the corner, certainly from 2022.

Based on qualitative and quantitative modeling analysis, two business cases will be examined:

- a multipurpose converged network approach relying on 25G for premium customers and mobile transport, in parallel with a primary pre-existing GPON or XGS-PON platform, to create more revenues, decrease OPEX, and accelerate ROI
- a full 25G greenfield deployment taking advantage of higher split ratios in the outside plant to lower CAPEX.

We will also examine some of the services that may drive capacity demand in the near- to mid-term, and therefore drive the need for more robust PON technologies sooner rather than later.

The models we propose will, by definition, be generic, but hopefully will allow service providers to see their relevance and inspire them to examine how 25G technology may impact their businesses for the better.

[1] AOI, Chorus, Chunghwa Telecom, Ciena, MACOM, MaxLinear, NBN Co., Nokia, Sumitomo Electric Industries, Ltd, and Tibit Communications

II. 25G on the approach

While other papers exist that describe the 25G solution designed by the 25GS-PON MSA in more depth and with more technical detail, in order to build business models on the basis of partial or full 25G implementation, we need to identify some of the salient characteristics that will make a difference to service providers' businesses.

- 25G PON can be implemented in two ways; either to deliver 25 Gbps downstream with a 10 Gbps upstream capacity, or a symmetric 25 Gbps throughput.
- The 25G PON chipset is designed to support lower latency features, such as Co-operative Transport Interface and multiple burst in a frame. These features are being standardized and will work in combination with the technology's high capacity to deliver mission-critical services.
- 25G PON is simultaneously compatible with existing GPON and XGS-PON solutions, which means that all three can co-exist in a single network implementation serving different end points. This attribute will facilitate the introduction of 25G PON, without forcing operators to decommission GPON or skip XGS-PON. In addition, 25G PON supports co-existence with future 50G PON.
- 25G PON does not require dedicated OLTs. There are solutions already in the market where 25G is available on the existing GPON and XGS-PON linecards.
- 25G PON does not require different splitters from GPON or XGS-PON, so it can be implemented with no outside plant change.
- The first 25G PON ONTs will address premium enterprise and mobile transport applications. It is expected that, within a few years, the cost of the ONTs will decrease, enabling 25G PON to be used for broader markets (e.g., small businesses, residential).

Figure 1. GPON/XGS-PON /25G PON frequencies

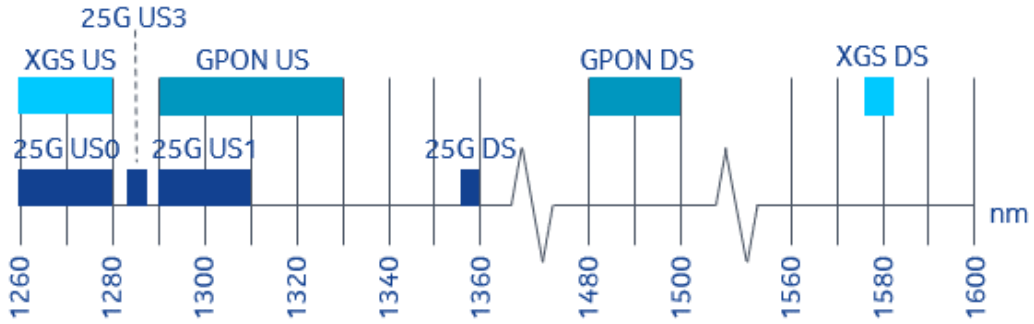

The above characteristics have a number of implications:

1. The high throughput and low latency characteristics of 25G PON make it a viable solution not just for mobile backhaul, but for fronthaul as well, which requires much higher throughputs. This also makes 25G PON a viable solution for distributed RAN implementations.

2. Whereas overhead prevents XGS-PON from delivering a true 10 Gbps end customer connection, 25G PON can, and it can even deliver a true 20 Gbps end customer connection.

3. The very high throughput offered by 25G PON can enable larger split ratios in greenfield deployments of FTTH than what is currently common practice. While networks deployed currently favor split ratios of 1:32 or 1:64, a 25G greenfield deployment could easily envisage a 1:128 ratio to more efficiently deliver massive gigabit services.

Obviously, these features could benefit service providers immediately, which is why the business models we will examine in part III make full use of them.

III. Practical Use Cases for 25G

1. 25G for multipurpose network convergence

For the sake of simplicity, we make a number of assumptions in order to theorize a model that will be applicable for most networks. Since most markets are fast moving towards a 1 Gbps service baseline, we assume that gigabit broadband is no longer a premium product. We also assume that most operators are currently implementing, or have plans to soon begin, network-wide XGS-PON migration.

Although 25G will be most attractive for mature markets, it could also be relevant to players in less advanced markets that have pushed hard for FTTH deployment. The model is particularly relevant to operators serving both the fixed and the mobile markets, but as will be addressed below, it can also be very significant for wholesale fiber network operators.

This approach addresses the following issues for service providers:

Mobile and 5G:

- **Disaggregated RAN enabler**. Most of the operators deploying 5G today consider splitting RAN functionality into physical network functions and virtual network functions. The virtualized functions can be centralized, so the network intelligence is shared across multiple cell sites to minimize CAPEX and OPEX at the cell site. There are 2 options for the split: higher layer and lower layer split, with different functionality for real-time processing. In higher layer split deployments, XGS-PON and 25G PON are used today. In lower layer split deployments, the signal is no longer processed at the cell site, so the unprocessed signal needs to be transported, and unprocessed signals can require capacity up to 20 Gbps. With current generations of PON, that's not an option, and connecting every cell site with P2P fiber would be very expensive. 25G PON, however, would be able to carry this unprocessed signal, thus enabling disaggregated RAN. Lower layer split will require low latency, requiring low latency features on a PON described above.

- **Cell site densification enabler.** While 5G deployment is already underway, most of it is happening at the macro-cell levels. A significant proportion of macro-cell sites are already fibered up (especially in developed markets), but service providers are still struggling to find a workable business case for 5G cell densification, in part because of the required fiber density. With a single 25G network covering residential, business and mobile, high capacity fiber backhaul becomes easily available for micro-cells closer to homes, where the FTTH network is deployed. 25G can become a key component to a viable cell densification 5G strategy, especially since the fiber capacity may already be there. From an investment

point of view, the cost of deployment would therefore be minimal. The unlocked revenue, customer satisfaction, and new business models would more than offset that small investment.

A more general consideration about long-term network planning is that the physical convergence of fixed and mobile networks, while not a reality today, is fast becoming a possibility. Currently, most mobile operators have a dedicated backhaul network for macro-cells, a combination of point-to-point fiber and microwave. When they have deployed micro-cells, it's mostly using microwave, and only in the best of cases is fiber used for backhaul. Those who are also serving the fixed market may have deployed FTTH, but very rarely are these networks merged. There is a dedicated fiber network for mobile, and in parallel, a residential FTTH network. In the best of cases, the residential FTTH network also serves the business market, but in some cases the high end of the business market is served by point-to-point fiber, in reality a third network. With 25G, these three networks could easily be merged into one for significant OPEX savings, as well as CAPEX going forward. Several techniques are available to operators to facilitate the network convergence, such as SDAN-based network slicing to segregate the consumer and mobile traffic, or separation of traffic per wavelength, for example GPON and XGS-PON wavelengths can be used for consumer traffic, and 25G PON wavelengths for mobile transport.

Enterprise:

Broadband for the enterprise sector is a key profitability vector for service providers; it's higher margin, associated with higher performance, and often linked to additional product upsells like LAN, mobile, and some vertical services. 25G can be a key tool to create value for enterprise customers.

- **Access differentiation:** Ever since fiber-to-the-home started connecting residential customers, service providers have had a hard time differentiating their enterprise offerings from cheaper residential offerings. In large business contracts, the main driver for price may be factors such as service level agreements, access redundancy, long distance and WAN services, etc. rather than raw capacity, whereas with small and medium businesses, it may very well be the bandwidth. If any home can get 1 Gbps or more for less that 100 EUR, it's a lot harder to convince a small- or medium-sized business to pay several times more for the same speed, even if the services are different in terms of service levels and other features. With 25G, a service provider could allocate certain speed tiers (like true 10 Gbps or even true 20 Gbps) specifically to enterprises and thus recreate a degree of capacity differentiation vs. residential that would translate into additional revenues. This would increase business customer satisfaction and lower churn while increasing (or at least not decreasing) ARPU at the same time.
- **Industry 4.0 enabler:** Industry 4.0 presents exciting opportunities linked to high-speed and low-latency connectivity delivered by 5G mobile solutions, but also through high-capacity fiber to the premises. Not only does this trend promise to redefine the manner in which many verticals (e.g., manufacturing, logistics, etc.) operate, but it also opens possibilities for entirely new functions; for example, advances in augmented and virtual reality combined with super-high-resolution video capability open up a new array of applications for design, inspection, testing, etc. Digital twin developments will allow real-time analyses in the cloud. All of these emerging functions will require huge throughputs while maintaining ultra-low latency.
- **Campus boosting:** Passive optical LAN has been making strides in the last few years, connecting facilities like stadiums, hotels, and university campuses with hundreds of access points or desks. However, PON being a shared capacity technology by definition, the costs of connecting large campuses can still be significant, which makes the selling of such solutions a little more difficult. With 25G PON, passive optical

LAN could be deployed much more economically and effectively while still delivering bandwidths above 1 Gbps to each end point.

Residential:

- **Trailblazing speeds:** While the question of why consumers would need speeds beyond 1 Gbps regularly comes up in the public discourse (as it did for 20Mbs, 100Mbps and 1 Gbps in the past 15 years), the fact is that speed sells. Some consumers are willing to pay a premium to be connected with speeds that seemed impossible just a short while ago. In this respect, delivering residential speeds of 5 Gbps or 10 Gbps may be a worthwhile strategy. This may not be a large portion of the customer base (certainly less than 10% initially), and it may not be a huge premium (probably in the 10-20% range over 1 Gbps prices), but this adds up. Additionally, it has been demonstrated by the early adopters of Gbps speeds that offering a gigabit service attracted new customers at lower speed tiers as the service provider demonstrates its superior technology. It is likely that the same magnet effect would apply to 5 Gbps or 10 Gbps speeds.

- **Next-generation applications:** With the emergence of augmented reality and virtual reality, low latency and high capacity in real time will be more crucial than ever. While the current generation of applications may still function with gigabit speeds, at the pace of innovation, it won't be long before some applications start lagging. While these will take time to become widely adopted, they will revolutionize many activities in the home, from gaming to virtual cultural outings to DIY. Similarly, while the recent pandemic lockdowns have demonstrated the importance of video communications, they have also shown the limitations of the existing solutions when it comes to image quality, sound quality, interactions, lag, document sharing, etc. While few of these applications may individually require speeds above 1 Gbps today, early adopters will soon find that higher speeds and lower latencies improve their experience considerably. See below for a representation of the future application space.

High-Capacity Services and Application Map

Asserting which emerging high-capacity services and applications will be successful is a difficult exercise. However, it is possible to identify some general trends and map the services likely to drive capacity demand. In this context, capacity is not only bandwidth, but also ultra-low latency as many of these services can only function well if near instantaneous. The following figure maps out some of the likely drivers in the next five years.

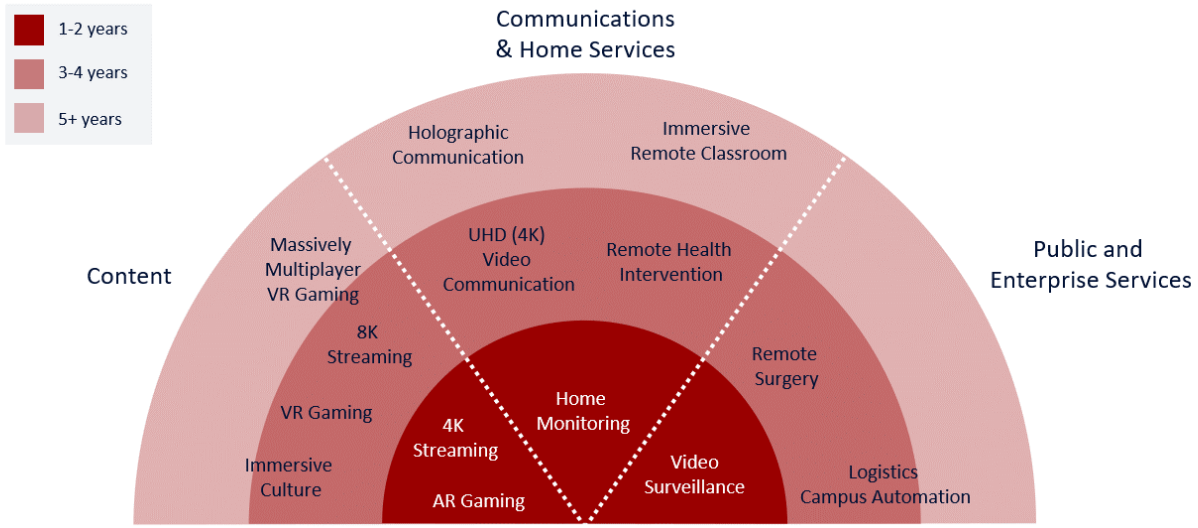

Source: Diffraction Analysis

Building the business case

This 25G implementation scenario is an incremental approach. It is also multi-faceted with residential, enterprise and mobile components. Building a generic business model on that basis would likely not be relevant to any service provider because the specifics would be too different in each case. Instead, for each area where 25G could enable new revenues (or significant cost savings), we will examine what the associated costs could be and how the revenues could be valued.

As a general consideration, it should be noted that, 25G being fully compatible with a parallel implementation of GPON and XGS-PON, the only thing that requires upgrading to enable 25G is the optics in the line card in the OLT and the ONT on the end user side. However, since a port in a standard point-to-multipoint fiber deployment generally covers 32 or 64 customers, upgrading or adding the 25G optics for a single customer can be a significant cost. One strategy is to dedicate a port to 25G customers and map those customers that want to be upgraded to that port.

In the following model, we examine the simple economics of extra CAPEX investment over GPON to upgrade to 25G PON and the extra revenue needed to offset that investment. Our view is that such an investment should be tactical, i.e., only done (at least early on) in select areas where the revenue potential is clearly identified. We are looking exclusively at broadband offers here, with no speculation on extra revenues that could be unlocked by offering value added services (e.g., VR, content, Netflix, etc.). Our ARPU and margin hypotheses are as follows:

- GPON generates 50 EUR/month of revenue for a 1 Gbps residential offer; 25G PON would allow for a true 10 Gbps premium residential offer sold at 75 EUR/month

- GPON generates 250 EUR/month of revenue for a 1 Gbps B2B offer; 25G PON would allow for a 20 Gbps B2B offer (or wholesale offer) sold at 400 EUR/month.
- We consider a 65% EBITDA margin on the fiber broadband, which is somewhat standard in the industry.

These are relatively conservative prices that should be attractive while maintaining the necessary differentiation between residential and B2B offers.

We outline 3 scenarios based on the relative success of each new offering. Here are our assumptions:

- A single architecture for all offerings and a 1:32 split level with 16 ports per card (i.e., 512 addressable customers).
- Costs are 2022 and 2025 projections sourced from a recent Omdia study[2], and payback times are calculated on the basis of differential revenue (i.e., only the extra revenue from 10G residential and 20G B2B/Wholesale offerings are taken into account).
- A starting take-up rate of 50%, 45% residential and 5% B2B/Wholesale.
- Our model does not account for the savings in shifting B2B customers from a dedicated P2P network over to a single 25G network serving all markets.

If we assume that a sound investment in access upgrade pays back in less than 12 months and that nothing above 24 months is worth considering, we conclude the following from this high-level modeling exercise:

- The business case to serve residential customers is limited for the time being. Based on 2022 prices, payback for a modest 5% adoption of residential 10 Gbps (new customer acquisition) is at 26 months, and only drops to 20 months with 2025 OLT/ONT costs. Even at an ARPU of 100 EUR/month for 10 Gbps residential, the payback time with 2025 costs is 15 months.
- On the B2B/Wholesale front, things are more promising. Here, limited adoption (2% migration to 20 Gbps) pays back in 15 months even though the cost of the OLT is absorbed by a small number of customers. A more substantial adoption (3% migration and 2% new customers) brings the payback at only 8 months based on 2022 costs. More B2B/Wholesale customers means the OLTs are paid back faster, and the revenue upswing substantial. The same scenario with 2025 prices brings the payback at a very attractive 6 months.

This shows that in the short term, courting the residential market with super-high-speed 10 Gbps offers is hard to pull off unless residentials are willing to pay a massive premium for such high speeds. Focusing on B2B/Wholesale for the near future seems like the best approach to make 25G work. This in turn means that if differentiation between residential and B2B on speeds is no longer necessary, B2B/Wholesale speeds could be delivered at the same price with 10 Gbps offers, keeping 20 Gbps for more premium uses or an upgrade path down the line. Residentials could be addressed with cheaper XGS-PON solutions with 5 Gbps still being a substantial improvement on existing 1 Gbps offers.

[2] Fiber and Copper Access Equipment Forecast Midyear Update 1H21 (July 2021)

Figure 2. 25G PON payback time

	Residential	B2B / Wholesale	
	Scenario 1	Scenario 2	Scenario 3
	+5% take rate from 10Gbps adoption	2% migration from 1Gbps to new B2B/Wholesale 20Gbps	3% migration + 2% adoption of B2B/Wholesale 20Gbps
2022 costs	Payback: 26 months	Payback: 15 months	Payback: 8 months
2025 costs	Payback: 20 months	Payback: 11 months	Payback: 6 months

As this high-level analysis shows – independent of extra revenue unlocked (or costs saved) from deploying and delivering 25G as highlighted in some of the use cases above – a tactical deployment of 25G to address B2B and wholesale (essentially mobile transport) demand where it exists is viable today. This in turn could unlock other use cases as outlined above when it comes to mobile network deployment and management or further B2B revenue development.

2. 25G greenfield deployment

Quite often in the realm of technologies, a late comer to the market can acquire a distinctive advantage by foregoing early generations of a given technology and adopting from the get-go a forward-looking, recent iteration of said technology. Indeed, in the realm of fiber-to-the-home deployments, we are seeing this already with some new market players deploying XGS-PON natively on their network instead of GPON. While slightly more expensive, this approach lets them deploy a platform that they know will be future proof for at least a few years, but also offer better service packages, and therefore gain traction on the market faster.

Is this approach equally valid with 25G PON today? Probably not in 2021, but it could become so soon.

As is often the case with technological leapfrogging, building a network with 25G would allow a service provider to offer speeds from day one that established competitors couldn't match. If, for example, a service provider offered 10 Gbps broadband as a baseline residential offer comparable to existing 1 Gbps offers, not only would that provider get buzz and probably quick take-up, but competitors would struggle to match that since it would imply deploying 25G themselves. (XGS-PON cannot deliver 10 Gbps because of overhead; it's limited to around 8 Gbps). That's a good position to be in when entering a new market. Furthermore, by deploying a single unified platform, a newcomer would operate a simpler, more streamlined system from the start while established operators would have to manage legacy systems, upgrades, etc.

But the biggest advantage that 25G could deliver if deployed as a brand new network actually has significant CAPEX savings impact: with 25G, a new entrant could decide to split his point-to-multipoint at much higher ratios than is currently the norm, for example 1:128 splits. With such split ratio, said new entrant would save costs on the outside plant per user connected. The main caveat is density: such high split ratios lower the available optical budget, so the distance from the OLT to the user will have to be short. In a typical south-east Asian city where the average residential high rise has hundreds of apartments, this kind of approach would obviously make sense. In less dense European cities, the parts of the urban landscape where the savings would stay significant might be more limited. In suburban or rural areas, the savings might be small or even non-existent simply because the building density is too low to take full advantage of these higher split ratios.

The biggest downside to such an approach is the cost of ONTs. Current 2022 estimates put the 25G ONT at 20 times the cost of a GPON ONT, which makes them suitable for high revenue B2B offers, but not for residential offerings. There simply isn't enough ARPU to go around to pay for it, even factoring in CAPEX savings from higher split ratios and higher market shares. However, the cost of 25G ONTs will be dropping fast with a 25% reduction in cost expected by 2025. By that point, the economic case might start to look relevant, at least in dense urban areas.

The prevalent case today remains that for a newcomer entering the market, the most sensible approach would be to deploy XGS-PON natively and selectively rely on 25G for wholesale, large enterprise and some SMBs. Within a couple of years however, a detailed cost analysis will need to be done to see where the tipping point is for a full blown 25G deployment to be undertaken by newcomers.

IV. Conclusions

25G technology is ready to be deployed today.

The economic case for tactical B2B/wholesale offerings relying on 25G seems to be sound. It would deliver a number of distinct advantages:

- It would allow renewed differentiation between residential and SMB offers with significantly higher speeds delivered to business customers;
- It would provide headspace for speed segmentation in B2B offerings that XGS-PON solutions can't match;
- It would deliver wholesale offers compatible with current and future mobile network transport solutions.

Because of the way 25GPON coexists with existing PON and XGS-SPON linecards, it can be implemented not only in areas of high density of business customers (e.g., financial districts, business parks, etc.), but even in more general areas where a few business customers could be enough to sustain a viable 25GPON implementation on a single port.

The residential economic case is harder to make in the near term because of the high cost of the ONTs. As these costs go down in the next few years, however, the case for a premium 25G-based residential offer will emerge, coincidental with market demand for such a service. Similarly, as ONT costs trend down, the economics around a new entrant launching a native 25G network deployment will likely be viable, especially in dense urban areas, given the high split ratio CAPEX savings.

The fact remains that 25G is not a 2025 or beyond technology, it's a 2022 technology that can immediately offer opportunities and pay dividends for service providers willing to stay ahead of the curve in delivering the best and most advanced services to their business customers today.