



Four misconceptions limiting the QoD API opportunity

Quality on Demand (QoD) could be a US\$3.4 billion opportunity by 2030, but telcos still risk misreading how it should be packaged, priced and taken to market. This article explores four common misconceptions around QoD and why operators need to package it around real customer outcomes, not just API access.

Krsna Singh, Research Analyst

Quality on Demand (QoD) is emerging as one of the most commercially promising network APIs, with a projected **US\$3.4 billion** opportunity by 2030. The pitch is simple: when an application needs better connectivity for a specific moment, it can request a temporary boost in network performance.

That could mean lower latency for a cloud gaming session, more stable bandwidth for a live video upload, or better reliability for a business app in a congested location. QoD sits as an API layer on top of the public network, using traffic prioritisation and orchestration to improve performance for specific sessions or traffic flows.

But while QoD is gaining momentum, telcos still risk misunderstanding how it should be positioned, packaged and sold. This article explores four common misconceptions we see coming up in the industry and what telcos should think instead.

Misconception 1: QoD is just a lighter version of network slicing

QoD and network slicing are often discussed together because both involve differentiated network performance. But they are not the same product. Network slicing creates a logically partitioned network environment with defined performance, security and resource characteristics. It is best suited to longer-term, highly controlled use cases where strict service levels, isolation or predictable resource allocation are required.

By contrast, QoD is more dynamic: it allows operators to apply a temporary quality profile to a specific application, session or task within the shared public network, rather than creating a dedicated slice. Because QoD operates within the shared public network, it offers flexibility and scalability, but its availability and performance still depend on network conditions and available capacity. This makes it best suited to short-term assurance use cases where performance needs to be improved for a limited period.

A factory might use network slicing for a permanent automation environment. By contrast, QoD might be used by a broadcaster for a 30-minute live upload from a congested event, or by a gaming platform to stabilise latency during a multiplayer session. Although these use cases may appear similar, their underlying requirements and resulting commercial models can differ significantly. This is a distinction we explore in more detail in our article on [Quality on Demand and network slicing APIs](#). For telcos, this means positioning QoD not as “slicing lite”, but as a flexible, on-demand tool for moments of need.

Misconception 2: QoD is only for ultra-premium, mission-critical use cases

Many early QoD trials focus on high-stakes scenarios such as drones, emergency services, industrial automation and autonomous systems, where connectivity failure can cause an immediate halt to operations or create safety risks. These use cases are important because they show what QoD can enable in demanding environments where network performance is critical. However, this does not mean that QoD is only relevant for ultra-premium, highly specialised or safety-critical scenarios.

The bigger commercial opportunity may lie in more repeatable, higher-volume use cases that are not necessarily “mission-critical”. QoD can be valuable wherever a short dip in performance creates a tangible problem, suggesting that it can address a broader set of enterprise challenges rather than being restricted to ultra-premium “hero” use cases. To open up this wider market, the key question telcos should ask may not be “is this mission critical?”, but “does poor performance create measurable cost, delay, risk or frustration?”

This opens up a much wider set of performance-sensitive and business-critical applications, where poor performance may not create safety risks but can still have a direct operational or financial impact. In media production and photojournalism, QoD could support more reliable uplink performance at crowded live events, helping teams transmit video faster and avoid overprovisioning. In retail and hospitality, it could support mobile point-of-sale terminals at packed stadiums, festivals or transport hubs, where failed payments directly hit revenue.

QoD could also support everyday productivity use cases. In enterprise IT, for example, it could help employees maintain stable video calls from congested locations such as airports, hotels or conference venues. This is where products such as *AT&T Turbo for Business* are interesting. While it is not clear whether Turbo for Business is enabled by network APIs today, it shows how prioritised connectivity can be packaged as a practical business service. AT&T positions this product as a way to prioritise business data traffic, helping users stay productive during periods of high network traffic. The significance of this example is that it is practical and enterprise-led, rather than a futuristic or experimental use case.

For telcos, the lesson is clear. QoD should not only be sold around the most exciting demo. The larger market may be in repeatable, everyday pain points: better video meetings, smoother field work, more reliable mobile point-of-sale, and fewer failed digital interactions.

Misconception 3: One QoD API means one commercial model

One of the biggest risks for telcos is treating QoD as a flat, one-size-fits-all price per API capability. Some use cases may be suited to per-call pricing, while others may fit time-based, session-based or event-based charging. There are three reasons why telcos need a more use-case led, nuanced commercial approach:

1. Different customer segments may be used to different pricing models. A business user trying to maintain stable connectivity at a congested airport may not think in terms of individual API calls; a day pass or monthly productivity add-on may be easier for the customer to understand and buy.
2. QoD's value varies depending on the use case, location, congestion level and performance requirement. A live broadcast upload from a packed stadium and a cloud gaming session at home may both need better performance, but they do not create the same value for the customer. A single flat price risks under-monetising high-value workflows while making lower-value use cases too expensive to adopt. Operators, therefore, need to link pricing to the outcome being delivered.
3. Finally, network resources are finite. If too many users can buy priority at the same time for the same low price, priority stops being meaningful. Pricing needs to help manage demand and protect network performance, not just monetise access. This means considering both the right price point, and the right dynamic pricing mechanisms, such as higher charges during peak usage periods or in congested locations.

Misconception 4: QoD will sell itself as a standalone API product

Telcos should not assume customers will buy QoD simply as a standalone technical feature. For many customers, "QoD API calls" will sound too technical and abstract. They want smoother live streams, fewer dropped meetings, more reliable payments, better customer experiences and less operational disruption.

This means telcos need to package QoD around outcomes, not API mechanics. A horizontal API catalogue is important because it creates consistency and makes QoD easier for developers and partners to access across operators. But a catalogue of capabilities is not the same thing as a sellable proposition. Enterprises are

unlikely to pay simply for API access; they will pay for guaranteed performance, assured outcomes and avoided losses.

Operators should therefore create use case-specific offers with vertical specific packaging: such as a cloud gaming profile, live broadcast profile or enterprise productivity profile, each with a clear value proposition for that customer segment. As discussed earlier, these offers may rely on the same underlying technology, but they should not necessarily be priced in the same way. This is how telcos can protect long-term value capture. If QoD is sold only as a raw horizontal API, it risks being treated as a commodity input. If it is packaged around clear customer outcomes, telcos have a stronger basis for premium pricing, customer stickiness and differentiation.

As one SVP of innovation and industry strategy at a global network intelligence platform provider told STL Partners: "From our market research, there is significant demand with a major caveat: it must be simple and easy for developers to adopt... and there also needs to be clear verifiable proof of how QoD improves app performance above and beyond best-efforts networking."

That proof will be critical. Telcos cannot simply claim that QoD improves performance. They need to show it in terms customers care about: fewer failed uploads, lower lag, fewer dropped sessions, faster task completion or higher transaction success rates. Doing this well will require more than publishing an API. Telcos will need to co-create with developers, application providers and enterprises: testing QoD in real workflows, experimenting with different performance profiles and building a clear ROI case.

In conclusion: don't sell the API product, sell the customer outcome

QoD is one of the clearest programmable network propositions because it links network performance directly to moments where quality matters. But telcos cannot assume the market will form on its own. As we discuss in *Network API pilots to profit: A commercial playbook for Quality on Demand*, operators will need to actively create demand by helping customers understand where QoD fits, what it improves and how its value can be measured.

That will require new sales and go-to-market capabilities. Telcos will need to engage earlier in customer workflows, support collaborative pilots and sell outcome-led propositions rather than technical API access. They should also recognise that go-to-market models need to be shaped around both the API and the use case. The channels and ecosystem partners that work for identity and anti-fraud APIs may not be relevant for QoD. Instead, operators should identify the established players, workflows and buying centres in each target sector, and build from there. Network API ecosystem and enablement players such as Nokia and Shabodi are already helping operators work more closely with customers and ecosystem partners to shape QoD requirements by sector, reinforcing the need for use-case-led, cross-ecosystem coordination.

This customer and partner led approach also means that telcos should not only think about QoD as a standalone revenue engine. In some cases, its bigger role may be to strengthen other parts of the operator's business, from enterprise mobility to private networks and industry-specific solutions. The opportunity is therefore not just whether QoD generates direct, net-new API revenue, but whether it makes wider propositions more valuable, differentiated and resilient.

Ultimately, QoD will succeed if telcos can package it around the moments it protects: the dropped call avoided, the live feed delivered, the payment completed and the customer experience saved.

Four misconceptions limiting the QoD API opportunity

Krsna Singh is a Research Analyst at STL Partners, where he focuses on key trends across the telecoms industry, including edge computing, private networks, enterprise services such as network APIs, and sustainability.

Get in touch with the author to learn more

krsna.singh@stlpartners.com

Or visit STL Partners' NaaS and APIs Hub

<https://stlpartners.com/naas-network-apis/>