Ericsson White Paper
EAB-25:001373 Uen
April 2025

ERICSSON

# Differentiated connectivity: Unleashing the full potential of 5G

# Contents

# Introduction

Mobile network performance has improved significantly over the past decades, creating the foundation for our digital society, which paved the way for the smartphone and application ecosystem. However, mobile connectivity is, in essence, still a best-effort mobile broadband (MBB) service that limits innovation in applications that could greatly benefit from consistent, reliable connection characteristics to deliver even higher value to users. For mobile connectivity to match the evolutionary potential enabled by technologies such as AI and cloud, the mobile industry needs to transform from the current linear business of providing a best-effort data pipeline, to offering an innovation platform that supports applications in fully utilizing connectivity with consistent performance characteristics. The key to achieving this industry transformation is differentiated connectivity.

In this paper, we first describe the current situation in the mobile industry, including challenges faced by application developers and communications service providers (CSPs). Then, we describe how we can achieve a virtuous circle of increased network capabilities, value-generating applications and resulting revenues. We also introduce performance levels as conceptual vehicles for the industry to move forward. For readers interested in more detail, we provide a seven-step guide on how a new service offering based on performance levels can be introduced in a mobile network. We conclude by highlighting the actions and alignment efforts needed for the industry to capture this opportunity.

# Current industry trends, outlooks and challenges for communications service providers

The introduction of MBB revolutionized communications and paved the way for the smartphone [1]. As new applications and rapid device development provided a steady growth in mobile data traffic, the CSP business model adapted from voice and SMS to mobile data subscriptions. Additional spectrum and 4G technology made it possible for CSPs to improve network capacity, stay ahead of the traffic curve and capitalize on the increasing data demand. With 5G, CSPs could double their spectrum asset thanks to the large 5G-deployed time-division duplex (TDD) mid-band carrier, and deliver significant increases in network coverage, capacity and user throughput with Massive MIMO [2].

During this decade, there has been a change of dynamics in many markets – network capacity is no longer the bottleneck it used to be. Data buckets have become unlimited or large enough to rarely be exceeded, smartphone development is incremental, and innovation in new applications has slowed down.

MBB has matured and the revenue growth previously experienced has slowed down. There are three main causes of these developments:

1. With one main product (MBB) with limited differentiation, consumers find it increasingly difficult to distinguish between brands, leading to greater competition on price.
2. Since MBB capacity is less of a concern, there is a near-zero marginal cost of adding new subscribers, which further increases the competition on price.
3. With unlimited or oversized data buckets, revenue growth no longer follows mobile traffic growth. With limited revenue growth and increasing operation costs, there is a risk that CSPs are incentivized to decrease service quality by underinvesting in their networks.

On the other hand, the acceleration of major cross-industry trends such as digitalization, electrification, and automation increase society's dependence on high-performance mobile connectivity. For example, network capabilities such as Quality of Service (QoS), consistent throughput and bounded latency are required for real-time cloud interactions, from gaming to real-time digital twinning and extended reality (XR) [3].

MBB will not be sufficient to cater for these demands, or to provide a platform for new innovations. Given the best-effort nature of mobile connectivity, application developers cannot predict the quality of the connection that their globally distributed users will receive. Since they cannot know in which cell a user will be, or predict or trust the connection within that cell, developers currently create apps assuming worst-case scenario throughputs and latencies.

A new approach is needed to break the current downward spiral and raise the value creation capability of cellular connectivity.

# Creating a virtuous circle through differentiated connectivity service offerings

Network traffic consists of an abundance of different data streams, all with unique characteristics and needs. By differentiating the network's performance and making it deterministic* with regards to a set of performance characteristics such as latency and throughput, it can be offered as specialized connectivity services with reliable performance, to meet the distinct and varied needs of each application and even each application data stream. Differentiated connectivity covers the full spectrum, including application behavior in client and server, as well as networking technologies such as radio scheduler algorithms, software-defined networking (SDN), diverse 5G Core network capabilities (including UE route selection policy (URSP) and APIs exposure), operations support systems (OSS)/business support systems (BSS) capabilities (such as orchestration and assurance), and network slicing. These technologies facilitate dynamic, flexible network configuration and management to drive a differentiated user experience.

* With deterministic, the intended meaning is a connection that fulfills a set of characteristics targets with high reliability.

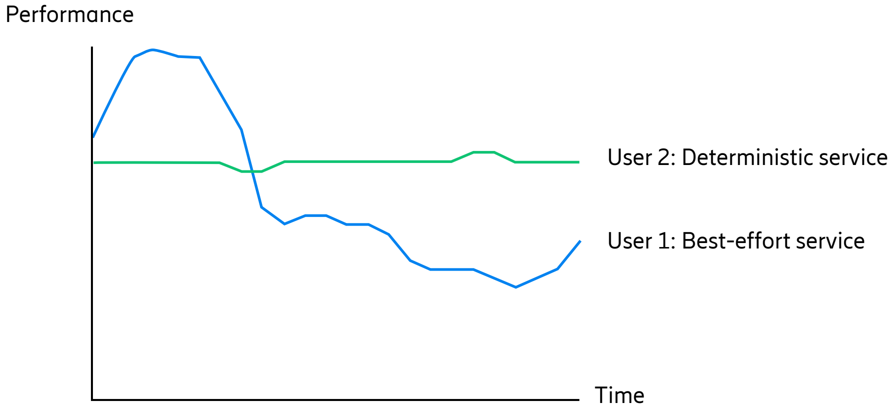


Figure 1: Comparison between a best-effort service and a deterministic service in a differentiated network

For developers, differentiated connectivity is not only of value because of a low latency or high throughput, but also because it is deterministic and reliable. When developers can trust that the performance, such as the throughput and latency, will not drop below a specified threshold, then such connectivity services have a monetizable value. An application can then be designed with specific connectivity requirements that it is known can be fulfilled. This is illustrated in Figure 1. To ensure the required level of trust, differentiated connectivity services will be provided with the likes of: monitoring tools, monthly performance reports, and Service Level Agreements (SLAs) to business customers, guaranteeing the high availability that they are used to with other services that their applications are dependent on.

Differentiated connectivity will address the full potential of digitalization, spurring innovation and competition and making performance an important differentiator in future networks. Even today, CSPs are taking steps toward this future with premium QoS plans at a price premium, including: AT&T US (turbo plan [4]), Singtel (5G express pass [5]), Telstra (speed and latency optimizer [6]), AIS Thailand (5G boosts [7]) and 3 Hong Kong (5G Live Connect plan for live streaming [8]).

Enhancing app experience will bring app developers a wider reach, lower churn, premium payments, and the possibility of creating new applications that cannot work in the best-effort paradigm. Instead of needing to develop for the lowest user connectivity performance scenario, developers will be empowered to develop for a defined and guaranteed performance across cells in a network and across networks globally.
In this new paradigm, performance will be the differentiator for user experience, and coverage of high-performance network capabilities will be a competitive advantage bringing new revenues to CSPs and a return on further network investments. In this market dynamic, CSPs who invest in network quality will be rewarded by end users and application developers, creating a virtuous circle of network investments.

**A broad range of applications benefit from deterministic differentiated connectivity**
Differentiated connectivity will bring value to both consumers and enterprises, in different user segments with different device types. We have defined three device/user segments to explain how differentiated connectivity can be implemented.

- **Platform devices:** These are consumer smartphones that feature applications distributed via an app store. An application such as outdoor mobile AR gaming would greatly benefit from differentiated connectivity to ensure the stable low latency and high throughput required for the real-time integration of high-resolution digital elements in the physical environment, assuring game developers that their users will have a great experience without delays or bad resolution.
- **Enterprise-managed devices**: These are typically smartphones, or, for example, a 5G router that has applications downloaded and managed through a company's application store and IT department. Business communications applications already generate multiple data streams with different needs today, and with current hybrid working environments, the need for stable and guaranteed connections is critical for productivity and employee satisfaction.
- **Full-stack systems**: These include devices with built-in connectivity, that feature apps controlled and installed by the manufacturer/integrator of the device or system. This could be, for example, a car or a TV-camera for broadcasting, requiring a trusted uplink capacity to ensure that they can reach their viewers with live coverage.

**Bringing differentiated connectivity to the market**
Differentiated connectivity can be sold either directly from CSPs to consumers or enterprise customers, or indirectly through integration with an application. In the indirect model, revenue would be realized from an end user perspective either as in-app purchases, or as part of a monthly app subscription fee. For platform devices, the most common model will most likely be integration with the application layer as this enables the required scale. Toward enterprises with enterprise-managed devices, there will likely be both models where enterprises buy their IT applications integrated with differentiated connectivity (likely for off-the-shelf applications), or in a direct CSP sales model (likely for in-house applications). Full-stack integrators will also use both direct and indirect models and will likely be the user/device group with the highest share of direct CSP sales.

The indirect sales model of selling connectivity through the application allows applications with specific connectivity needs to include the cost of connectivity in the service fee to their customers. In the examples given above, this could be included as an app subscription fee, an in-app multi-player pass for the AR game, or as part of the monthly volume licensing fee that a business pays for the communications application. Other examples could be a telehealth provider including guaranteed connectivity in their application within the fee to ensure that patients will always have uninterrupted high-quality calls with their doctor, or a stock-trading app that includes deterministic connectivity services to ensure all orders are sent at exactly the right time, paid as part of the stock transaction fee.

**Transforming the value chain for delivering differentiated connectivity**
Integration of differentiated connectivity with applications would require developers to sign individual contracts with CSPs around the world. This could be the right approach for certain cases (such as large enterprises with enterprise-managed devices, or with full-stack devices). However, to reach the mass market of developers and end users, a scalable, easy way to expose advanced network capabilities is needed. This is enabled by the introduction of an aggregator as an intermediary [9]. Figure 2 illustrates how this indirect sales model would be structured.

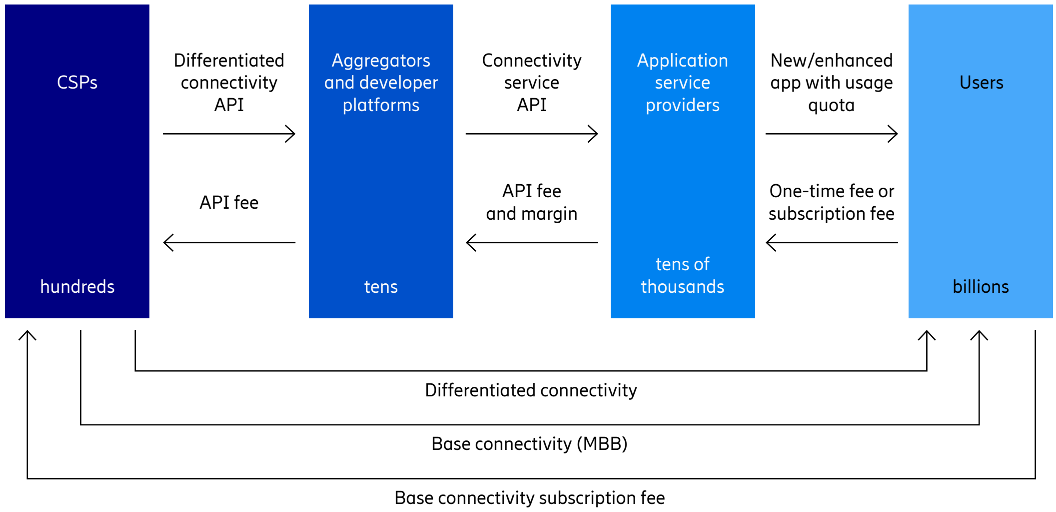


Figure 2: Value chain for delivering differentiated connectivity, coexisting with the traditional MBB value chain.

In the indirect model, the value chain changes for delivering applications and connectivity to end users. From a model where the CSPs and the application service providers independently deliver parts of the experience to the consumer, the future will see these players jointly delivering an application experience. The future delivery model will not replace the current one; they will coexist.

In the new differentiated world, revenue models for connectivity will change. The regular MBB subscriptions will remain, but they need to be complemented with new models that better meet the dynamic nature of differentiated connectivity. Therefore, pay-per-use models will dominate upstream in the value chain where the CSP charges the aggregator for the differentiated connectivity usage per consumption unit, such as time unit (hour) or data unit (megabyte). In turn, the aggregator packages connectivity from several CSPs into global APIs and provides these to the developer platforms at a margin. These platforms then package the global APIs into developer services and sell them to global app developers.

# Introducing performance levels

Developers designing applications for best-effort MBB have a notion of what 4G and 5G networks can handle. Every user who has tried to use modern applications and suddenly ended up with 3G-only coverage knows that applications do not cope with performance below a certain level. In the same manner, when differentiated connectivity beyond best effort is introduced, the application service provider community needs clear guidelines on what can be expected from different "levels" of differentiated connectivity and how their application behaviors fit with those.

Early exploration will likely create a scenario where industry players introduce a range of different connectivity services, unique for each CSP and without coherence between markets, making application design difficult for application service providers. Therefore, we propose that the industry aligns on a select few guidelines and levels, which simplifies the approach for CSPs, enabling them to focus their service offering creation and network investments to attract maximum interest from a broad application developer community with global reach and scale. Industry-aligned guidelines and levels also simplify the approach for aggregators and developer platforms to create what global application service providers need most, namely coherent global network API offerings based on reasonably similar connectivity offerings across multiple CSPs in different markets.

**Performance levels: Industry-aligned specifications of network connections**
A performance level specifies the characteristics of a bi-directional network connection and is defined by KPIs (such as latency and throughput) that application traffic reaches when a network connection with that level is used. Higher levels with lower latency and a higher throughput should be specified over time when networks become increasingly capable. Figure 3 (lower part) depicts an example where one performance level "1.0" has been introduced in the network, along with an additional "2.0" level for especially demanding application flows.

We envisage that the industry will explore different performance level-based offerings, and over time, will align on a few selected performance levels widely known and available across markets, enabling application service providers to design demanding applications for those that scale beyond a single CSP. As a complement, there will be offerings based on CSP-unique performance levels, particularly to serve the full stack Internet of Things (IoT) device/user segment where customization can be more essential than global application scale.

**Creating a CSP connectivity service offering using performance levels**
When a CSP wants to offer a performance level-based 5G connectivity service to its users for use with certain application flows, that CSP needs to design a corresponding connectivity service offering and dimension the network for such traffic. These service offerings by the CSP are sometimes referred to as "slices" [10], but this article uses the more generic term "service offering", since "slice" often denotes specific technical realization in terms of slicing signaling and identifiers specified by 3GPP [11].

Figure 3 shows the relation between applications, aggregator offerings, CSP connectivity service offerings and performance levels. Performance levels are defined as connection KPI specifications. The CSP uses network features and configurations to create a network realization with connection characteristics corresponding to the performance level objectives for throughput and/or latency. Other non performance-related characteristics, such as security, can be added to form a network slice realization. Then, a full connectivity service offering needs to be complemented with the target area of coverage and support services, such as developer instructions and customer support. The offering will depend on the chosen sales model (direct or via an aggregator), with implications on service pricing, expected volume, target customers, promotion, delivery terms and naming convention. That way, a single performance level can be the basis for multiple connectivity service offerings from a single CSP (as shown in Figure 3), and thus also for adequate differentiation between CSPs, spurring competition. CSPs can also compete by adding service offerings based on higher performance levels and with wider coverage across their market. The service offering could start at targeted areas associated to the specific need of an application, while the CSP keeps adding coverage as the requirement expands.

To ensure trust in the deterministic characteristics of the service offering, the CSP can pursue several approaches. First, performance transparency is important. The CSP can offer visualized expected and real-time statistics of their service in monitoring dashboards and monthly reports. As another layer of assurance, they can also provide SLAs guaranteeing high availability, in particular when selling to enterprises. Second, CSPs could provide free-tier pricing and beta releases of new offerings to allow customers to see the value firsthand before they put it into production. Finally, external validation of the service is important, with, for example, third-party certification and customer reviews and referrals.

Similar to most other mobile connectivity service offerings, differentiated connectivity offerings are typically only valid for 3GPP access – performance level KPIs are not guaranteed if the device connects to Wi-Fi. Full stack IoT devices and enterprise-managed devices can be configured to camp only on 3GPP during use times of differentiated connectivity offerings, securing service objective fulfillment. Consumer smartphones, however, are under the control of end users who will need to be educated on how to relate to Wi-Fi connectivity, to get the desired performance when using demanding applications.

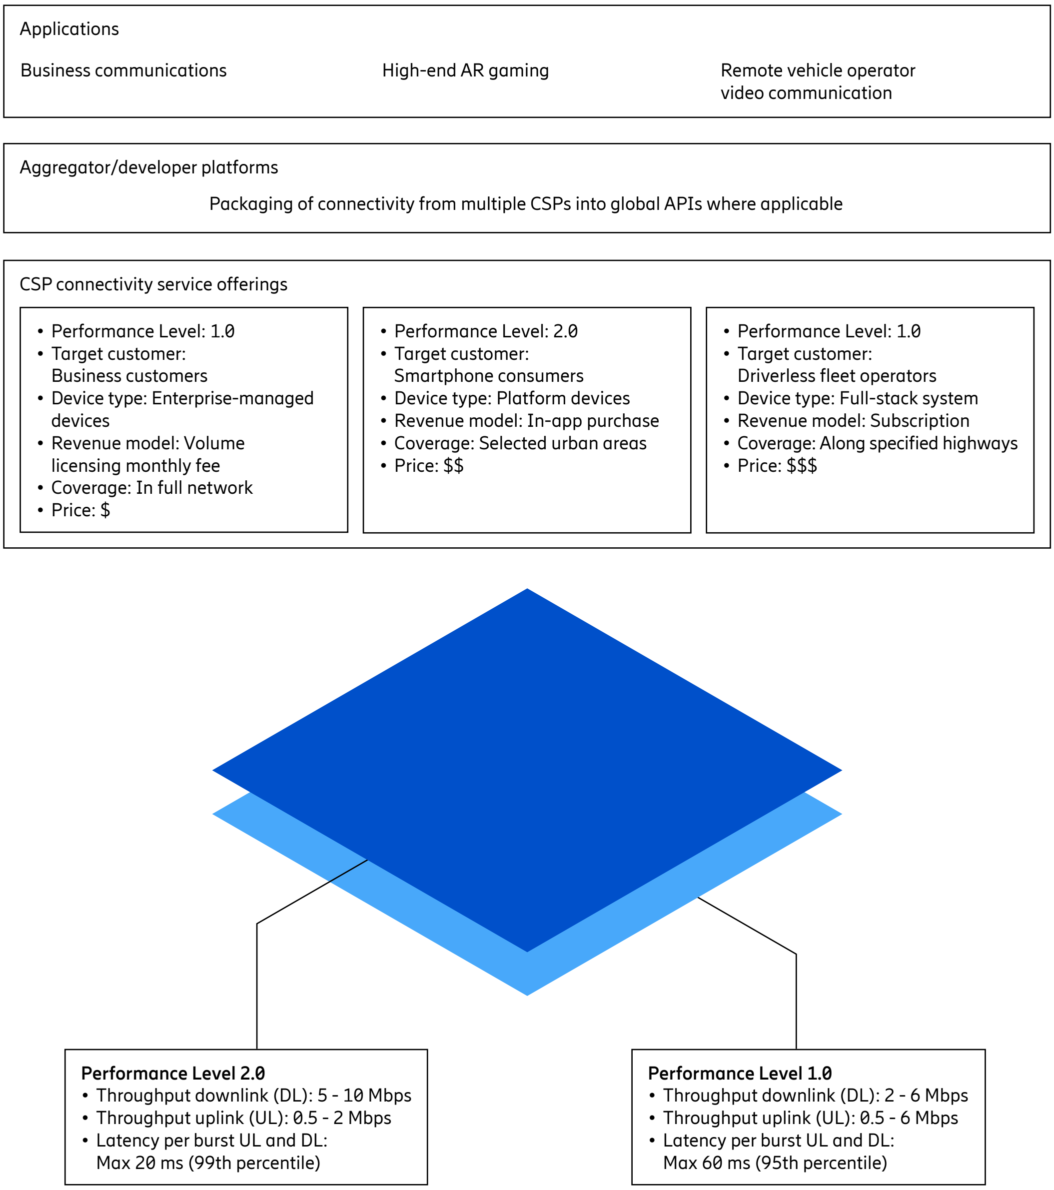


Figure 3: Example relationship between applications, aggregator/developer platform offerings, CSP connectivity service offerings and performance levels.

# Creating a differentiated connectivity service offering in seven steps

Key technical concepts are needed to realize differentiated connectivity. Most of these concepts either already exist in today's networks and devices, or are on the way to being released in standards and products.

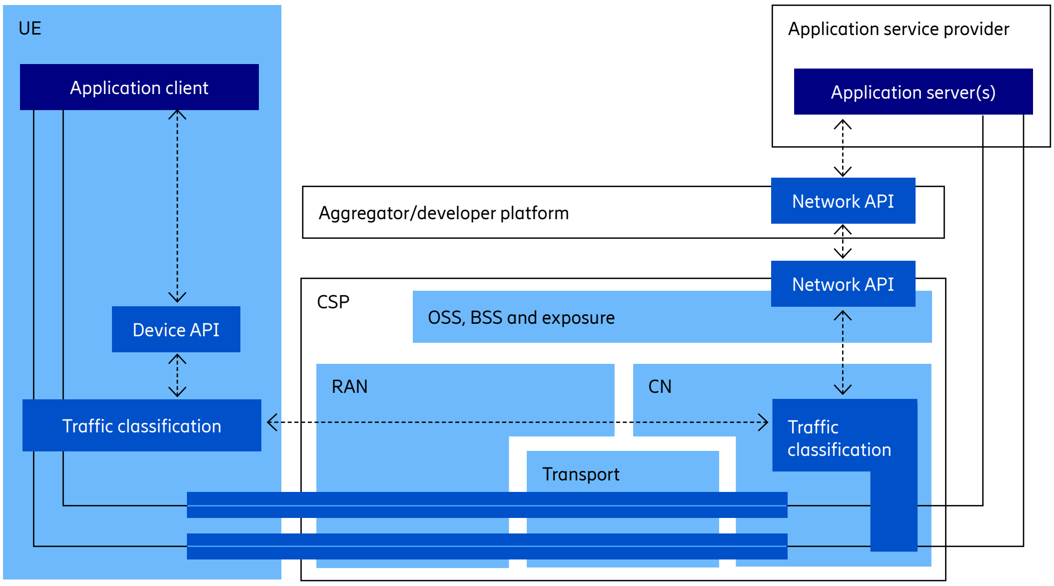


Figure 4: Functional architecture overview, depicting the application client and server, the CSP's mobile network, and the aggregator/developer platform offering network APIs and the UE.

Figure 4 illustrates the high-level functional architecture of ecosystem entities involved in creating, setting up and managing differentiated connectivity service offerings. To enable differentiated handling of flows through the network, especially in RAN, traffic classification functionality that maps different applications and application flows for a specific UE to different network resources (network slices, Packet Data Unit (PDU) sessions, QoS flows and Radio Bearers) in both uplink (UL) and downlink (DL) is key. Traffic classification based on Network-Initiated Quality of Service (NI-QoS) has been standardized in 3GPP since Release 7 [12] and is used in commercial operations mainly for Voice Over LTE (VoLTE)/IP Multimedia Subsystem (IMS). It can be an excellent basis for network APIs, especially for full-stack systems as it is designed based on traffic classification that is dynamic and finely granular. Wider adoption for platform devices and enterprise-managed devices is hindered primarily by initiatives to encrypt traffic end-to-end over mobile networks [13]. For such devices, URSP offers a standardized approach [14] to enable CSPs to influence the mapping of application traffic to network resources. Such policies provide the device QoS functionality with traffic steering rules (see examples [15,16]), or in other words, which PDU session a certain traffic type is to use. As separate PDU sessions may have different QoS treatments, these rules can then control the QoS treatment as well.

The process of creating and launching a new differentiated connectivity offering can be broken down into seven steps. The following sections detail these steps one by one, including the interactions with the corresponding processes run by the application providers and aggregators consuming the differentiated connectivity offerings.

**1. Service offering development**

The first step in creating a differentiated connectivity service offering is to define and detail it, build a marketing plan, and then create a business case considering all the use cases and use places that can benefit from the new connectivity service. As described in the previous chapter, the performance level first needs to be complemented with service targets, such as coverage, on top of the performance level objectives such as throughput and latency, as illustrated in Figure 3. The CSP should also understand the target customers for the product with an understanding of how the specific offering would appeal to them. Testing the offering with the intended target market is a good way to validate service value at this stage.

The marketing plan will contain guidance on how the product should be positioned, priced and promoted to reach the target customers. Considerations about the revenue model should also be part of this planning stage. The CSP also needs to consider how to best take the service to their intended customers, either in a direct or indirect sales model as described in earlier chapters. This will also serve as input for compiling the service business case, which should include a sales forecast and projections on revenue and costs of developing and launching the differentiated connectivity service offering.

To understand the business case, it is important to align the potential offering with the current network realities by assessing the network's ability to deliver the service and the production cost of the service. As input for the next step (network design), the detailed offering needs to include forecasts on service uptake for network dimensioning, and business guidance on network resource allocation. This is needed to guide the network organization on the business value of fulfilling or exceeding the service performance objectives.

**2. Network upgrade/design/provisioning**

The next step is to carry out the necessary network design to support the required specifications and prepare the network for the introduction of performance levels.

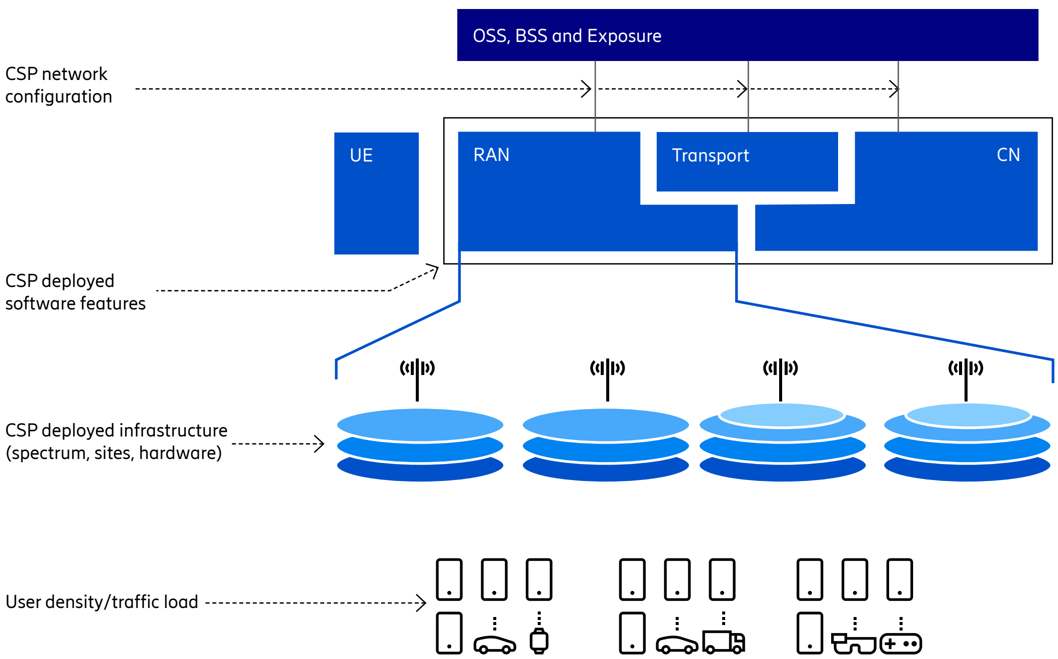


Figure 5: Factors for CSP consideration in planning for a differentiated connectivity service

As illustrated in Figure 5, the ability to provide the service depends both on CSP-deployed infrastructure in terms of spectrum, sites and hardware, but also on the software features implemented in the network, how the network is configured, and the user density and traffic load.

If a CSP wants to provide the new service in areas where the resources of the deployed infrastructure are insufficient, additional network resources need to be installed. This can involve, for example, upgrading to more capable radios, deployment of new spectrum and radios on existing sites as well as deployment on new sites.

Once the deployed infrastructure is sufficient, the right software features and configurations need to be established across the RAN, transport, Core and management domains. This includes specifying the end-to-end realization outlined in Figure 4 in detail, including software features and configurations to secure the desired performance considering both capacity and coverage for both DL and UL, network identifiers such as 5QI(s), Single-Network Slice Selection Assistance Information (S-NSSAIs), URSP rule templates, if applicable, and NI-QoS related packet filters such as for Low Latency, Low Loss and Scalable Throughput (L4S) [17] and/or IP 5 tuples. It also includes defining resource policies and QoS configurations throughout RAN, transport and Core, including offering this service and others in parallel to achieve the desired network behavior in case of congestion (where all KPI targets of all connections cannot be fulfilled). This will be familiar for CSPs already offering multiple connectivity services, such as MBB and Fixed Wireless Access (FWA).

The network design also needs to factor in BSS impacts, to ensure the new service can be purchased, activated, assured and billed. The design also needs to reckon with the network API exposure infrastructure impacts to make the offering available via APIs toward aggregators and application service providers. These system impacts will differ depending on whether the service is sold directly to an enterprise, which could be customer-tailored, or in an indirect sales channel, which leaves little room for unique solutions.

We foresee that an evolution toward intent-driven networks, with intents as a new management paradigm to guide and automate the network behavior, is essential to manage complexity and the network production cost, both in the provisioning and assurance of the new service [18]. Getting the end-to-end realization in place while meeting all the service targets and constraints and balancing between all the other services supported by the network can be done with traditional manual configuration methods as long as the number of services is moderate. In the mid term, once differentiated connectivity has taken off at scale with a multitude of service offerings, programmability and intent-driven automation will be key.

The network design typically needs to be iterated with the service offering responsible in order to strike a balance between cost of implementation and revenue potential. Once the design has been agreed, it can be implemented and provisioned in the network to support the new offering.

**3. Service offering launch**
Once implemented in the network, the service can be launched and promoted to the target customers according to the plan outlined in step 1. As part of the launch, the CSP also needs to launch any supporting services that help their customers use the service. This could be, for example, technical documentation, training material, and customer support channels. These need to be in place regardless of how the service is brought to the market.

The right agreements must also be in place for the launch. In the current business landscape, developer platforms assume no responsibility for the availability of underlying services in telecommunications networks. This is expected to change with the introduction of APIs for differentiated connectivity. CSPs should therefore provide the services with SLAs on service availability so that all tiers of the delivery chain can, in turn, issue their own SLAs on that underlying service availability. For application service provider customers, one SLA with a developer platform could then cover service availability in multiple networks.

**4. Applications: Preparatory steps**
The shift to deterministic connectivity with pre-defined performance levels means that applications can be designed to deliver a predictable service experience, enabling new application "modes" beyond those used for best-effort connectivity. Application service providers need to discover and realize the value of a differentiated connectivity offering sold via developer platforms or the CSPs directly.

Application service providers need to first plan what the service offering, which includes predictable connectivity, would look like. This will include designing user experience journeys by considering how to involve end users in the sign-up, activation and accounting of new connectivity offerings.

Applications also need to implement support for traffic classification, which for example could be based on 5G network slicing with URSP including device and network API integration, the ability to set up multiple PDU sessions, and the ability to send the right traffic on the right PDU session as illustrated in Figure 4.

**5. Subscriber sign-up step**
Once networks and applications are prepared and the offerings have been launched, subscribers can sign up for the service. For direct sales of premium connectivity, where the premium connectivity service is provided as an extension of an existing subscription or as an additional subscription, the activation results in user provisioning of the necessary subscription data and in case of URSP, device provisioning of updated URSP rules. For enterprise sales, an enterprise IT administrator can define how the enterprise applications are mapped to the CSP-provided connectivity services in the device via a Mobile Device Management solution.

For sales via network APIs, the premium connectivity is not directly visible to the CSP's subscribers, so there is no specific sign-up toward the CSP. Instead, subscribers, who are also users of the application service provider's applications, will sign-up for a service from the application service provider that will utilize premium connectivity.

**6. Session runtime: Initiation of premium connectivity**
When using network APIs to initiate the premium connectivity service, as illustrated in Figure 4, a key part of the process is to ensure that the appropriate RAN and Core network configurations for the service that the application initiates can be identified. Additionally, it must be ensured that the traffic classification in both UE and the CSP's network is set up consistently to map the traffic to the correct network resources down to the radio bearer. The application requests the premium connectivity to be activated either via a CSP network API, or an aggregated API via an aggregator. In case of an aggregated API, the aggregator resolves which CSP the device is served by and requests the service from that CSP. The request needs to uniquely identify which level(s) the application intends to use, as well as any applicable device information, to enable the CSP to set up the traffic classification correctly and identify the corresponding S-NSSAI, data network name (DNN) and 5QI to use in communication between RAN, Core and UE. In the CSP's network, the correct subscriber data is provisioned in the Core network to entitle the use of the service and in case of URSP triggers, the appropriate URSP rules are provisioned to the device.

Once this is done, the application can establish the premium connectivity service via the device QoS functionality. To enable this, it is key that the service that the application requests via the device QoS functionality can be derived deterministically from the service that the application requests via the network API. In case of URSP, the device QoS functionality then ensures a PDU session is established with the correct S-NSSAI and DNN as provisioned in the URSP rules, while the Core network, in turn, can provide the S-NSSAI and 5QI to RAN to close the loop.

**7. Service assurance**
The final key step in launching a differentiated connectivity offering is the service assurance [19]. As CSP, it is essential to ensure that the target KPIs of all connectivity services are fulfilled and to take remedial action if performance falls below acceptable levels.

Introduction of multiple differentiated connectivity services will make operations of a mobile network in the traditional way increasingly more complex, which needs to be addressed by a higher degree of automation, making the mobile networks more autonomous. While multiple levels of closed loop automation constitute the heart of an automation solution, these closed loops need to be able to translate the business objectives defined in step 1 into the correct automated actions. One way to realize this automation is through an intent-driven automation architecture, where the network resource allocation policy will serve as guidance to automation in the form of intents.

With intents as an input, Intent Management Functions (IMFs) will be responsible for the service assurance with the objective to fulfill the intents in an optimal way given the deployed resources through closed loop actions. IMFs will be deployed at different levels, working with different scopes and to different time scales. An IMF in the radio node will be responsible for the resources within that radio node, and will be able to take actions related to things such as beamforming, scheduling and handover at timescales down to the ms level to assure intent fulfillment within the associated coverage area. An IMF in the RAN management level will, on the other hand, orchestrate multiple radio nodes over an area and will therefore have the overall responsibility to assure the RAN service over a larger RAN area. Finally, an end-to-end IMF will take care of the highest level intent, ensuring that all parts of the network (RAN, transport network, Core network) contribute to the overall intents, assuring the end-to-end network connectivity service performance.

CSPs also need to provide service support by having an incident management function. Depending on the go-to-market model, they need to prepare for providing tier-1, tier-2, and/or tier-3 support. Typical incident management functions must include incident logging and verification, incident categorization and prioritization and incident resolution.

# Industry action and alignment needs

The transformation through differentiated connectivity will enable new value to be offered to consumers and enterprises, benefiting not only the end users but all players in the industry, including application service providers, CSPs, device vendors, telecommunications vendors, smartphone platform vendors, aggregators and developer platforms. Realizing the industry transformation will require an aligned vision among the key players, and the willingness to put in efforts and investments to start making the vision a reality. Telecommunication vendors play an important role in coordinating and pushing the initial moves to start this evolution, while each of the key players needs to start acting now to embrace, trigger and scale this industry transformation.

**Immediate actions needed by industry players**

**CSP actions**
CSPs are the key stakeholders for enabling differentiated connectivity and transforming networks to innovation platforms. For immediate actions to take, CSPs should start building business cases for offering differentiated connectivity and align internally on the vision of the new revenue stream and business model; define product commercial packaging as to how to offer connectivity with different performance levels, and plan networks' readiness for offering and exposing connectivity with deterministic performance. CSPs also need to be prepared to work with early adopting application service providers to test the end-to-end business and technical flow to prepare the commercial launch of the new performance levels.

**Application service provider actions**
It is important for the application service providers to start identifying new experiences and new values that can be created by the support of differentiated connectivity. A useful question to consider is, "Have there been cases in the past where we had to compromise on functionality or experience that could have been developed into the app or the device, due to the assumed limited network performance?" Another way to think about it is, "Imagine what new things can be achieved if all the smartphones using the app, or the IoT devices, were now connected to an always-high throughput and low latency reliable network, as if they were all in the same local area network."

In the initial phase of exploration, application service providers should also proactively engage with leading CSPs to start implementing new features and experiences and test the new value creation. Application service providers focusing on smartphone applications should be particularly observant of those performance level-based offerings that start to scale beyond single CSPs.

**Device and device platform vendor actions**
Device and device platform vendors need to make sure the devices support key features that enable differentiated connectivity. The app development APIs need to support mapping by CSP policies (URSP) to the new performance levels on network, to allow application developers to easily implement new features and traffic flows using differentiated connectivity.

**Realizing the industry vision at a global scale**
While the key players in the industry should take immediate actions to enable and experiment with differentiated connectivity and start realizing its value, it is also important to start the alignment on how it will be scaled at a global level. All the key players must understand the importance of having standardized APIs, and the key roles of aggregators and developer platforms, to allow tens of thousands of apps to be developed easily, with the ability to reach billions of users, supported by networks with differentiated connectivity from hundreds of CSPs around the world. To achieve this, CSPs, device platforms, aggregators and developer platforms need to be actively engaged in discussions through standardization fora such as CAMARA and GSMA.

# Conclusion

The telecom industry is in urgent need of a transformation to realize the true value of wireless connectivity. This is what will turn connectivity into a platform providing deterministic performance that supports application service providers in making further innovations, in combination with other innovation enablers such as AI and cloud-based computing.

The performance levels introduced in this whitepaper are intended to make it easier for application service providers and developers to leverage the 5G standalone network capabilities and drive adoption of differentiated connectivity services. For that to happen, an industry aligned on the specifications needed to expose the wireless connection characteristics is required. This will foster services innovation and enable CSPs to offer such services more widely and with greater variety of differentiation. It will also be pivotal for the telco industry to move into performance-based business models for connectivity.

There are already clear indications in the industry that differentiated connectivity is needed, and many CSPs around the world are already trying to explore its potential. This paper introduces the steps a CSP should take to create differentiated connectivity services.

To trigger and accelerate the industry transformation, key players in the industry need to start acting now, including jointly testing value creation through the new performance levels and the new business models, and aligning on the implementation through APIs and standardization to achieve global scale.

# Glossary

| | |
|---|---|
| **3GPP** | 3rd Generation Partnership Project |
| **5QI** | 5G QoS Identifier |
| **AI** | Artificial Intelligence |
| **APIs** | Application Programming Interfaces |
| **BSS** | Business support systems |
| **CN** | Core Network |
| **CSPs** | Communications Service Providers |
| **DL** | Downlink |
| **GSMA** | Global System for Mobile Communications Association |
| **IMFs** | Intent Management Functions |
| **IMS** | IP Multimedia Subsystem |
| **IoT** | Internet of Things |
| **KPIs** | Key Performance Indicators |
| **MBB** | Mobile Broadband |
| **MIMO** | Multiple Input Multiple Output |
| **NI-QoS** | Network Initiated Quality of Service |
| **OSS** | Operations Support Systems |
| **PDU** | Packet Data Unit |
| **QoS** | Quality of Service |
| **RAN** | Radio Access Network |
| **SLAs** | Service Level Agreements |
| **S-NSSAI** | Single-Network Slice Selection Assistance Information |
| **UE** | User Equipment |
| **UL** | Uplink |
| **URSP** | UE route selection policy |

# References

1. History of mobile internet – transformative tech - Ericsson
2. Meeting 5G network requirements with Massive MIMO - Ericsson
3. Technology Review Magazine 2023 – Spotlight on extended reality - Ericsson
4. https://www.att.com/plans/att-turbo/
5. https://www.singtel.com/personal/products-services/mobile/5g/5g-express-pass
6. https://www.telstra.com.au/optimisers/mobile-speed-optimiser
7. https://www.ais.th/en/consumers/lifestyle/new-myais/mynetwork/mynetwork-mobile
8. https://web.three.com.hk/plans/5gliveconnect/index-en.html
9. Global network API platform to monetize 5G - Ericsson
10. T-Mobile DevEdge | 5G Network Slicing
11. 3GPP 23.501
12. "QoS control in the 3GPP evolved packet system", Hannes Ekström, IEEE Communications Magazine, Volume 47, Issue 2, February 2009
13. A collaborative approach to encrypted traffic - Ericsson
14. 3GPP 24.526, Section 4.2
15. Apple iOS Network Slicing
16. Android 5G Network Slicing
17. Enabling time-critical applications over 5G with rate adaptation - Ericsson
18. 5G monetization via intent-based network operations - Ericsson
19. Service quality monitoring in the digital economy - Ericsson
20. Autonomous networks with multi-layer, intent-based operation - Ericsson

# Further reading

1 5G Advanced, your network for the next wave of 5G - Ericsson

2 5G standalone (5G SA) experience 5G without limits - Ericsson

# Authors

**Håkan Olofsson** has worked in the mobile industry for 30 years, with a particular focus on its radio access network aspects. He joined Ericsson in 1994 and has served in a variety of capacities, mostly dealing with strategic technology development and the evolution from 2G all the way to 5G. His background also includes a managing consultant role at strategy firm Northstream in Sweden and France. He is currently head of Systems, Concepts and Algorithms at Ericsson, and a board member at ITRL, KTH and Vinnova Drive Sweden. Olofsson holds an M.Sc. in physics engineering from Uppsala University, Sweden.

**Karl Olsson** joined Ericsson in 2020 and currently serves as a Group Strategy Manager. His work focuses primarily on strategy development and business development for emerging technologies. Olsson holds an M.Sc. in Industrial Engineering and Management from the faculty of engineering at Lund University, Sweden.

**Fredrik Alriksson** is an expert in system architecture for the Internet of Things (IoT) and new services at Business Unit Networks, where he leads strategic technology and concept development related to the IoT and industry verticals. He joined Ericsson in 1999 and has worked in R&D with architecture evolution, covering a broad set of technology areas including RAN, Core and IP Multimedia Subsystem. Alriksson holds an M.Sc. in electrical engineering from KTH Royal Institute of Technology in Stockholm, Sweden.

**Abhey Seth** works as Head of Product Management (acting). He is responsible for bringing out advanced network capabilities to the application developer community via network APIs and driving business growth in the enterprise space. He also manages industry partnerships, platform strategy and solution development. In his 14 years with Ericsson, Abhey has held various roles, such as Sales Readiness Manager for Enterprise and IoT communication solutions, Solution Architect for Communication Services, and Pre-sales engineer for IMS networks. Abhey has deep technical expertise in APIs, Core and IMS networks. He holds an Electronics and Communication engineering degree and has worked with global teams across Ericsson in India and Sweden.

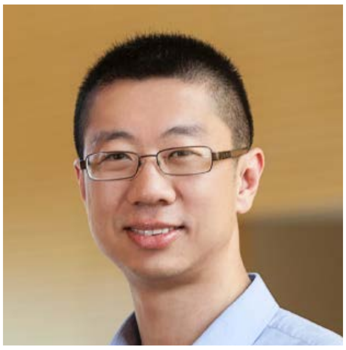

**Bo Zhong** currently serves as a Strategy Manager in Ericsson Business Area Networks. He has 10 years' experience in the Telecom industry and has worked mainly on combining new applications and services with wireless technologies. Before his current role, he worked both in Ericsson R&D, and as a driver of Innovation and Intrapreneurship in Ericsson's global innovation program Ericsson ONE. He is an experienced coach for Design Thinking and business model innovation. Bo Zhong holds a PhD degree in Telecommunication, from the University of London, UK.

**Jan Backman** is an expert in Packet Core Mobility Architecture. He started working for Ericsson in 1997 and has over 25 years of experience in the Telecom industry and especially in Packet Core. Jan's experience spans driver development, chief product architect, systems management as well as end-to-end architecture responsibilities. He has a master's in computer science with a focus on computer communication from Luleå University in Sweden.