

L4S Solution Brief

**Turning congestion into low-latency,
predictable experiences**



ERICSSON

Enabling predictable performance in 5G SA networks with L4S

As mobile networks evolve toward 5G standalone (SA), application experience is no longer defined solely by coverage and peak data rates, but by how effectively the network manages latency, packet loss and throughput stability under stress situations.

Contents

- 02 Enabling predictable performance in 5G SA networks with L4S
- 04 How L4S works
- 08 L4S in action
- 10 Standardization, the ecosystem and Ericsson's portfolio
- 11 Summary

In mobile environments, network conditions vary continuously due to user mobility, fluctuating load and radio constraints, making a consistent experience increasingly difficult to achieve.

Low Latency, Low Loss, and Scalable Throughput (L4S) is an Internet Engineering Task Force (IETF)-standardized congestion control framework that enables applications to adapt smoothly to changing network conditions. Through fast and precise congestion signaling, L4S reduces delays and packet loss, delivering a smoother experience for latency-sensitive applications, both those that are part of today's best-effort connectivity model, and those yet to come.

Mobile network congestion is like congestion on highways. Traffic conditions vary over time, and the average travel speed is directly affected by factors such as the vehicle type, number of vehicles, quality of the road and weather conditions. Drivers usually react only to nearby vehicles, often resulting in stop-and-go traffic, elevated stress levels and an increased risk of accidents.

A smarter approach would detect early signs of congestion and signal vehicles to slow down gradually, regardless of the underlying cause. L4S works in the same way: It provides an early "slow down" signal as congestion begins to form, without explaining why or prescribing a specific rate, allowing applications to adapt smoothly before queues and delays build up (which is analogous to the Adaptive Cruise Control in a car on the highway). L4S is a key functionality for communications service providers in their journey toward differentiated connectivity. By providing the elasticity and fast congestion control required to maintain low latency and minimal loss under varying radio conditions, L4S enables differentiated connectivity services to scale across users, applications and network load scenarios, while

supporting stronger performance guarantees for selected use cases.

Application requirements are not uniform

Application requirements vary significantly depending on their sensitivity to latency, packet loss and bandwidth fluctuations.

Some applications can tolerate delay or throughput variation, while others experience severe degradation if network conditions are unstable. Examples include:

- live and interactive video applications, such as augmented reality (AR), or remote collaboration tools, which require low latency to maintain immersion
- cloud gaming and extended reality (XR) services, which depend on stable throughput and fast feedback loops to ensure responsiveness
- remote control and automation use cases, such as delivery drones, remote driving, or industrial operations, where latency and packet loss directly impact safety and accuracy
- field-based professional services, such as media and content creation including photojournalism, where unreliable packet delivery can have immediate business consequences
- multi-modal AI applications, such as AI-assisted field operations, require continuous context awareness, stable throughput and fast feedback loops to maintain real-time responsiveness and interaction quality
- physical AI applications, such as robotics, require continuous real-time uplink of video and sensor data, where low latency and minimal jitter are critical to ensure safe and accurate remote control

For these applications, traditional best-effort connectivity is insufficient. These applications require network behavior that supports fast adaptation while maintaining a high quality of experience (QoE).



Network challenges impacting application performance

Achieving consistent application performance in mobile networks is challenging due to several factors:

- network congestion in crowded areas, large venues, or during seasonal and live events
- radio coverage limitations, especially at cell edges or in complex indoor environments
- dynamic traffic patterns, where sudden spikes in demand can impact application performance
- operational constraints, such as imperfect network planning or temporary failures

While quality of service (QoS) mechanisms can prioritize traffic, they are not sufficient under high load when used alone. Even prioritized users may experience degraded performance when available capacity becomes constrained. What is required is a way to ensure graceful performance degradation, rather than abrupt quality drops.

What L4S brings to differentiated connectivity

L4S is a congestion-signaling mechanism that enables applications to receive faster and more precise feedback from the network than what is possible with traditional loss-based approaches. This allows applications to react quickly to emerging congestion, adjusting their sending rate before queues build up and latency increases.

In 5G SA networks, L4S works together with QoS mechanisms to maintain low latency and low packet loss while enabling scalable throughput adaptation. As network load increases, L4S allows applications to smoothly adjust bitrate – and when needed, resolution or frame rate – avoiding freezes, excessive distortion, or sudden quality drops. The result is a more stable and predictable user experience, even under challenging network conditions.

What this paper covers

This paper explains how L4S works, illustrates its benefits through practical application examples, and discusses standardization progress and ecosystem maturity. It also outlines Ericsson's portfolio approach to L4S, showing how service providers can use it as a foundational capability to enable differentiated connectivity with better guaranteed latency and packet loss performance for more sensitive applications.

How L4S works

L4S requires support from both the network and endpoints to function as intended.

L4S can be viewed as a form of soft contract between applications and the network. The network indicates congestion by marking packets when queues build, while applications implement congestion control mechanisms that aim to keep the network queues short. If either the network or the application does not adhere to these principles, L4S will not deliver the intended benefits.

The interaction between the application and a 5G network is illustrated in Figure 1. The sender indicates support for L4S in the IP header of the application data. This indicates that the sender can adapt its transmission bitrate on the proportion of packets marked with “congestion experienced” (CE) by congested network nodes. Congested network nodes can then mark packets with CE as a first step. Then, the receiver detects these marked packets, and feeds back the explicit congestion notification (ECN) status of each packet to the sender, which then enables closed-loop congestion control.

L4S should not be considered a QoS enhancement – rather, it complements existing QoS mechanisms. For example,

ECN is a standard IP-layer capability that allows network nodes to signal congestion by marking packets, instead of dropping them. It defines how congestion is indicated (via ECN bits and CE markings) but, by itself, ECN does not define how endpoints should react to congestion, nor does it guarantee low latency. Traditional ECN is typically used with loss-based or hybrid congestion control, which still allows queues to build.

L4S, on the other hand, introduces an alternative way that uses ECN as its primary congestion signal, but adds rules on how that signal is generated and interpreted. The network marks packets early on, as soon as queues begin to grow. Applications that implement L4S-capable congestion control algorithms react proportionally to the quantity of CE-marked packets. The combined behavior of network and endpoints keeps queues consistently very small, enabling low latency and low loss. In this sense, L4S can be viewed as a QoE enhancement, but not as a QoS feature.

L4S aligns well with 3GPP access networks, which already offer a rich assortment of QoS capabilities. It also introduces a clear separation of

responsibilities across applications and transport protocol layers, and the network elements, as illustrated in Figure 2. Applications and transport protocols are responsible for rate adaptation and congestion control, while network elements perform congestion marking. This separation ensures that the design and configuration of L4S congestion control algorithms are independent of the lower-layer L4S marking behavior, QoS configuration or access technology used.

As a result, applications that support L4S over 5G networks can also work with other access technologies, such as cabled networks based on the Data Over Cable Service Interface Specification (DOCSIS) standard and vice versa. This access-agnostic behavior is a key enabler for broader adoption, as it increases the incentive for application service providers (ASPs) and third-party developers to invest in L4S support.

L4S functionality is distributed across both the network and the applications, and can be grouped into three main components: the application sender, the network nodes and the application receiver.

Figure 1: Application and network interaction

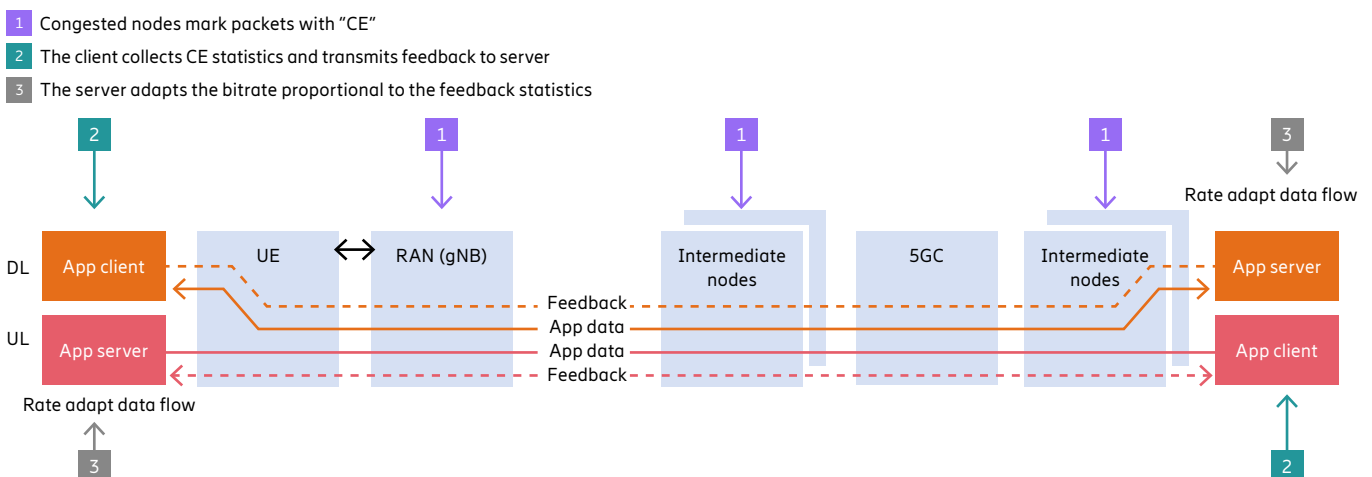
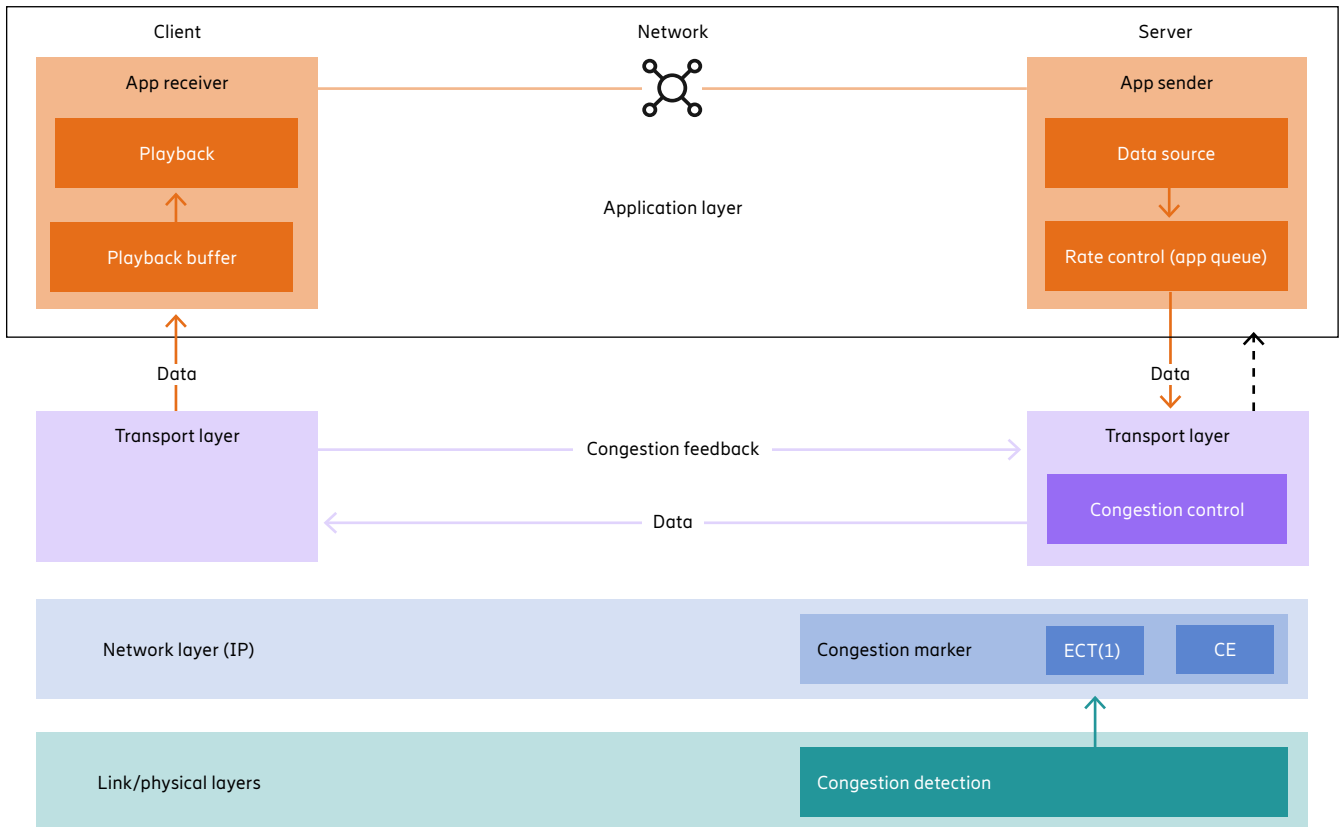


Figure 2: L4S handling in the mobile network



Application sender

The application sender is responsible for rate adaptation and adjusts its transmission rate based on feedback from the application receiver.

To enable L4S operations, the sender sets the ECN bits in the IP header of each transmitted packet to ECT-Capable Transport (1). This indicates to L4S-capable network nodes that these packets may be re-marked to CE when congestion occurs. By doing so, the sender signals its ability to react promptly to fine-grained congestion feedback rather than reacting later once packet loss or other heuristics, like delay, begin to occur.

Data packets are transmitted using standard transport protocols such as Transmission Control Protocol (TCP), Quick UDP Internet Connections (QUIC) or Real-time Transport Protocol over User Datagram Protocol (RTP over UDP), depending on the application type.

Network nodes

Network nodes on the path between sender and receiver typically maintain queues and may drop packets when the incoming bitrate exceeds the capacity of an outgoing link. An L4S-capable node, such as a 5G gNB, instead marks L4S-capable packets when queues begin to build. This is done by re-marking packets from ECT(1) to CE, providing an early and explicit congestion signal to the endpoints. In 5G SA networks, L4S can be used in combination with existing QoS mechanisms for differentiated connectivity, described later in this paper. The specific algorithm used for L4S marking is typically proprietary and dependent on the characteristics of the access technology.

Application receiver

The application receiver collects information from received packets and sends feedback to the application sender using dedicated feedback messages. The feedback mechanism depends on the transport protocol in use. TCP and QUIC use dedicated acknowledgement (ACKs), while RTP over UDP relies on RTCP for feedback signaling. This feedback includes information about congestion markings and enables the sender to adjust its transmission behavior accordingly.

Figure 3: ECN codepoints

Binary codepoint	Codepoint name	Meaning
00	Not-ECT	Not ECN-capable transport
01	ECT(1)	L4S-capable transport
10	ECT(0)	ECN-capable transport
11	CE	Congestion experienced

Note: The ECN field in the IP header uses the codepoints listed in Figure 3.

Congestion control and rate adaptation

L4S enables the transport layer to keep network queues short and network delays low. This is achieved with the use of a scalable congestion controller that manages the amount of data in flight in the network based on the fraction of packets that are CE marked. The data in flight is the amount of data along the path between sender and receiver.

The application layer is responsible for ensuring that the end-to-end latency remains low. The actual rate adaptation strategy depends on the type of application or media, such as video streaming, HTTP-based services or AI workloads.

For example, in video applications, the target bitrate can be calculated based on the amount of data that can be in flight in the network and the measured round-trip time (RTT):

$$\text{target_rate} = \frac{\text{data_in_flight}}{\text{RTT}} \text{ [bps]}$$

The throughput is given by the relation data_in_flight proportional to $1/\text{pMark}$, among other factors like RTT and packet size, where data_in_flight is in [bits] and RTT is in [s].

In addition to the calculation of the target bitrate above, applications may also rely on push-back signals from the transport layer to adjust the amount of data transmitted. Such adjustments can take different forms, including changes in compression level, updates to the frequency of images or pruning of data objects. Regardless of the

method used for congestion control and rate adaptation, the objective is to keep end-to-end latency low, including delays introduced in the application layer.

L4S implementation in the RAN and core network

L4S in 5G SA networks relies on the existing 5G QoS framework to enable predictable handling of latency-sensitive traffic while preserving end-to-end, application-driven congestion control.

In the 5G Core, the User Plane Function (UPF) can be configured to inspect the IP packets to identify L4S-capable traffic. This detection is based on standard ECN/L4S semantics at the IP layer. Once identified, these packets are mapped to a dedicated QoS flow associated with a specific 5G QoS identifier (5QI). The approach enables consistently low latency and predictable performance compared to traditional traffic management methods. In Ericsson's 5G Core solution, this can optionally be combined with 5-tuple identification, providing differentiation of L4S per service.

For uplink traffic, the Session Management Function (SMF) provides Traffic Flow Template (TFT) rules to the user equipment (UE), ensuring that L4S-capable traffic is correctly classified and steered to the appropriate QoS flow. This mechanism establishes a consistent treatment of L4S traffic across the core and access network domains. By reusing the standardized 5G QoS framework, L4S can be

deployed alongside the QoS features that are available.

As illustrated in Figure 4, L4S-capable traffic in the downlink is carried as a dedicated QoS flow within a common packet data unit (PDU) session, alongside other traffic such as mobile broadband. This allows L4S traffic to be isolated from potentially bursty non-L4S flows, preventing uncontrolled queue build-up, while avoiding the operational complexity of separate PDU sessions or application specific handling in the core network.

Congestion detection and L4S marking are performed in the 5G RAN, where real-time information about radio conditions, scheduling behavior and instantaneous cell load is available. By placing congestion marking close to the air interface, the network can generate early and accurate congestion signals that reflect actual radio constraints. These signals are conveyed end-to-end between the application endpoints, allowing the application and transport layers to remain responsible for rate adaptation and congestion control.

While L4S enables early congestion signaling and fast rate adaptation, its effectiveness in mobile networks also depends on ensuring that radio-layer effects do not distort the congestion signals seen by the application. Sudden changes in radio conditions, or short-term capacity reductions can introduce delays and jitter that are not directly caused by congestion, but may still affect how applications interpret congestion signals.

Figure 4: Downlink L4S and mobile broadband QoS flows mapped to a common PDU session in a 5G SA network

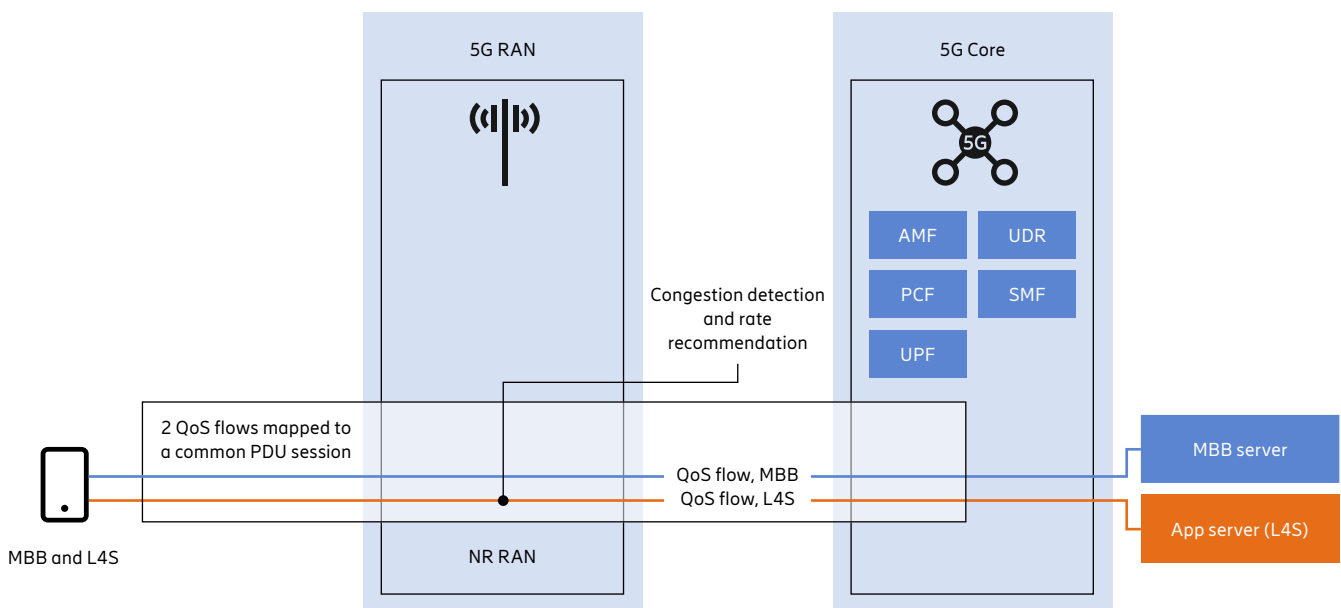
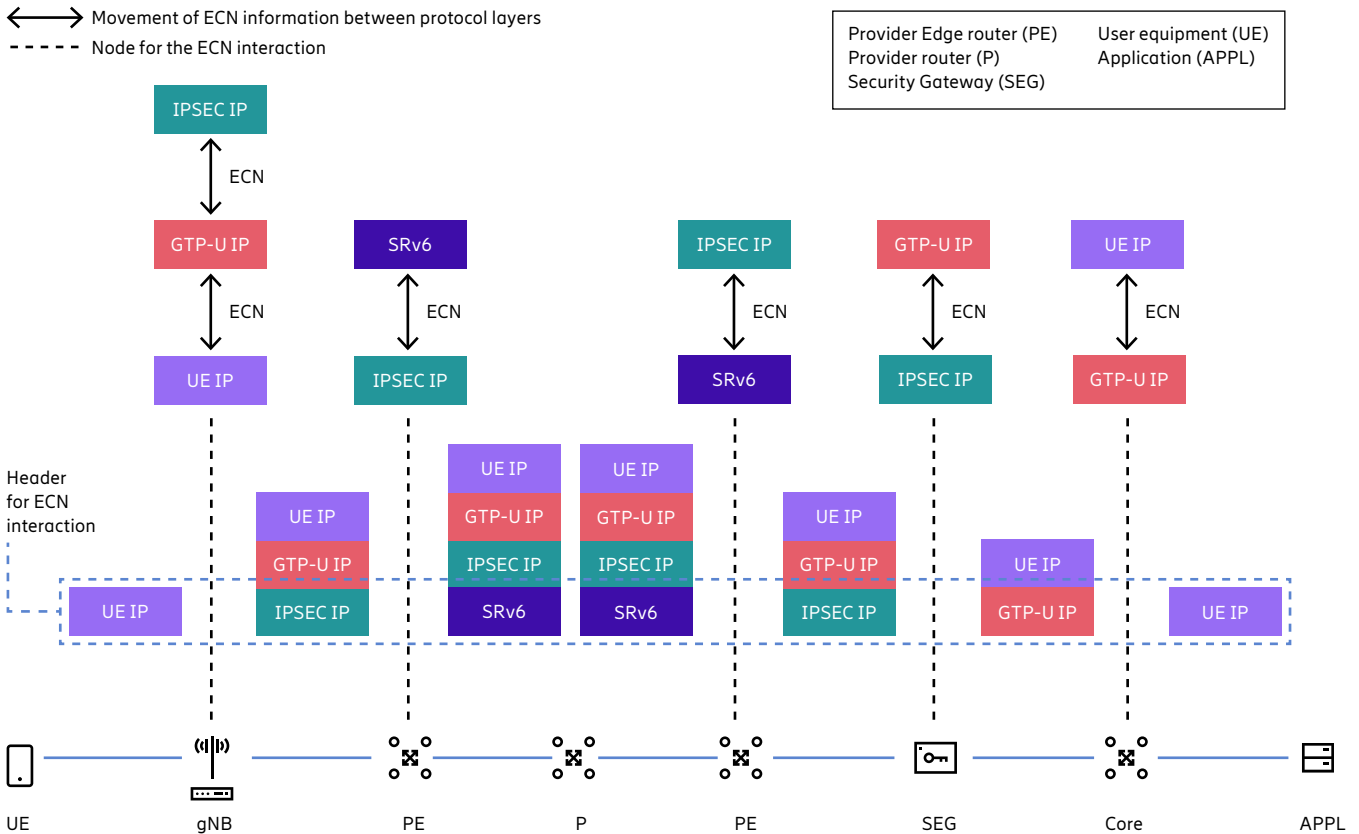


Figure 5: ECN congestion marking in transport network



Ericsson 5G Advanced RAN software addresses this by complementing L4S with radio scheduling and prioritization, delivered as part of the 5G Advanced RAN differentiated connectivity subscription. This results in:

- More timely uplink scheduling reduces delay and jitter unrelated to congestion. Features such as shorter scheduling request (SR) periodicity, uplink configured grant (UL-CG), or uplink pre-scheduling enable more predictable uplink transmissions, helping ensure that L4S congestion signals remain accurate and preventing unnecessary rate reductions at the transport or application layers.
- Latency priority scheduling (LPS) dynamically increases scheduling priority when queue delay exceeds a defined threshold, for example, due to a sudden reduction in available link capacity. This compensates for scenarios in which applications, such as video encoders, cannot instantly react to congestion signals. By smoothing short-term latency spikes in the 5G RAN, this mechanism allows L4S-enabled applications sufficient time to adapt without compromising user experience.
- Rate-controlled scheduling (RCS) increases scheduling priority when the effective bitrate falls below a configured minimum threshold. This is particularly beneficial for applications that require

guaranteed baseline quality, such as video streaming or remote-controlled vehicles where a minimum video quality is needed for safety reasons. By increasing the likelihood that minimum application requirements are met even under congestion, RCS allows L4S-enabled applications to remain responsive while maintaining essential service quality.

L4S congestion marking in the RAN is performed in the gNB for both uplink and downlink traffic. By basing congestion decisions on a wide range of radio performance metrics that are only available at the gNB level, the network can generate more accurate and timely congestion signals. This tight integration between L4S congestion signaling and 5G Advanced RAN capabilities is crucial for delivering consistently low latency under load, rather than being limited to ideal network conditions.

L4S implementation in the transport network

As congested transport networks are relatively common, L4S support in this domain can be needed to ensure full end-to-end support, as L4S will not work as intended if there are congestion points in the network that do not congestion mark the L4S-enabled traffic flows.

L4S support in transport networks requires that L4S-enabled flows are handled according to the principles outlined in the L4S standards. This means identifying and placing L4S-enabled traffic flows in a separate queue to measure queue build-up and congestion-marked packets that have been queued too long. Congestion marking in the ECN field will be done in the outermost IP header, which is why copying the ECN field between IP headers when the packets are encapsulated or decapsulated is needed (see Figure 5).

Key takeaway

L4S keeps latency and packet loss low by signaling congestion early and explicitly. Network nodes mark packets as queues begin to build, and L4S-capable endpoints adjust their sending rate in proportion to the markings. This closed-loop interaction

keeps queues short and enables predictable, graceful performance under load. In combination with QoS features, L4S enables opportunities for differentiated connectivity offerings that can offer high throughput and low latency.

L4S in action

In collaboration with Ericsson, SoftBank and MasOrange have implemented L4S to prepare their networks for handling advanced capabilities and experiences.

Live case 1: SoftBank Corp.

To prepare its live 5G SA network to support XR services, SoftBank Corp. ("SoftBank") implemented L4S together with 5G Advanced capabilities such as Critical IoT (L4S, UL-CG) and differentiated connectivity (RCS), in collaboration with Ericsson and Qualcomm Technologies, enabling ultra-low latency XR services on a live 5G SA network.¹

Goals and needs

Applications such as XR place stringent performance demands on mobile networks. SoftBank aimed to support real-time XR services requiring continuous, stable and ultra-low-latency connectivity. XR performance is directly impacted by latency fluctuations and packet loss, even when average network speeds are high, making traditional best-effort delivery insufficient. The objective was, therefore, to deliver predictable, low-latency communication suitable for immersive XR experiences on a commercial 5G SA network. In parallel, SoftBank required a solution that enabled service differentiation, allowing latency-sensitive traffic to be optimized without affecting other users or services sharing the same network infrastructure.

Challenges

Mobile networks are typically designed to balance resources fairly across users, which can introduce latency spikes under load. For XR applications, such behavior leads to unstable performance and degraded QoE. The challenge was to reduce latency and packet loss in the wireless link. This had to be achieved in a live network environment, where performance optimization needed to be tightly controlled and isolated to the relevant devices and application flows.

Solution

To address these challenges, SoftBank deployed a combination of L4S and 3GPP-standardized 5G Advanced features. L4S enabled congestion to be managed more efficiently and responsively. Configure Grant let devices transmit uplink data immediately using pre-defined schedules to reduce latency, and RCS-enabled base stations dynamically managed resources to maintain configured throughput. Network slicing was applied to ensure that these capabilities were limited to the XR devices and application, enabling targeted optimization without impacting other network traffic.

Achievements and results

The solution delivered a reduction of approximately 90 percent in wireless link latency compared to scenarios not using 5G and 5G Advanced features such as L4S. It enabled continuous and stable

SoftBank achieved a reduction of approximately 90 percent in wireless link latency compared to scenarios not using 5G and 5G Advanced features such as L4S.

90%

low-latency communication, meeting the performance requirements of XR-content streaming using smart glasses connected via a smartphone. This customer example demonstrates how L4S, combined with 5G Advanced capabilities, can deliver differentiated connectivity for latency-critical services on a live commercial network, supporting new classes of real-time applications.



¹ Ericsson, "SoftBank Corp., Ericsson, Qualcomm Technologies trial 5G and 5G Advanced capabilities including L4S on 5G SA commercial network" (December 23, 2025).



Live case 2: MasOrange – advancing 5G Advanced with L4S-enabled experience²

More predictable and responsive user experience is expected as 5G moves beyond best-effort connectivity. MasOrange and Ericsson are advancing Spain's mobile experience by introducing 5G Advanced with a strong focus on ultra-low and consistent latency, enabled through L4S.

Goals and needs

As consumer services increasingly depend on real-time interaction, MasOrange sought to deliver consistently low and deterministic latency across its network. The goal was to support emerging use cases such as real-time video, cloud gaming and AR/VR, where even small variations in latency can degrade the user experience. To do this, MasOrange needed network capabilities that could handle congestion intelligently while maintaining performance for time-sensitive traffic, ensuring a high and reliable QoE for consumers.

Challenges

Traditional mobile networks are optimized for throughput but can struggle to maintain low and stable latency under load, especially in dense urban environments. For MasOrange, the challenge was to introduce advanced capabilities that would minimize delay and packet loss even during periods of congestion, without compromising overall network efficiency. Supporting interactive and immersive applications at scale required a more advanced approach to traffic management than what was possible with earlier generations of mobile technology.

Solutions and results

Ericsson enabled L4S as part of its 5G Advanced portfolio, allowing MasOrange to support time-critical applications with consistently low latency and reduced packet loss, even under varying network conditions. Together with other 5G Advanced features,

In Spain, 20 cities benefit from L4S-enabled 5G Advanced, which delivers more predictable low-latency experiences.

20

L4S is already benefiting consumers in 20 key Spanish cities, including Madrid and Barcelona, by improving the responsiveness of services such as real-time video, AR/VR and cloud gaming. By enabling more predictable performance over 5G SA, L4S helps MasOrange unlock new service possibilities and represents a key step toward delivering a fuller, more immersive 5G experience in Spain.

² Ericsson, "Spain steps into 5G Advanced with MasOrange and Ericsson" (February 16 2026).

Standardization, the ecosystem and Ericsson’s portfolio

The wide deployment of L4S in other access technologies is an enabler for the development of applications that support L4S.

L4S standardization was initially driven in IETF and has since been standardized in 3GPP and in DOCSIS. Standardization work is also ongoing in the Institute of Electrical and Electronics Engineers (IEEE) for Wi-Fi. In addition, 3GPP brings forward a path toward differentiated connectivity, in which L4S is an integral part.

As mentioned previously, L4S requires support both in networks and in applications. L4S support in 5G has been explained above. L4S requires end-to-end support in the mobile network and Ericsson supports that in its 5G RAN, Core and Transport portfolios.

With regard to L4S support in applications, the ecosystem continues to expand. As a result, any list of applications with L4S support risks quickly becoming outdated. Figure 6 highlights two well-known examples. A common characteristic among these applications is that L4S is typically introduced without extensive publicity. Instead, it functions as an underlying capability that enhances application performance when supported by the network. This is made possible because L4S does not require any specialized API signaling between application and

network, beyond setting the ECN ECT(1) codepoint in the IP header. Additionally, because transport protocols like TCP, QUIC and RTP/UDP have support for L4S, this means that applications built on these transport protocols can enable L4S.

Finally, CAMARA integrates L4S support into its global API framework, within the Quality-on-Demand (QoD) API family. This opens up further potential for service providers to adapt and capitalize on L4S’ features.

Ericsson’s L4S solution benefits

- Fully 3GPP-compliant implementation: L4S is implemented using a standardized 5G QoS framework, ensuring full alignment with 3GPP specifications and interoperability within 5G SA networks.
- QoS-based isolation of L4S traffic: L4S flows are mapped to dedicated QoS bearers, isolating latency-sensitive traffic from potentially bursty non-L4S traffic and preventing uncontrolled queue build-up under load.
- Fine-grained, latency-optimized traffic handling: Per packet identification and steering of L4S-capable traffic in the Packet Core enables precise

mapping to dedicated QoS flows and latency-optimized radio bearers. This granular approach delivers predictable low latency and represents a key differentiation compared with traditional, flow-based traffic management.

- Integration with differentiated connectivity: By leveraging the existing 5G QoS framework, L4S can be combined with additional QoS and scheduling features, making it an integral part of differentiated connectivity offerings rather than a standalone mechanism.
- Network-controlled uplink and downlink congestion signaling: L4S congestion marking is performed in the RAN for both downlink and uplink traffic, enabling consistent and reliable congestion signaling without requiring L4S support in UE.
- Service provider control of uplink behavior: Network-based uplink marking and QoS control provide service providers with the ability to selectively manage and enhance uplink performance for L4S traffic, which is critical for interactive and latency-sensitive applications.

Figure 6: Examples of well-known applications with L4S support

Application	Description	Links
Apple FaceTime	Peer-to-peer communication	Testing and Debugging L4S in Your App
WebRTC (libwebrtc)	Google Meet teleconferencing	L4S Interoperability Events

Summary

Stable and predictable performance will be vital as performance demands and network congestion increase. As SoftBank and MasOrange have shown, L4S is a key enabler of advanced capabilities.

L4S is gaining strong traction, with growing ecosystem support across applications and access technologies like 5G, DOCSIS and Wi-Fi. As networks move beyond best-effort connectivity, L4S is a key capability for differentiated connectivity in 5G as it ensures graceful performance degradation, rather than abrupt quality drops when capacity or coverage is constrained.

However, as the industry moves closer toward experience-based capabilities, and use cases become more advanced, the challenges for service providers and their networks is set to increase in lockstep. Not only will the demand for a stable and reliable QoE rise, the difficulty in securing it will grow too, presenting a persistent challenge. The ability of L4S to provide the elasticity needed to adapt quickly and control network congestion under varying conditions has already proven its value among the early adopters.

Ericsson's end-to-end portfolio supports L4S in 5G Advanced RAN and core networks. With the delivery of complementary capabilities such as radio scheduling and prioritization, it further reduces the delay impact on applications, that are caused by other factors like sudden changes in radio conditions. Ericsson's portfolio also brings competence in addressing L4S in the transport domain – positioning service providers to benefit fully from this capability and enable differentiated connectivity that meets the demands of next-generation applications and services. QoS mechanisms alone will not be enough to meet these demands, as without L4S, even prioritized users can be subject to performance degradation.

Implications for service providers:

- L4S is a critical capability for moving beyond best-effort connectivity toward experience-based differentiation in 5G SA networks.
- By combining L4S with 5G Advanced features such as latency priority scheduling and rate-controlled scheduling, service providers can deliver predictable low-latency performance for selected applications even under congestion.
- Investing in end-to-end L4S support across RAN, Core and Transport enables new premium service offerings for latency-sensitive use cases and creates opportunities for monetization based on performance guarantees, rather than raw bandwidth alone.

Implications for ASPs and developers:

- Enabling L4S in latency-sensitive applications unlocks more stable and predictable performance across both mobile and fixed access networks.
- L4S allows applications to adapt more gracefully to congestion, improving QoE without requiring tight coupling or proprietary interfaces with the network.
- As L4S support becomes more widespread across access technologies, early adoption provides a competitive advantage in delivering high-quality, real-time user experiences.
- L4S' congestion signaling provides ASPs with fast and precise feedback from the network, helping them adopt a more agile, reactive and efficient posture in the evolving ecosystem.

As 5G SA networks evolve beyond best-effort delivery, the industry is shifting toward connectivity defined by experience – where predictable latency, low packet loss and stable throughput matter as much as peak speed for latency-sensitive applications such as XR, cloud gaming and real-time video. As shown by SoftBank and MasOrange, and their partnerships with Ericsson, L4S is a key building block for enabling this shift. By introducing early, explicit congestion signaling and scalable congestion control, L4S allows networks and applications to work together to deliver consistently high user experience in real-world conditions.

With a growing ecosystem, Ericsson is making the transition today with proven live network deployment – preparing networks to support next-generation, latency-sensitive services and unlocking new opportunities for experience-based differentiation in 5G SA.

About Ericsson

Ericsson's high-performing networks provide connectivity for billions of people every day. For nearly 150 years, we've been pioneers in creating technology for communication. We offer mobile communication and connectivity solutions for service providers and enterprises. Together with our customers and partners, we make the digital world of tomorrow a reality.

www.ericsson.com