

# Intent-driven networks

# Contents

Introduction	3
Differentiated connectivity services	4
Intents and utility functions	6
Intent-driven autonomous networks	8
Intents in RAN	12
RAN application examples	14
Technology feasibility	16
Intent-driven operations	18
Industry alignment and standardization	20
Conclusion	22
Glossary	23
References	24
Further reading	25
Authors	26

# Introduction

“Intents” allow the expression of objectives that a system must achieve, without telling the system how to achieve them. Intents can be used across a communications service provider’s organization, allowing business objectives to be shared with network operations, where they may be autonomously translated into network resource-related actions, even in the radio network functions (NFs).

Transforming network operations to become intent driven is a key step in the journey to autonomous networks, where a gradual evolution is expected over time. Operating mobile networks with an increased use of AI and automation will enable service providers to control costs and address new business opportunities. It will also speed up the go-to-market process of new connectivity service offerings and reduce operational costs, in addition to improving both network performance and connectivity service performance. The architecture for intent-driven and autonomous networks follows industry standards such as TM Forum, 3GPP and Open RAN.

In a recent report on differentiated 5G connectivity<sup>[1]</sup> the monetization potential of 5G standalone (SA) was analyzed. Differentiated connectivity services beyond best-effort mobile broadband are anticipated to enable increased revenue growth and profitability, providing new value for users. Consequently, networks must be able to assure connectivity services in line with service providers' business objectives. Introducing multiple differentiated services sharing the same spectrum will make traditional mobile network operations increasingly more complex. This challenge must be addressed by integrating a higher degree of automation to maintain sustainable operational costs, leading to autonomous networks.

Ericsson is driving the industry to align around several types of connectivity services as a base for differentiation <sup>[1]</sup>.

# Differentiated connectivity services

For the industry to move beyond best-effort mobile broadband services, members of the application service provider (ASP) community need clear guidelines on what can be expected from differentiated connectivity and how their application behaviors fit with those guidelines. Network connectivity services define which connectivity KPIs to expect from a network connection, such as throughput and latency. The characteristics of the connectivity could further be described according to if it is suited for an application service flow that is buffered or not, with a fixed or adaptive bitrate. Differentiation between services could be maintained even in situations with a surplus of network capacity, for example, to save energy. In situations with capacity shortage, the differentiation should be guided by the business objectives of the service provider. The industry would benefit from aligning around a set of connectivity performance levels. One such proposal is presented in the differentiated connectivity report previously mentioned [\[1\]](#).

To create a full network connectivity service offering, a service provider must also specify targets for availability, geographical validity and other complementary characteristics of the offering.

The requirements on the network connectivity service infer requirements on sub-domains such as radio access network (RAN) and the adherent RAN connectivity service.

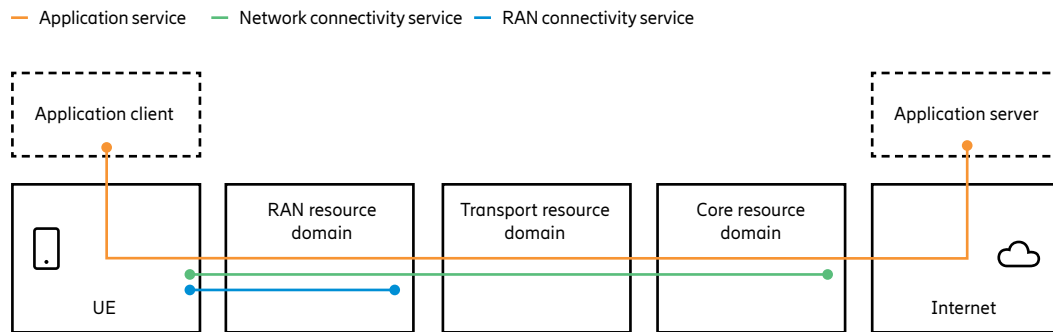


Figure 1: Application service versus connectivity services

Managing multiple differentiated services that share the same spectrum, while also balancing connectivity service performance with energy saving, will become a challenge. This would lead to unsustainable operational costs if a service provider continued to operate its mobile network in the traditional way. The thousands of individual parameters configured today will grow with the number of connectivity services and will require tuning per service. To manage the complexity, the network should be given objectives to achieve, rather than configuring the network for a wanted behavior. This can be achieved using intents.

# Intents and utility functions

“Intent” is defined slightly differently across standardization organizations, but the consensus is that intent is “declarative information about requirements.” This means that intent can be used to express the requirements that need to be met by a network, or by a subsystem within that network. Furthermore, intent excludes a description of how the requirements are supposed to be met. Decisions about actions to be taken or configurations to be applied in networks are made autonomously by the system that receives the intent. Thus, it provides a clear separation of concerns between subsystems involved in operations, and substantially alleviates the operational burden of manually configuring complex network parameter settings. Intents may define requirements for performance, availability, geographical and temporal validity and more. Requirements of connectivity service performance are expressed here by specifying thresholds based on performance metrics such as throughput, latency, or a combination thereof. Compliance with the requirements is therefore mostly binary – that is, the network is either compliant, or it is not.

While it is generally understood that improvements to the performance metrics are preferential, the autonomous network has no means to quantify the degree of preference. This is particularly interesting when multiple performance metrics are used in combination and the autonomous network needs to decide how much resource will be allocated for each metric’s individual improvement. The concept of “utility” helps the network to make the right decisions.

In autonomous networks, utility refers to the value or usefulness of a particular system state for a service provider and its customers. It is described as “knowledge about what makes an outcome or situation preferential” [2]. This information can be conveyed within intents in the form of utility functions. They associate every possible performance metric outcome with a utility score that rates the value of, and preference for, this outcome – reflecting the business priorities of the service provider.

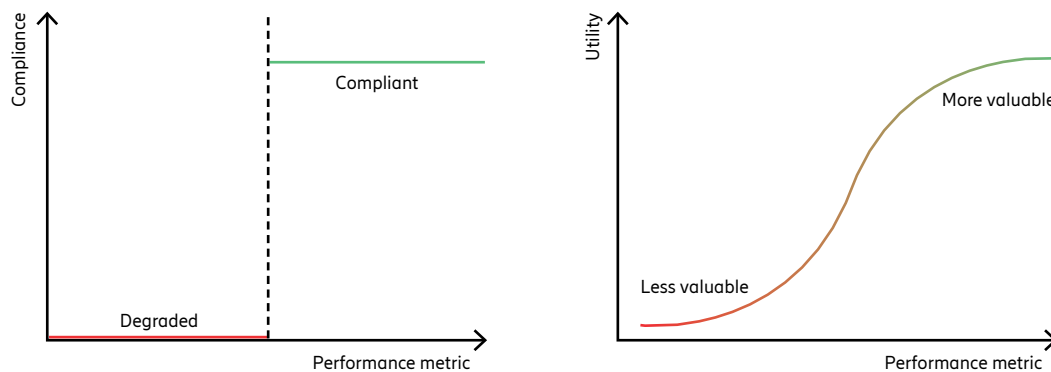


Figure 2: Compliance thresholds versus utility functions

Using this information, the autonomous network knows the relative value of its state. Furthermore, it can assess if an action would improve or degrade the value of the system state by scoring the expected outcome of the action being considered, and comparing this with the current state's utility. This directly facilitates optimization processes aimed at finding the actions that would maximize overall system utility, corresponding to the most valuable state achievable with the available resources and operational actions from a business perspective.

The same utility-based assessment can also balance conflicting requirements. An action that is improving certain performance metrics at the expense of others constitutes a conflict. Utility scoring tells the autonomous network whether the preferential effect of this action would outweigh its collateral degradation.

Utility information provides the means for analysis and decision making that is also highly adaptive to changes in system behavior and intent requirements. It does not rely on human-designed policy rules based on pre-defined situations. Instead, utility information enables a generic decision-finding algorithm capable of dealing with any situation. This reduces human involvement and enables the network to self-adapt to new intents and new services, thereby increasing the level of autonomy and reducing operational costs.

The shape and scale of utility functions are determined by the issuer of the intent. The issuer knows the business reason for choosing requirements and therefore it understands what business value a particular outcome would have. Utility functions carry this notion of business value into network operations. Human technicians have similar knowledge, allowing them to make the right decision even when facing new situations; for example, they know the service provider's strategy, goals and priorities set by management. With utility information, similar capabilities are provided in the operation of autonomous networks and enable business-value-aware decision making across the operational layers.

# Intent-driven autonomous networks

A network is operated by multiple systems and subsystems that might also be provided by multiple vendors. Conflicts will arise as concerns addressed across multiple subsystems are interconnected, and so operational actions will therefore lead to collateral effects. For example, optimization of energy consumption may impact other aspects, such as service performance. Resolution of such conflicts was previously based on human decision making and direct involvement in operations, or through human-curated policies. Autonomous networks need similar capabilities for controlling and managing conflicts, but without detailed human input. One solution to this challenge, which will encourage a gradual evolution from human-driven automation toward intent-driven and self-adaptive autonomy, is the organization of network operations with autonomous domains. An autonomous network which consists of multiple autonomous domains has the capability to self-adapt to changing strategies, requirements, traffic conditions and other situations.

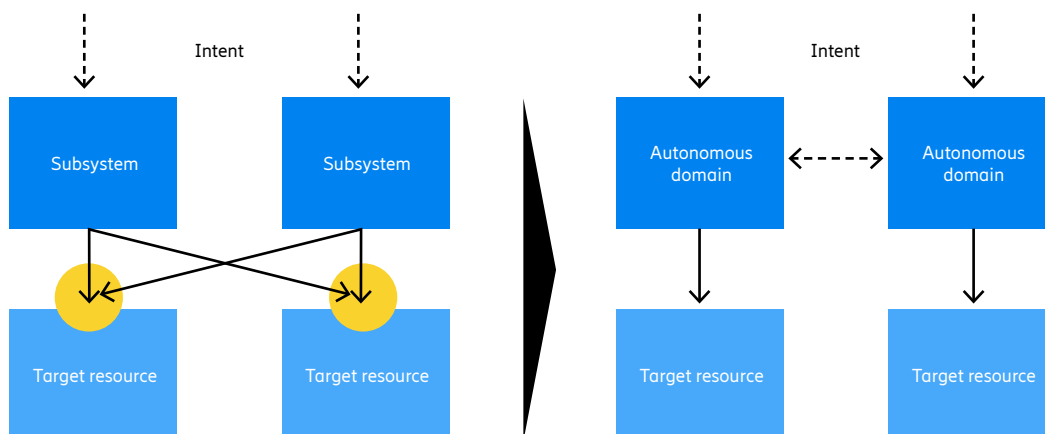


Figure 3: Conflicts inferred by uncoordinated subsystems (left) can be handled in an autonomous network (right)



An autonomous domain identifies a specific scope within an autonomous network. Autonomous domains are defined by describing a functional scope and assigning sets or types of resources. A subsystem being assigned to an autonomous domain implies exclusive and disjoint responsibility for performing certain actions concerning the included resources. This exclusivity ensures a clearly defined and automatically detectable allocation of authority, without interference from other subsystems.

An autonomous domain can be intent aware. This means that the autonomous domain would implement an intent management function (IMF) and interact with other autonomous domains by exchanging requirements. As an intent owner, the IMF sends intents to an intent handler, realized by another IMF, as shown in Figure 4. An IMF may act as both an intent owner and intent handler. If one IMF needs the resources controlled by another IMF to behave in a particular way, the first IMF can use intent to request that its requirements are considered. This means that instead of dealing with directly interfering and conflicting actions, subsystems cooperate through intent management. Features of the intent models and interfaces, such as requirement negotiation, decisions by utility maximization and intent reporting are available for autonomous conflict resolution. Autonomous domains are found on all operation layers in the network, as depicted in Figure 5, interacting within and across the layers.

An IMF is responsible for handling intents within its autonomous domain. As an intent handler, the IMF implements monitoring of the system state and the translation of intent requirements into solution strategies for actions to achieve compliance with the received intents. An action may be the direct deployment and configuration of services and resources. Furthermore, an IMF can determine what behavior it needs from other autonomous domains to comply to its intent. As an intent owner, an IMF can act by expressing these requirements with intents sent to the respective IMF for the addressed autonomous domains. This process of translating received intents, including utility functions and intent forwarding is a key principle in end-to-end, intent-based operations.

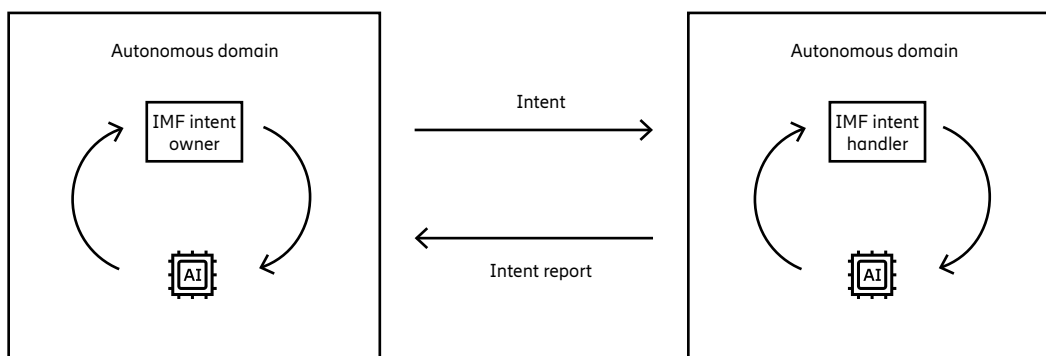


Figure 4: IMF interaction between AI-driven autonomous domains

The realization of an IMF can vary significantly depending on the nature and scope of its autonomous domain. For example, the intent management in the business layer differs significantly from intent management in the resource layer.

The implementation of an IMF might reuse existing policies, thus preserving investment. However, higher degrees of autonomy would inevitably require a combination of AI techniques taking roles in the operational processes and control loops. For example: models for anomaly detection can identify issues; prediction models enable proactive operation by finding potential issues ahead of time; and prediction capabilities embedded in digital twins can help with safe exploration of possible solutions.

AI can drive the control loops needed in intent handling. For example, multi-agent reinforcement learning refers to models collaboratively and iteratively optimizing operational parameters. The intent is, in this respect, determining and dynamically managing the objectives of the process and the reward used in the machine-learning process.

Furthermore, generative techniques, such as generative adversarial networks or transformer-based approaches, have the potential to create new solutions and find optimized configurations, even if the network is facing previously unexplored situations. Large language models can provide an intuitive experience for users in interactions with the network.

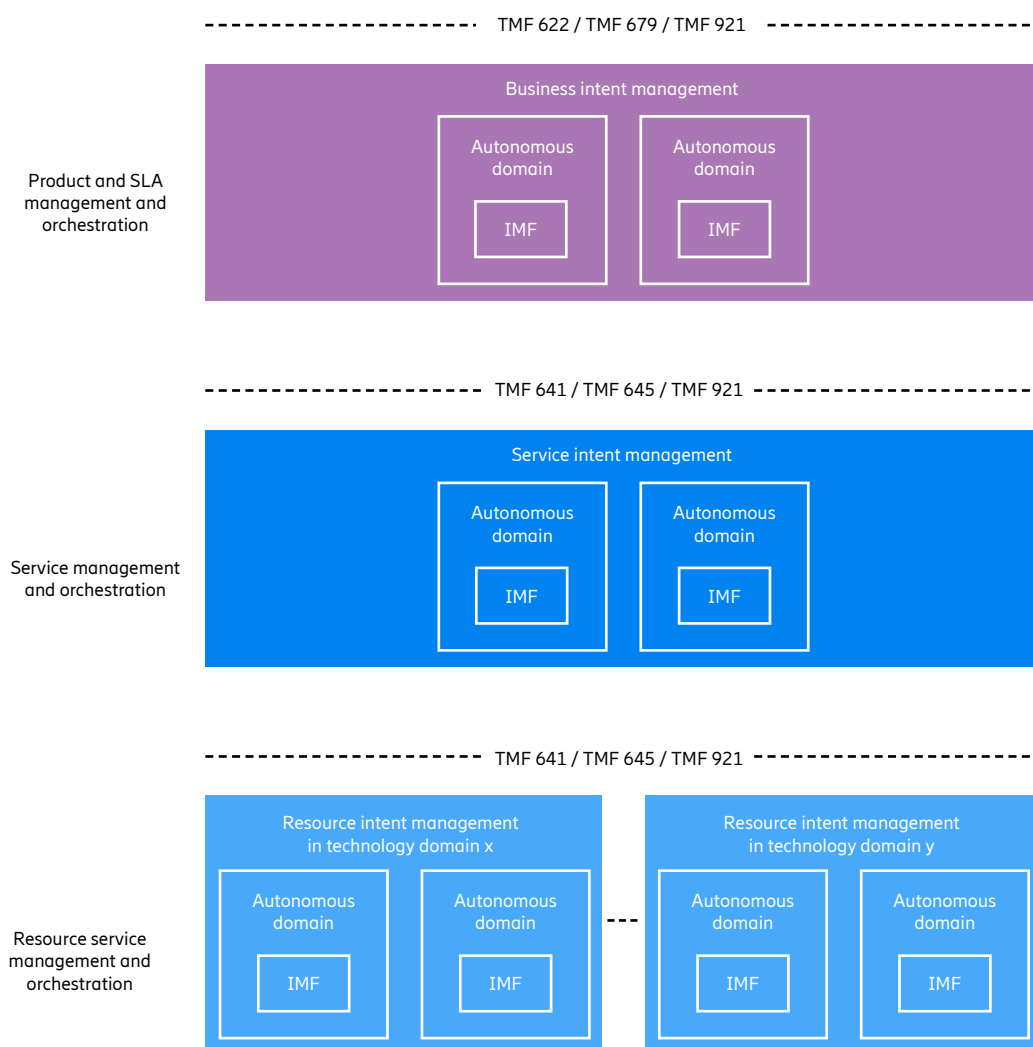


Figure 5: Service provider operations stack for autonomous networks enablement covering business, service and resource operations layers

The network operations may be described with three layers – business operations, service operations and resource operations [2] – interacting through open and standardized interfaces, as indicated in Figure 5 and Figure 6. Product and service level agreement (SLA) management and orchestration at the business operations layer receives commercial orders from customers as product intents, for example relating to a network connectivity service offering. The service provider here defines the value of fulfilling the requirements in the product intent as a utility function.

The product intent is decomposed to a customer-facing service (CFS) intent and forwarded to an IMF at the service operations layer according to its autonomous domain. At the service operations layer, the CFS intents are further decomposed into resource facing service (RFS) intents, driving the deployment and configuration of the network resources to meet the requested service requirements for the related technology domains, for example RAN, Core and transport.

As intent handlers, the IMFs at the resource operations layer will receive intents according to their defined autonomous domains and may deploy capabilities like network slicing and quality of service (QoS) to realize the network connectivity service. Throughout this chain of intent decomposition, the requirements and utility function are translated to reflect the metrics available in respective autonomous domains.

IMFs will be deployed at all these different operations layers, working with different scopes and on different time scales. As an intent handler, it will be responsible for the service assurance with the objective to fulfill the intents in an optimal way through closed-loop actions relevant for the operations layer. For example, an intent handler in the service operations layer will take care of the service intent, ensuring that all parts of the network (RAN, transport and core) contribute to the overall intents. This assures the network connectivity service performance, while an intent handler in the RAN resource operations layer will orchestrate multiple radio nodes over an area and will therefore have the overall responsibility to assure the RAN connectivity service over a larger RAN area.

The associated granular observability offers the IMF and service provider real-time awareness about how different services perform and how intents are met on each operations layer. The observability per connectivity service, slice, per user equipment (UE) and geography also help control any potential financial risks associated with missing the target performance levels for the connectivity services.

# Intents in RAN

Intents in the RAN resource operations layer enter at the RAN resource service management and orchestration (SMO) level, and will be handled by the RAN IMF of the associated autonomous RAN domain. In an Open RAN, SMO-based solution, RAN IMFs may be realized as rApps. Depending on the capabilities of the RAN NFs, the RAN IMF decides whether to process and forward a decomposed intent to the autonomous domain defined by the RAN NF, or to realize the requirements in the intent through traditional configuration of RAN NF features. Forwarding the intent implies that the RAN NF will take actions to assure the intent within its autonomous domain.

While the intent handler in the RAN management level will orchestrate multiple radio nodes to assure the RAN service over a larger RAN area, the intent handler in the RAN NF will be responsible for its resources and will be able to take actions related to the associated coverage area.

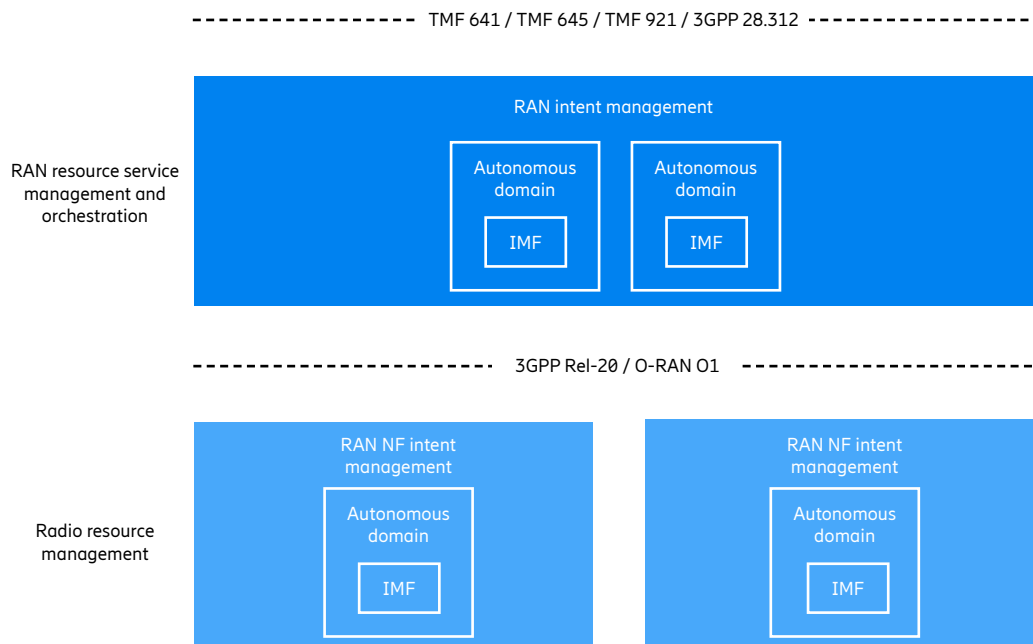


Figure 6: Service provider operations stack for autonomous RAN enablement on the resource operations layer

The management interface to the RAN NF, enhanced with intent handling capabilities, may be used to propagate intents to the RAN NF, defining its own autonomous domain. Intents to the RAN NF related to connectivity are associated with identifiers – like 5G QoS identifier (5QI) and single-network slice selection assistance information (S-NSSAI), to identify the RAN connectivity service flows in the scope of the intent. The utility functions provided with the intents are used to steer the radio resource management (RRM) to maximize the utility aggregated across RAN connectivity service flows in the RAN NF. The business value of the service provider is therefore ultimately used to guide the traffic treatment in the RAN NF.

Instead of manually activating energy-saving modes and policies, intents can introduce the concern to save energy. These intents may be defined by the service provider on the RAN resource operations layer as a resource-related intent to the RAN IMF. This allows the use of associated utility functions defining the energy saving preference versus RAN connectivity service performance. Additionally, depending on the capabilities of the RAN NF, the energy performance intents may be forwarded to the RAN NF by the RAN IMF, or the RAN NF may be configured to meet the energy requirements in the intent.

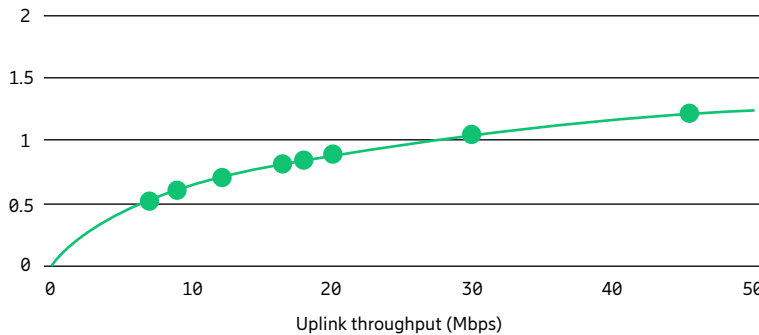
In the RAN NF, functions like beamforming, link adaptation, radio scheduling, carrier aggregation, mobility and Low Latency, Low Loss, and Scalable throughput (L4S) can become utility driven. The utility, derived from performance metrics, will then guide the network's decisions to ensure resource utilization according to the service providers' business objectives.

It is envisioned that there will be a gradual evolution toward a fully intent-driven RAN. Even then, there may still be a need for some level of configuration of the RAN NF, for example, related to initial deployment. Intent management will therefore co-exist with configuration management in the RAN NF, requiring less configuration effort.

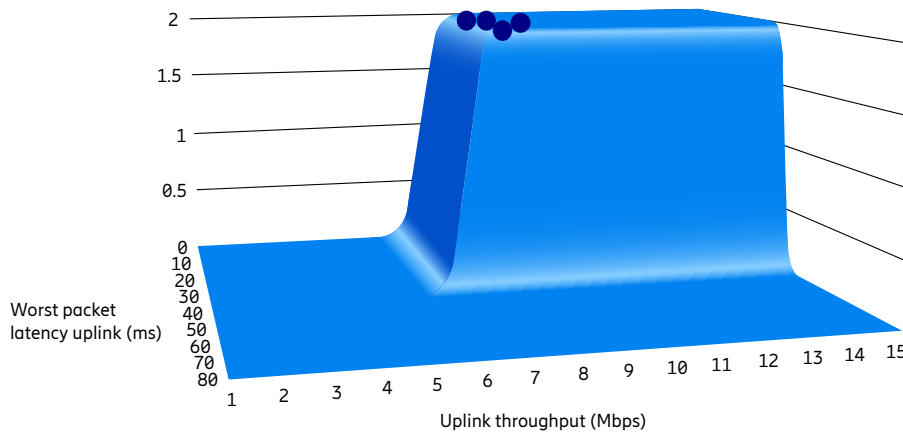
# RAN application examples

A service provider's product offering that includes a real-time connectivity service to or from devices can be used to illustrate how intents can be used in the RAN. This service provider product can be used for tele-driving of cars (video feed from camera in the car), cloud gaming (video feed from server to console), extended reality and more. The intent requirements in this example are formulated as 7 Mbps minimum uplink throughput and a bounded one-way uplink packet latency of 50 ms for the network connectivity service. Percolating down the hierarchies of Figure 5 and Figure 6, the associated RAN intent reads "7 Mbps minimum uplink throughput and a bounded RAN one-way uplink latency of 40 ms for traffic associated with S-NSSAI = 50," where RAN has been allocated 40 ms of the total latency budget of 50 ms. Here, the S-NSSAI has been used as the connectivity service identifier, with the assumption that the 5G Core is configured to mark traffic associated with the real-time connectivity service with that S-NSSAI, either by means of subscription identity, user equipment route selection policy (URSP) signaling or a dynamic application programming interface (API) call. In parallel, the service provider offers a default connectivity service suitable for mobile broadband traffic on another S-NSSAI with performance levels and intent associated with uplink throughput levels, but no latency requirements. To further refine the example, consider the scenario of an area of a mobile network in which the service provider offers a mobile broadband product and an enterprise product for tele-driving of cars. The mobile broadband product is realized by a default connectivity service with throughput intents, and the enterprise product is realized by a real-time connectivity service with throughput/latency intents as above. Based on business considerations, the service provider assigns a relative priority to the real-time connectivity service and the default connectivity service encoded in utility functions, as illustrated in Figure 7.

(a) Utility for mobile broadband



(b) Utility for mobile broadband



(c) Total utility of the system is the sum of the individual utilities.

$$U_{\text{Total}}(\vec{R}, \vec{L}) = \sum_{i=1}^N U_i(\vec{R}_i, \vec{L}_i)$$

Figure 7: Utility functions for mobile broadband (a) and tele-driving (b), with total utility (c) indicated as a sum over the individual connections

In line with the declarative approach of intents, there are many different RAN NF implementations possible where the RRM algorithms (such as algorithms for mobility, frequency selection, radio scheduling and link adaptation) are designed to gauge towards the intents with priorities given by the utility function. One set of algorithms that has been demonstrated to work in lab and field settings focuses on the radio scheduler. In this implementation, the IETF protocol L4S [3] is used to control the bounded latency of the video service. The throughput levels to the individual users are controlled in the scheduler, either directly with scheduling priorities or indirectly via flow-control (throughput control) mechanisms of L4S for the video service, and transmission control protocol (TCP) for the mobile broadband flows. The last piece of the puzzle is the scheduler’s knowledge of the individual UE’s radio conditions that translate spectrum used into throughput provided. With the bounded latency intent under control by means of L4S, the scheduler can focus on its dynamic scheduling priorities and flow control signaling to control the throughput levels to maximize the total utility function at any given moment of time, both during times of congestion and during times of surplus resources.

# Technology feasibility

The ability for a RAN to steer toward intents and utility functions was recently demonstrated in both lab and live network environments. Figure 8 shows a test run in a lab system running on commercial product code using the principles above. In the test setup, there is a real-time connectivity service with throughput-latency intent of the structure above, but with the more aggressive throughput intent of 60 Mbps and bounded one-way latency of 40 ms. Based on what would correspond to a service provider's business decision, the utility function used in the test run is shaped in a way to increase the real-time connectivity service bitrate above 60 Mbps if there is capacity headroom to do so. In the case of tele-driving cars, the rationale for this provisioning of extra bandwidth would be, for example, to enable the car to activate a second camera to improve the user experience whenever this comes without consequences to other users in the area.

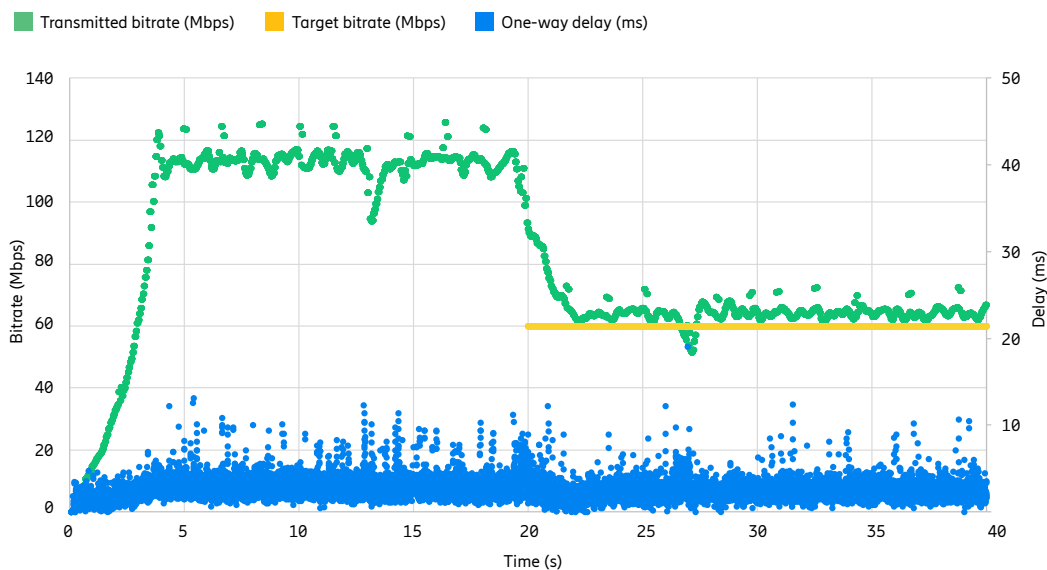


Figure 8: System logs following the intents for the real-time connectivity service



The blue dots in Figure 8 represent measurement of the one-way bounded latency intent. The figure shows that this is consistently below 40 ms (below 20 ms most of the time). For the first 20 s, the real-time connectivity service UE is the only UE in the system and the utility function is maximized by giving that UE the full cell capacity of 130 Mbps. At  $T = 20$  s, a mobile broadband user enters the system. Based on the utility functions, the system calculates the target bitrates for the two UEs that maximize the total utility. With the utility functions used in the test run, the maximum occurs at a target bitrate of 60 Mbps for the real-time connectivity service UE. Consequently, the system adjusts to a new target bitrate of 60 Mbps for the UE and uses L4S-packet markings to regulate down the data stream to fit on the 60 Mbps with no latency spikes. Finally, the green lines measure the momentaneous data rate of the application and show how the video application adjusts well to the assigned radio bandwidth, due to the L4S control loop.

# Intent-driven operations

Intent-driven operations are an integral part of the already-established processes for service providers' operations. These processes can be roughly divided into two categories, as shown in Figure 9. Processes for strategy, planning and product definition determine how the service providers would set up their business and approach customers. This is accompanied by processes for operation and support, which execute and automate the delivery of value to customers.

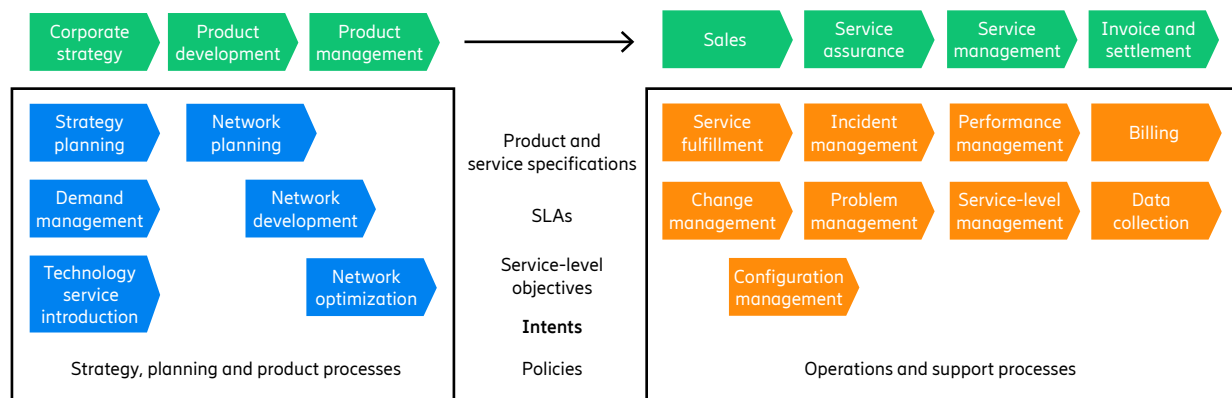


Figure 9: Intents as key input from strategic planning into operations and support processes

Within this framework of processes, artifacts created in strategy and planning are the input and guidance for operations. This includes examples such as product and service specifications in combination with SLAs. They represent orderable business objects exposing the network capabilities to customers. Service-level objectives manage detailed service requirements and define what automation needs to achieve and assure. Policies allow detailed control of operational decisions through prescribed decision and action-taking procedures.

In this environment, intent introduces common objectives across processes. It informs the operation and support processes about their requirements originating from strategic considerations and customer needs. Features of intent management, such as utility and requirement negotiation, enable dynamic operations and adaptivity to new conditions. It reduces the lead time typically caused by adaptation of more static artifacts used in process automation, such as policies. It also raises the autonomous decision-making capabilities implemented in the processes. As a result, the processes become less human dependent and more self-adaptive when managing conflicts, employing collaborative and autonomous decision making with clear separation of concerns and allocation of responsibilities.

Intent-based operations do not change the fundamental processes, but gradually evolve their implementation by adding capabilities in a backward-compatible way. This means the trust gained by current automation concepts will evolve rather than become disrupted. However, detailed decisions that are a concern for manual planning processes today can, in the future, be transitioned into automated operation processes in which autonomous systems consider them when choosing the configuration of the network. This will enable continuous, self-adaptive optimization while keeping low-level decisions out of strategic planning processes.

# Industry alignment and standardization

Intent-driven networks are a paradigm in the scope of many standards development organizations (SDOs), such as TM Forum, 3GPP and the O-RAN Alliance. The declarative and requirements-focused way of managing networks is identified as a solution to operate the increasingly complex networks of the future, where an intent interface may have the opportunity of being truly multi-vendor, avoiding unnecessary exposure of internal product systemization.

The Autonomous Networks project in TM Forum [\[2\]](#) defines a generic reference architecture of autonomous networks. It introduces autonomous domains with a clear allocation of responsibility to sub-systems and processes. TM Forum has defined the IMF and its roles within the autonomous domains, including their interaction through exchanging intent. In this respect, the TM Forum intent ontology (TIO) provides the vocabulary for expressing intent with an extensible information model that is not constrained to a single technology domain.

Furthermore, TM Forum introduces intent-management capabilities across a range of APIs. Most notably, intent can be associated with product and service orders, exposed through catalogs, and used to determine decisions in fulfillment and assurance processes. These APIs are already widely used in service automation today. Introducing intent in the established APIs adds capabilities to established processes in a backward-compatible way. As such, it facilitates, but does not enforce, a transition from automation to intent-driven autonomy. It allows this path to be taken on a per-use-case basis when there is a clear business benefit.

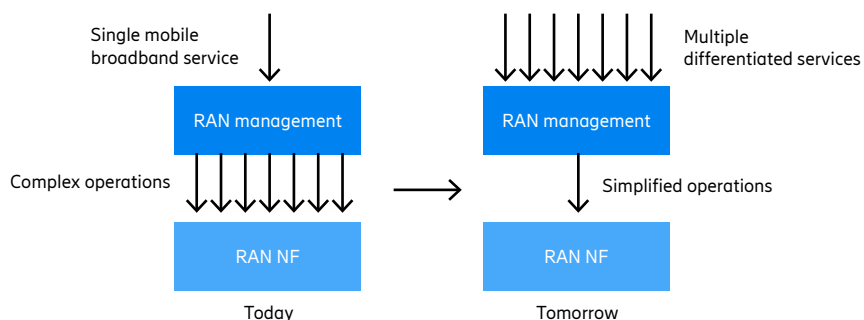


Figure 10: 3GPP standardization for simplified operations

In 3GPP, Rel-19 intent models and procedures on the RAN-management level are being defined [4], still requiring the behavior of the NFs to be configured by detailed parameter settings. The O-RAN Alliance follows the definition of an intent API to the SMO.

3GPP Rel-20 provides an opportunity to also define an intent interface to the RAN NFs for 6G. This would allow simplified operations of the RAN NF and a higher level of automation both on the management layer and in the RAN NF, supporting a multiplicity of differentiated connectivity services.

Ericsson is driving standardization activities in TM Forum, 3GPP and Open RAN that need to be embraced to ensure an open ecosystem for autonomous networks. Ericsson is leading efforts to align the industry around programmable networks and the support for differentiated connectivity [1]. The Ericsson initiative for a service provider partnership to combine and sell network APIs on a global scale is a further step to enable the next level of digitalized, high-performance and programmable networks [5].

# Conclusion

Intent-driven networks are a key enabler for the journey to open, programmable, autonomous networks. It will facilitate profitable service provider connectivity offerings, based on differentiated connectivity beyond best-effort mobile broadband. In an autonomous network, intents are applied on all operations layers, from the business support system to the NF, guided by the business value for the service provider through utility functions.

The autonomous network has the capability to self-adapt to changing strategies and requirements given by intents, and to varying traffic conditions and situations. Governed by an IMF, it may be realized by AI technologies, real-time closed loops, open APIs and metrics exposure to assure connectivity service experience.

The use of autonomous domains provides a clear separation of concerns, introduces modularity, facilitates integration and allows multi-vendor deployments by alleviating the need for conflict mitigation between the domains.

An intent-driven network does not change the operational processes needed to run the service provider enterprise, but provides common objectives across processes facilitating operations.

The industry, including ASPs, service providers, device vendors, smartphone platform vendors, API aggregators and more, will benefit from realizing the industry vision through differentiated connectivity. However, willingness from all parties to put in efforts and investments will be required to achieve this vision. An aligned vision would influence the adoption by standards-development organizations, ensuring an open ecosystem. To achieve the vision of autonomous networks, service providers must embrace intent-driven operations for differentiated connectivity. This will allow sustainable operations costs when expanding the service offering beyond best effort.

Ericsson's position is that intent-driven networks are fundamental for service providers to unlock the business potential of new products based on services beyond best effort. Intent-driven networks are required to shorten the time to market for these products by removing the operational complexity of configuring thousands of RAN-feature parameters to optimize performance across these multiple new services within the same spectrum. As an example, we have demonstrated that an intent-driven RAN solution can retain a bounded latency for an adaptive video service during varying traffic load conditions by including latency in the utility function. Our expectation is that the first intent-driven solutions will reach the market in 2025–2027.

# Glossary

3GPP	Third Generation Partnership Program
AI	Artificial intelligence
API	Application programmers interface
ASP	Application service provider
CFS	Customer-facing service
IMF	Intent management function
L4S	Low Latency, Low Loss, and Scalable throughput
S-NSSAI	Single-network slice selection assistance information
RAN	Radio Access Network
RFS	Resource facing service
SLA	Service level agreement
SMO	Service management and orchestration
TIO	TM Forum intent ontology
TMF	Telecom Management Forum, TM Forum

# References

1. [Capture the value of differentiated 5G connectivity - Ericsson](#)
2. [Intent in Autonomous Networks - TM Forum Introductory Guide IG1253](#)
3. [RFC 9330 - Low Latency, Low Loss, and Scalable Throughput \(L4S\) Internet Service: Architecture](#)
4. [3GPP TS 28.312 - Management and orchestration; Intent driven management services for mobile networks](#)
5. [Global telecom leaders join forces to redefine the industry with network APIs](#)



# Further reading

1. [Driving 5G monetization through intent-based network operations – Ericsson](#)
2. [Enterprise Service Orchestration – Ericsson](#)
3. [Cognitive framework for better 5G network LCM – Ericsson](#)
4. [Service management and intent-driven autonomous networks – Ericsson](#)
5. [Autonomous networks with multi-layer, intent-based operation – Ericsson](#)
6. [Translating learned data into business intent – Ericsson](#)

# Authors



**Paul Stjernholm** joined Ericsson in 1995 and has been dedicated to the mobile telecom industry since 1G. In his current role as a concepts researcher, his work focuses on RAN management strategies, automation and standardization. Stjernholm holds an M.Sc. in applied physics and electrical engineering from Linköping University, Sweden.



**Jörg Niemöller** is an analytics and customer experience expert in Solution Area Cognitive Network Solutions. He joined Ericsson in 1998 and spent several years at Ericsson Research, gaining experience in machine-reasoning technologies and knowledge modeling. He is currently leading the introduction of autonomous networks and standardization of intent-driven autonomous operation. Niemöller holds a Ph.D. in computer science from Tilburg University, the Netherlands, and a diploma degree in electrical engineering from the TU Dortmund University, Germany.



**Ann-Christine Eriksson** is a senior specialist in RAN and service layer interaction at Business Area Networks. She joined Ericsson in 1988 and has since worked with research and development of mobile networks. In her current role, she focuses on defining the evolution of RAN architecture for management and intent-driven autonomous networks. Eriksson holds an M.Sc. in physical engineering from KTH Royal Institute of Technology in Stockholm, Sweden.



**Mathias Sintorn** is an expert in traffic handling and service performance at Business Area Networks. He joined Ericsson in 1998 and has since worked intensively with system development, standardization and strategy for 3G, 4G and 5G RAN as well as managed services of 3G (including the rollout of High-Speed Downlink Packet Access). In his current role, he defines the long-term evolution of the RAN architecture, specifically in the area of RAN automation. Sintorn holds an M.Sc. in engineering physics from Uppsala University.



**Mohammed Ziauddin Mujale** is a Master Engineer in Technology Evolution and Architecture Authority for the BSS/OSS area. He joined Ericsson in 2011 and has since worked intensively on systems and architectures for OSS managing the mobile networks. His specialization is in service management and network slicing in mobile networks. Previously, he has worked as Product Manager for service management products and has the customer perspective on the OSS domain. Mujale holds an Engineering degree from Bangalore University and has specialized in systems engineering and management at the Indian Institute of Science (IISc), Bengaluru.



**Erik Westerberg** is a Senior Expert and the Chief network architect for Ericsson's RAN systems. Erik has a Ph.D. in physics from Stockholm University and joined Ericsson from the Massachusetts Institute of Technology in 1996. During his 28 years with Ericsson, Westerberg has contributed to four generations of mobile systems in standards, product systemization and in strategic network evolution projects together with service providers. In his present role, Westerberg oversees the 5G and 6G architecture with special focus on automation and intents.