Ericsson White Paper
BMAS-21:031168 Uen
October 2021

ERICSSON

# Driving 5G monetization through intent-based network operations

# Introduction

While communications service providers around the world will soon be racing to monetize on their 5G investments, they will nonetheless face a significant challenge in operating such complex networks and ecosystems. In fact, overcoming this may be the key differentiator between reaping the promise of 5G and falling short on meeting their business objectives.

Our 'AI and Automation in Operations Report' provides insight into how operations have become a key driver of business strategy; it is a well-known fact that service providers need to transform their operations. Not only are 5G enterprise customers' requirements the key drivers for this transformation, but service providers' own business requirements are equally important. Many processes and tasks are still being performed manually and will all need to be fully automated.

AI and automation are key technologies to overcome some of the main challenges associated with operations for 5G and beyond. This paper talks about the unique challenges in 5G network operations, explains how some of the traditional approaches may be cost prohibitive and have difficulty scaling, and explores how an intent-based approach can help overcome these challenges to deliver an agile, automated and cost-effective way to manage 5G network operations. Adoption of an intent-based approach will also aid service providers' transformation from traditional telecom operators to digital service providers.

**Contents**

# Opportunities and challenges

Industry digitization provides service providers with new options to grow and expand their businesses beyond connectivity. The 5G-IoT landscape has immense potential, but it is difficult to traverse. According to Ericsson's '5G for business: a 2030 market compass' report, the aggregate opportunity in the enterprise arena is estimated to exceed USD 700 billion by 2030.

**Figure 1: Global service creator, 5G for business addressable market, 2020–2030 CAGR**

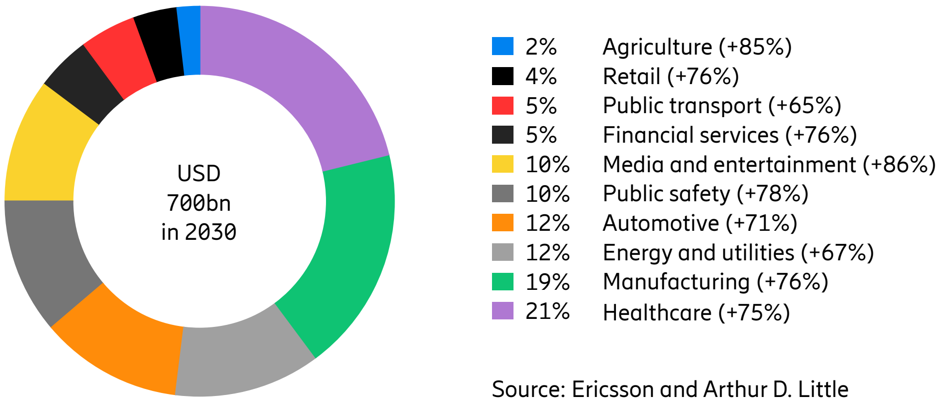


Source: Ericsson and Arthur D. Little

As shown in Ericsson ConsumerLab's report 'Harnessing the 5G consumer potential', the revenue opportunity in the consumer space is also clearly visible. The research predicts that over the next decade service providers could earn USD 131 billion from 5G use cases and digital service cases and could gain a 34 percent higher 5G average revenue per user (ARPU).

5G will not only bring a plethora of opportunities for service providers in terms of use cases, innovative business models, B2B/enterprise customers and more, but also it's own unique challenges as well. How service providers react and overcome challenges such as complexity of network deployments, monetizing 5G investments, adoption of AI and automation and network operation transformations will differentiate the leaders from the followers.
Some service providers have the notion that the cost of overcoming these challenges by investing in innovative technologies is high and the results are unclear. This is compounded by the lack of a strong business case and a clear roadmap for the networks. However, early movers will have a clear advantage over competitors in their transition towards future zero-touch networks. Therefore, service providers need to evaluate new business models and use cases, carefully considering their value chain positioning to maximize opportunities.

Due to its ultra-fast connectivity, flexibility and reliability, 5G will allow enterprises to use cellular connectivity as their primary access technology, making 5G one of the biggest platforms for innovation in enterprise, just like 4G was for the individual consumers. Hence with 5G, not only is the operational transformation of service providers' public networks important, but supporting numerous private and enterprise networks, along with institutional and public sector transformation programs, remains a compelling use case driving revenues.

While good connectivity and coverage remain the most important characteristics for service providers' public networks, there are stringent service-level agreements (SLAs) in place for private and enterprise networks, with far more important characteristics than simply connectivity and coverage. Operating such a large number of smaller – albeit strict – SLA-based networks require a huge transformational shift for service providers. Enterprise networks differ a lot from service providers' public networks with extremely high benchmarks for some of the following characteristics:

- reliability
- availability (of the order of five nines)
- security
- integrity (latency, throughput and so on.)

Industry verticals such as airports, oil and gas, ports, mining, healthcare, Industrial internet of things (IIoT) and public safety all require millisecond decision-making and full closed-loop automated procedures for their efficient operations. Imagine handling such networks via traditional operations approaches – for example, alarms handling, ticketing and resolution of trouble tickets, work order distribution and management – with less than 20 percent of the processes being automated. Existing network operations involving traditional AI approaches will be unable to handle the complexity and scale of industrial and enterprise 5G operations.

Another challenge in operating enterprise 5G networks is the complexity of underlying use cases. Enhanced mobile broadband (eMBB) and fixed wireless access (FWA) are two main use cases for service providers' public networks. But for service providers' enterprise networks, use cases such as ultra-reliable low-latency communication (URLLC), IIoT, network slicing and software-defined wide area network (SD-WAN) remain key, with stringent SLAs for each.

Different services and use cases have different SLAs, requirements, KPI benchmarks and so on, so a one-size-fits-all "network operations" strategy will not work with 5G operations, especially for the enterprise segment. It is almost impossible to automate all these use cases with traditional AI and network operations methodologies.

5G adoption is far more rapid than in previous generations and is likely to break the 20 percent level for share of mobile subscriptions 2 years faster than 4G (Harnessing the 5G consumer potential). By the end of 2021, 5G subscriptions are expected to reach 580 million (Ericsson Mobility Report June 2021), and by 2026, it is estimated that 5G networks will carry more than half of the world's smartphone traffic. This, coupled with the announcement last year that China plans to invest at least USD 1.4 trillion into building out 5G infrastructure over the next 5 years, shows the rapidly increasing pace of 5G adoption. It also highlights the urgent need for service providers to ramp up their operations and investment to not only meet the predicted demands, but to anticipate future trends and react accordingly.

# The changing role of operations

Industry digitization provides service providers with new options to grow and expand their businesses beyond connectivity. With the growing focus on customer experience over the last few years, operations have become central to service providers' strategies, as indicated by our 'Ericsson Technology Review report'. During the last few years there has been a paradigm shift from a "network-centric" approach to a "customer- and service-centric" approach, resulting in transformation across the technology, tools and processes in which operations are performed.

Network operations are no longer seen as a support function within service providers' organizations but are closely tied to delivering business goals and ensuring high customer experience. There is no doubt that the network's performance is central to overall network operations strategy and that remains one of the operation's top priorities. However, operations are no longer seen as just helping to keep the network running. There is so much more to operations as part of service providers' overall strategy as we move beyond 5G operations. Below are some of the key operation roles and their importance in the context of overall service providers' business strategies:

1. Delivering a high customer experience – customer experience remains the top business strategy for service providers. Operations have a key role to play in realizing this strategic priority of the wider business.
2. Helping service providers in capturing new businesses – with 5G, service providers have a chance to venture into new businesses, industry verticals and use cases, for example, mining, industrial automation, airports and ports, to name just a few. Operations will be at the forefront of innovation if service providers are to succeed in these new businesses.
3. Securing new business models and driving revenue growth – private/enterprise networks have stringent SLAs, enhanced resiliency requirements and full closed-loop automation. Operations will be a key differentiator in driving revenue growth in these critical areas for service providers, helping them secure new business models.
4. Delivering high network performance – a customer's perception about the network is mainly driven by network performance, with at least one-third of Net Promoter Score (NPS) derived from it. Operations teams are not only responsible for improving quality of service and overall efficiency, but also work closely with cross-functional units such as marketing and sales to further improve the NPS.
5. Delivering on overall service providers' cost-efficiency programs – operations are associated with service providers' overall cost-efficiency programs. However, cost efficiency against a tide of increasing complexity is not something which can be achieved quickly. It requires constant effort to eliminate repeated and costly manual processes, implement AI-based use cases and minimize faults.

Operations will continue to play a key role in the adoption of AI and automation in telecom networks. Service providers' customer experience improvement, as well as overall business strategies, are closely linked to network operations. In fact, how quickly and effectively service providers can monetize their 5G investments will mainly be driven by the innovations in operations.

# How AI has transformed operations

Telecom network operations have long been reactive and were served by network operations centers (NOCs) which used to act on an alarm or anomaly in the network. These kinds of operations were aptly described as reactive NOC operations, as an NOC engineer would react once an alarm was seen in the network and then act upon it accordingly. The bulk of operations processes and actions were performed manually.

In the next phase, some of the manual processes were automated, resulting in semi-manual operations. Aided by thousands of automation rules, plenty of operational optimizations were achieved both in back- and front-office organizations, thereby improving efficiency. Reuse of these automation rules drastically reduced the time to market for new use cases.

Gradually, with the introduction of AI in operations, these reactive, semi-manual NOCs were transformed to predictive operations where network violations were anticipated and resolved before they happened. Figure 2 compares a reactive and alarm-driven workflow to a data-driven predictive workflow using AI and machine learning (ML).

**Figure 2: Typical reactive and proactive operations workflow**

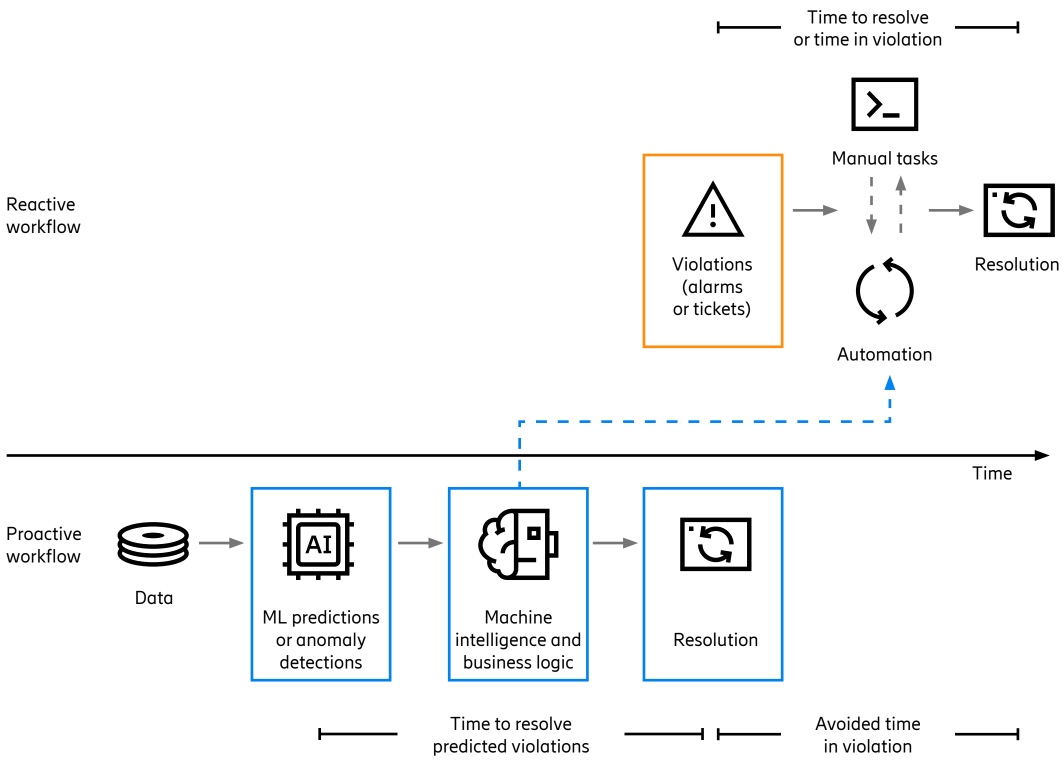

A typical reactive workflow starts with operations receiving an alarm or trouble ticket (for example, a customer complaint). Operations then initiate troubleshooting – which is often a combination of semi-manual and automated tasks – before resolutions are found and executed. The time it takes to resolve an issue also represents the time the issue had been in violation.

Based on the prediction, AI-driven tasks such as pattern matching and automated root cause analysis are initiated to remedy the problem. In some cases, fallouts can occur where AI cannot adequately decide and must hand off manual investigations. Nonetheless, with well-designed use cases, the vast majority of the issues can be resolved.

With AI-driven proactive workflows, not only can potential issues be detected ahead of time, but most of them can also be resolved before actual violations occur. This not only reduces network outages and violations but helps reduce operational cost through AI-driven automations.

Here, AI boosted the already-developed automation capabilities resulting in a huge reduction in Priority-1 incidents and improved NPS scores. Semi-manual operations were changed to data-driven operations and this first phase of AI adoption in operations has resulted in the aforementioned transformations.

By using AI to predict, prevent and resolve events, we can reduce opex despite having to manage increased complexity. Plenty of innovative use cases were developed augmenting human capabilities and delivering results that had never been seen before. Some of the most important use cases in different operations categories are summarized below:

| 22% | 21% | 29% |
|---|---|---|
| Service uptime improved by 22 percent | Field visits reduced by 21 percent | Customer complaint index improved by 29 percent |

Source: Observed results for a Tier 1 service provider in Asia

Early AI systems work on finding patterns in large data sets. If there are few or no structured inputs to find patterns, ML systems can't solve a new problem as there will be no apparent relation to its prior knowledge. The limitations posed by ML modeling are solved by applying machine reasoning to our systems.

Also, with predictive operations, we often require manual intervention to solve them and hence they may not always be self-healing. However, as we transition towards 5G networks that many enterprises and services are dependent on, maintaining service performance and SLAs will undoubtedly require a new level of AI-driven automation. This means AI must guide operations to fully automate the decisions to self-correct and self-heal, with as little human intervention or limitations as possible. Next, let's look at why we fall short in reaching that ambition with the current and traditional AI approaches.

# How the traditional approach falls short

Today, operations implement AI as individual use cases that are often siloed and purpose-built to their own requirements. The process from business requirements as input to actuations as output is represented by a rigid and policy-based workflow as illustrated in Figure 3.

**Figure 3: Policy-based workflow**

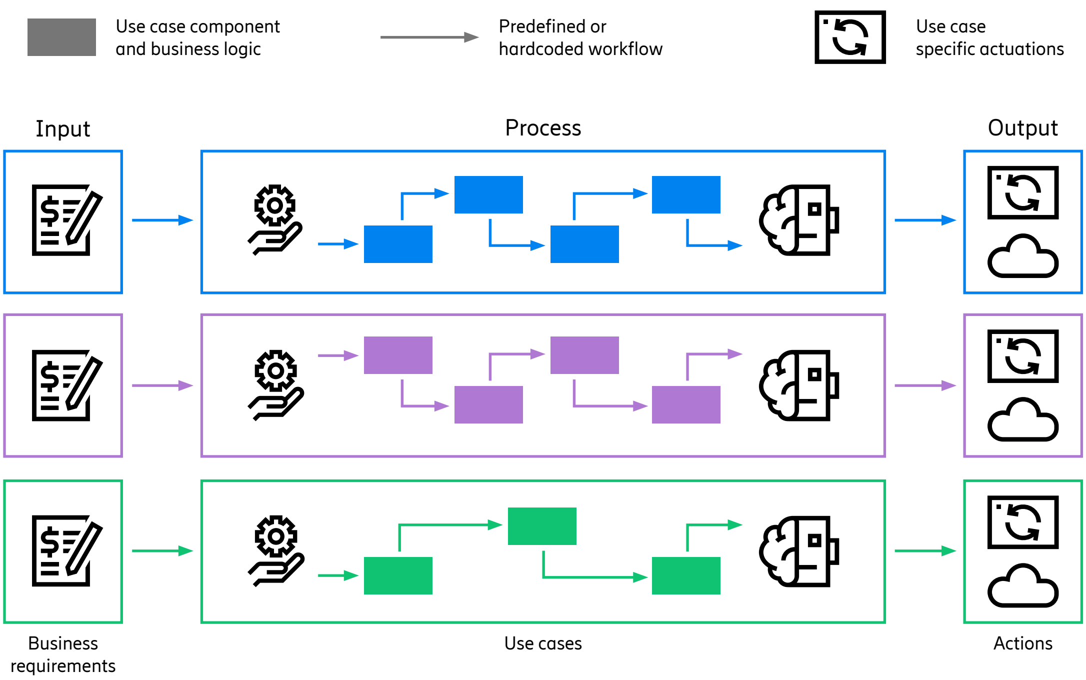


The process starts by first translating business requirements into technical requirements and then building a use case that delivers those requirements. The use case consists of carefully designed workflows, business logic and policies that work together based on the requirements.

While this approach has proven effective on a per-project basis, it will fall short in meeting the business expectations at scale due to the following reasons:

- Agility – whether it is requirement translation, use case implementation or policy execution, the rigid and often hardcoded logic will not easily adapt to new or changing business requirements.
- Time to market – with every new offering or change, the process requires new workflow, business logic and policy development that leads to longer turnaround and time-to-revenue.
- Conflict in business priorities – with siloed use cases focusing on their own requirements and unaware of others, conflict and contention for resources will arise without cross-use case business priorities and arbitration. For example, one use case may decide to actuate a policy to meet its own requirement but negatively impact another use case; or a use case wants to acquire additional resource to provide a service that is of a lower priority than another.
- Cost and efficiency – with a siloed use case approach, a vast amount of knowledge including logic, facts and data-driven insights are boxed-in and not visible or shared between use cases. This will lead to duplication of effort and higher overall cost between projects.

# Intent-based operations

Intent-based operations will transform the operations of tomorrow to align with new and changing business needs and deliver business monetization objectives. An intent-based operation is an automated process that onboards, manages and delivers business requirements provided by the business stakeholders. The process ensures that an overall desired outcome, comprised of the sum of all requirements, can be achieved using an autonomous and self-organizing operation.

**Figure 4: Intent-based operations process**

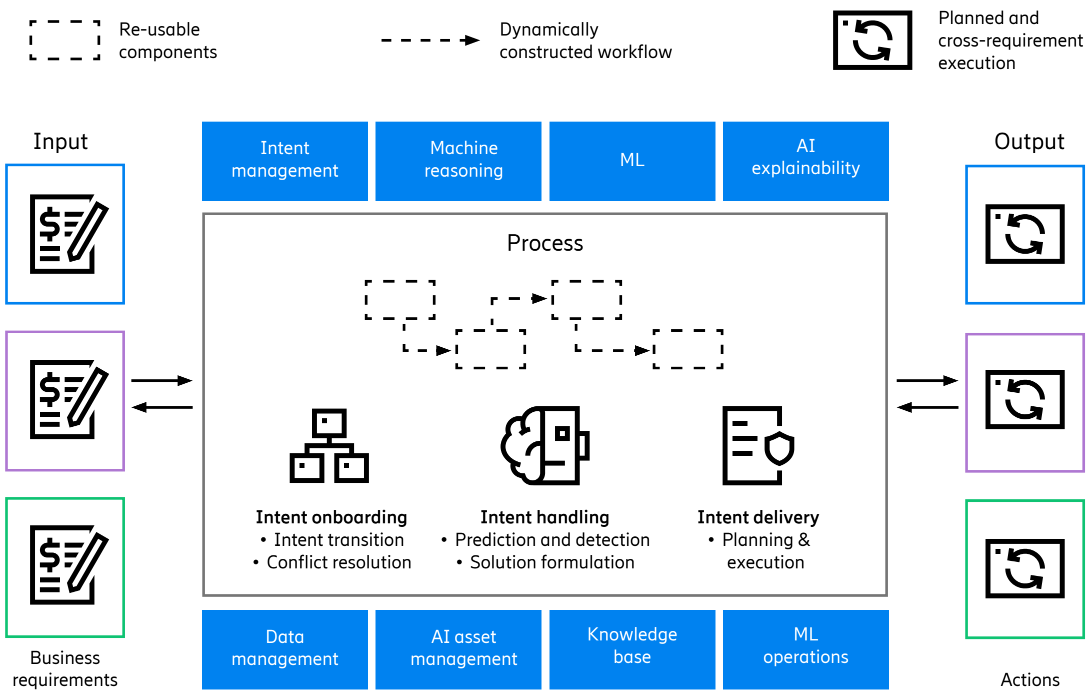


Business requirements refer to business objectives that need to be achieved. The business objectives are desired outcomes that can be measured, and can include any context, scope or constraints when delivering the outcome. In the context of operations, for example:

- Business requirement 1 - deliver an SLA of 25ms latency in the 5G private network for healthcare enterprise A.
- Business requirement 2 - reduce energy consumption by at least 10 percent and carbon footprint by 5 percent for all Ericsson radio sites in markets A, B and C.

As service providers continue to onboard new customers and launch new products and initiatives, it is easy to imagine the growing number and accelerated pace of not only new business requirements, but also changes to existing ones to better align with market needs and business priorities. The value of intent-based operations is to automate the management and delivery of business requirements in an agile, transparent and scalable way. Let us look at the key workflows:

**Intent onboarding**
Intent onboarding is a process where business requirements are translated and onboarded onto operations as "intent".. Intent, based on the definition of TM Forum Autonomous Network Project 2020, is "the formal specification of all expectations including requirements, goals, and constraints given to a technical system." In essence, intent captures the objectives described in the business requirement in a well-defined model so that the requirements, goals and constraints are well understood. Figure 5 illustrates how a business requirement is onboarded as intent by this process:

**Figure 5: Intent onboarding process**

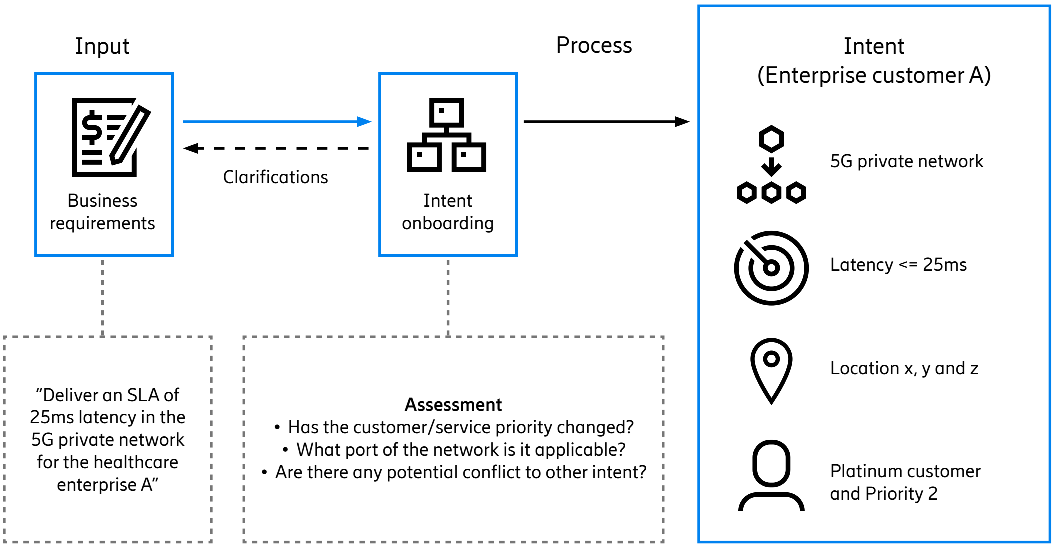


Intent onboarding assesses the business requirement to deliver an SLA of 25ms latency in the 5G private network for healthcare enterprise customer A and requests any clarifications necessary. During the assessment phase, priorities, scope and conflicts are analyzed to ensure that the requirement is well understood with no missing context, logic or ambiguities. In some instances, conflicts may arise because different requirements may impact the requirement of others. For example, one requirement may be to reduce radio site energy consumption while another to increase cell coverage. Therefore, a conflict resolution framework is required as part of the intent onboarding process to support configurable logic that will flag any potential conflicts before a requirement can be onboarded as intent.

**Intent handling**
After an intent is onboarded, intent handling is a process that manages the analysis of its state and if it has deviated from expectations. It then formulates a solution that would return it to a state that meets the intent.

**Figure 6: Intent handling process**

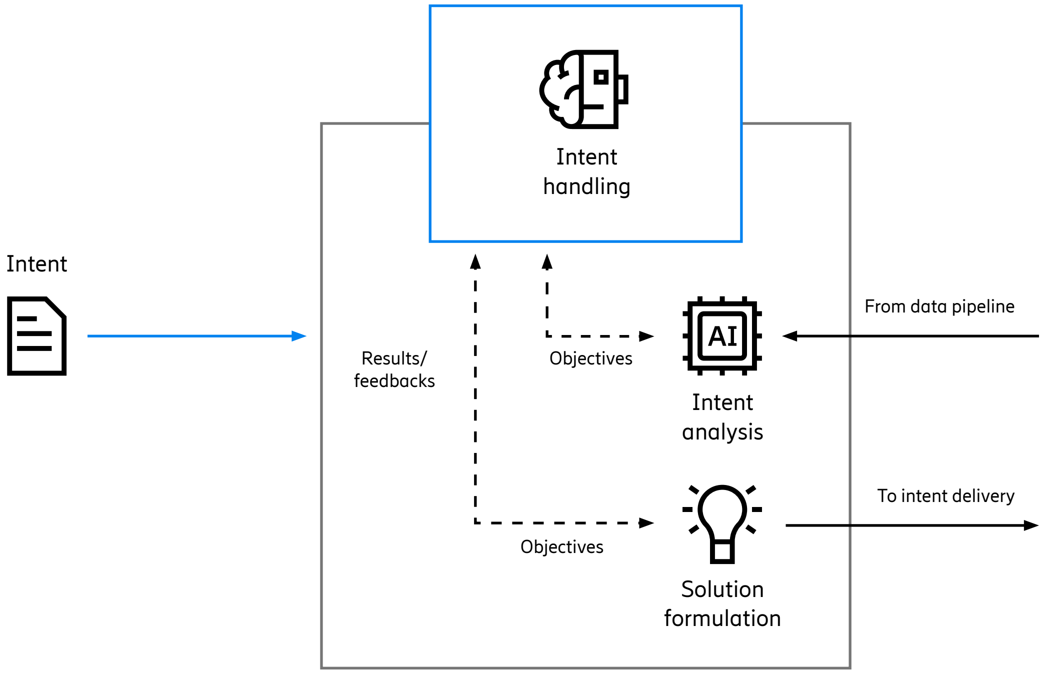


As shown in Figure 6, the intent handling process has two workflows – intent analysis and solution formulation. Intent analysis is responsible for analyzing the state of the intent and determining if the expectations have been met. For example, enterprise customer A is receiving less than 25ms of latency at locations X, Y and Z. Intent analysis will use data pipelines to ingest and process the necessary data and compare the current state with the desired state. ML techniques will be used in the analysis to predict if the desired state will be deviated from in the future.

If a deviation is predicted or detected, solution formulation will construct a workflow with the necessary tasks to find a resolution that corrects the problem and can return the service to the desired state. These tasks may include isolating the root cause of the issue, discovering and qualifying existing policies or producing a new solution.

Since intent is declarative in nature, it specifies what needs to be accomplished instead of how to implement a solution. Therefore, intent handling does not rely on predefined workflows, processes and policies. In fact, it is designed to adapt and solve using relevant and reusable assets – such as AI techniques and knowledge base graphs – that can best accomplish a certain task. This makes intent-based operations more autonomous and adaptive than the rigid and hard-coded approach. The output of an intent handling process is a solution that is the most feasible with which to deliver the intent. This solution may simply be an existing and predefined policy that had been qualified, or it may be a brand-new solution that had been constructed on-the-fly in the absence of a predefined and qualified policy.

**Intent delivery**
Intent delivery is a process that plans the execution of all solutions from multiple intent handling streams. Each solution provides a set of actions specific to its intent handling and includes contextual information such as priorities and task execution time as well as time to violation.

**Figure 7: Intent delivery process**

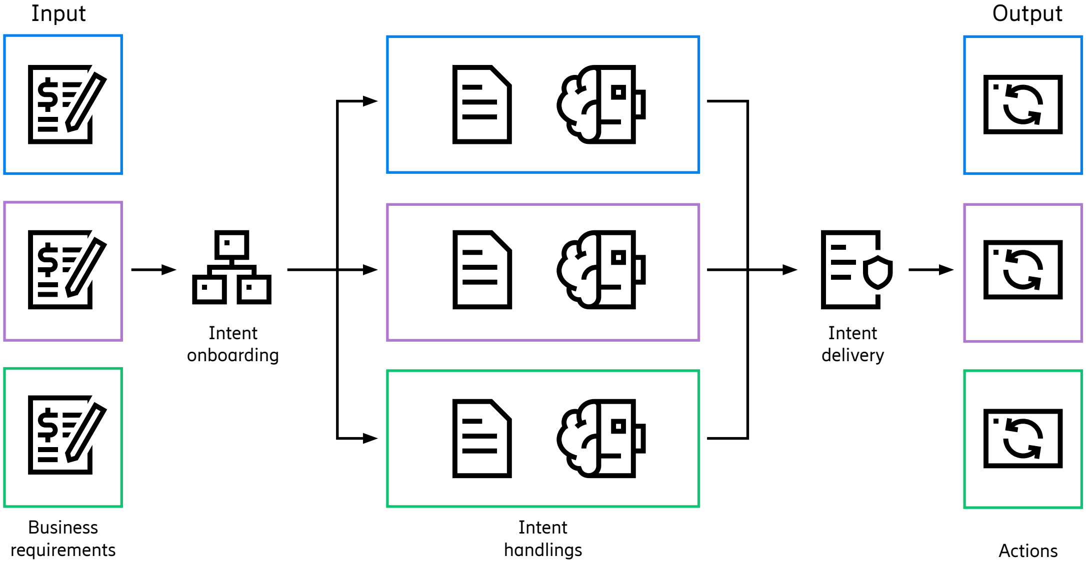


Intent delivery coordinates the actions from each solution to ensure the most optimal sequence of instructions towards the orchestration/actuation system(s).

This coordination needs to consider several factors, including:
- Priorities: Which solution serves an intent with a higher priority.
- Time to violation: How soon will an intent be violated.
- Task execution time: How long a certain task will take.
- Conflicts: Do actions from different solutions impact one another – for example, two solutions may contact actions that compete for the same resources at the same time. Alternatively, one may reverse the action of the other.

When conflicts occur and priorities between intents are clearly stated, intent delivery will prioritize one solution and instruct intent handling of the impact solution to reformulate a new solution (with additional constraints). The output of intent delivery is a set of well-planned and coordinated actions towards an orchestration/actuation system that will best deliver all the intents being managed.

**Powered by an autonomous and self-organizing system**
Intent-based operations rely on an autonomous and self-organizing platform that provides the capability needed to operate an intent.

**Figure 8: Autonomous and self-organizing system**

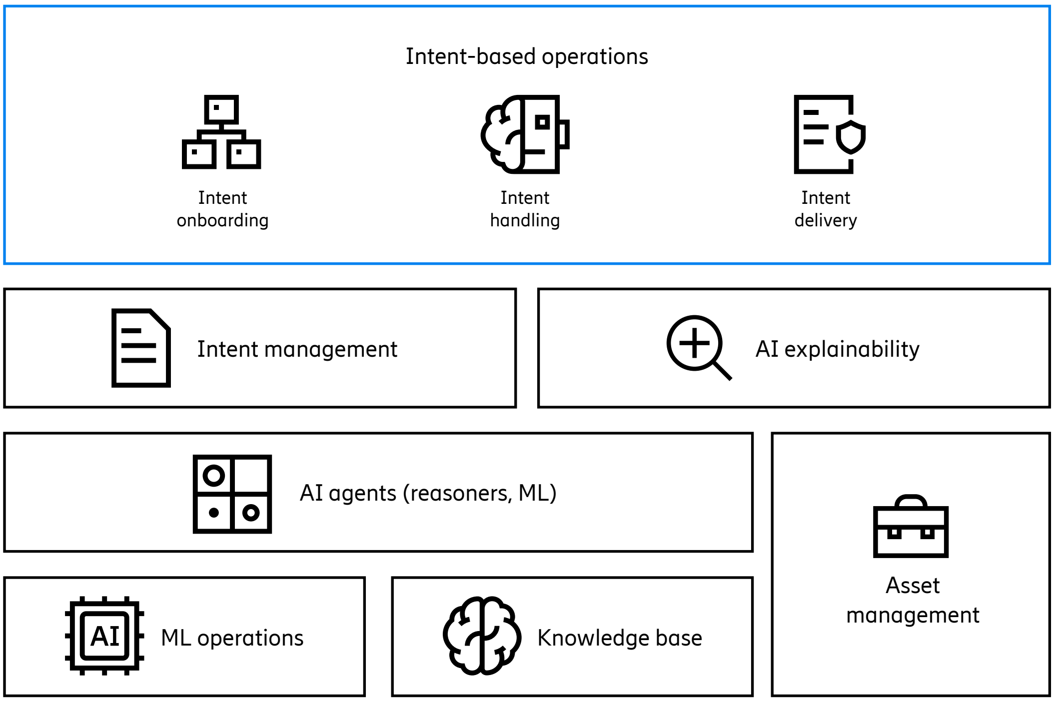


Figure 8 illustrates the main building blocks in the system that is used to support an intent-based operation. More specifically, knowledge base, AI agents and intent management are some of the most critical components.

**Knowledge base**
The effectiveness of an intent-based operation first and foremost depends on the knowledge base in the system. The knowledge base stores curated knowledge objects that represent everything the system knows. It can include persistent knowledge such as product details, business logic, rules, facts, procedural and even folk knowledge. The knowledge base can also include time-sensitive objects such as real-time inventory, ontologies and data-driven events, insights or predictions. Intent-based operations rely on the knowledge base to provide the necessary knowledge to fulfill different tasks.

**AI agents**
An AI agent is a "worker" programmed to execute a certain task. For example, to predict the likelihood of a KPI violation or find the root cause of a service degradation. The AI agent may have reasoning capabilities that, when given access to a knowledge graph, can draw conclusions based on certain reasoning techniques such as backward chaining or constraint solving. Intent-based operations rely on AI agents to execute the necessary tasks to complete various goals and objectives in an operational flow.

**Intent management**
Intent management provides the necessary framework, including data models, for an intent-based operation to operate on the system. The framework provides a structure as to how an intent is understood and processed throughout the system. It also enables a self-organizing methodology where necessary components are discovered and summoned to deliver a particular objective – before deciding the next step depending on the outcome – without any hard-coded logic.

# Real-life scenario

Let's look at a real-life scenario on how service providers may begin to onboard customers with new offerings and how an intent-based operation can help deliver the business objectives.

Imagine a service provider had built up an 5G infrastructure within a specified geographical area. The infrastructure consists of 5G macro radio sites and small cells; distributed radio units (RU) with centralized distributed units (DU) and centralized units (CU); central offices (CO) serving as edge and regional offices (ROs) hosting 5G core virtual network functions (VNF) and breakouts to the internet and hyperscale cloud provider (HCP) clouds; and finally, the 5G transport network.

Say the service provider first offers a 5G FWA service providing service providing broadband internet services targeting consumers. The business intent for the FWA service is to provide a certain peak download throughput (for example, 300Mbps) to a number of households (for example, 200,000) across the area, and at a certain internal cost to the service. This intent is onboarded by the intent-based operations given the network utilization, available capacity and the projected cost to support the service.

Next, the service provider would like to offer a smart warehouse solution to an industry warehouse enterprise customer in a particular location where the warehouse is. The smart warehouse solution provides guaranteed latency and throughput as SLAs to IoT devices such as robots, drones and sensors. Internally, operations would like to curb the maximum energy consumption to meet cost and sustainability targets. Intent-based operations decide that in order to onboard and deliver this intent, the workload from the FWA service at the same location must be re-distributed. In fact, the most optimal solution is to prioritize the industry warehouse enterprise traffic and re-route a portion of the FWA traffic to a nearby CO while meeting the intent of both services.

At this point, the services are running as intended. However, the service provider wants to offer a first responder service to public health and safety governance agencies such as emergency medical technicians/paramedics, police officers and fire departments. This solution provides an SLA that guarantees latency, throughput, security and availability and is offered to enterprises with a higher priority than others. While the network capacity isn't optimal to onboard such a high demand and profile service, intent-based operations nonetheless propose a solution where the service can be onboarded if the constraint in the intent of the other two services are relaxed. In fact, the proposal was to allow the increase of cost for FWA and energy consumption for smart warehouse while meeting the rest of the intent. The business and operational stakeholders sign off on the proposal and the new intent is successfully onboarded.

**Figure 9: Real-life scenario for an intent-based operation**

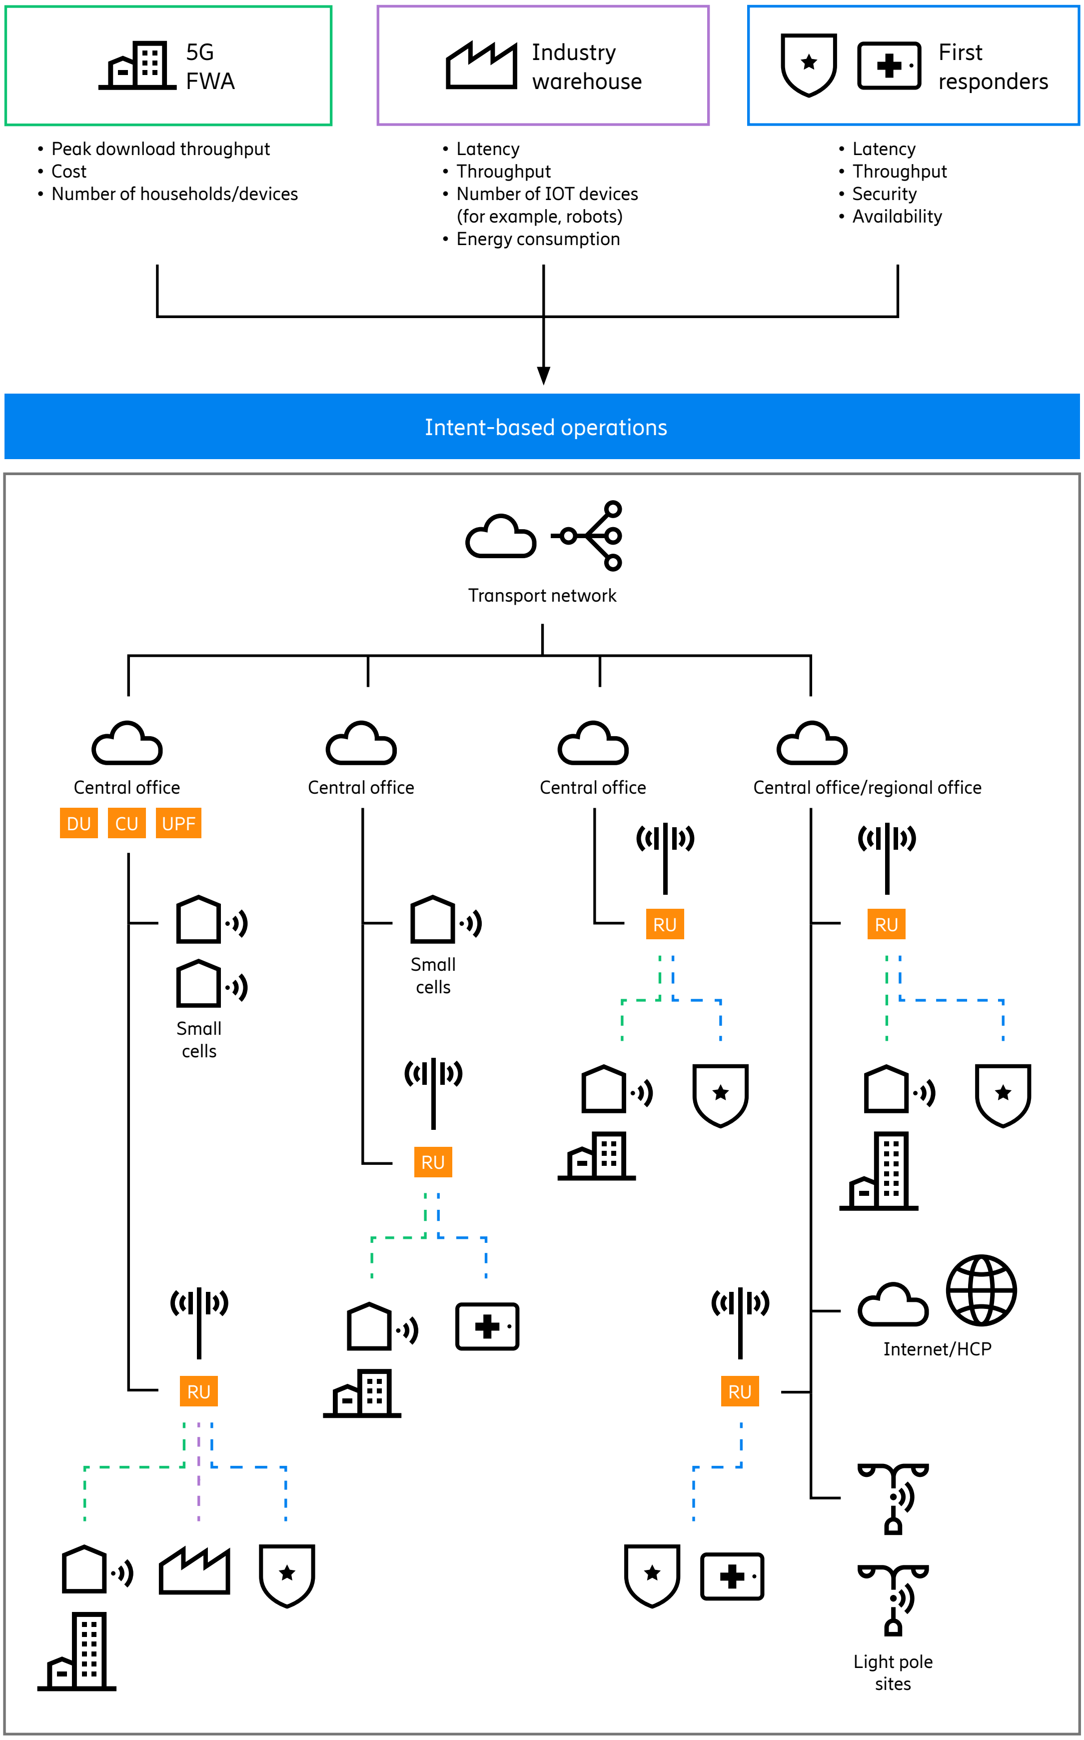

# Conclusion

As the pace of new and changing business requirements increases, agility and adaptability of operations will become critical to supporting the needs of the business. Intent-based operations enable operations to adapt and scale to new business objectives without costly and cumbersome development. It also provides the structure and clarity operations requires aligning resource allocations and processes with the business objectives and priorities.

Additionally, intent-based operations will deliver the transparency business stakeholders lack today on how well their objectives and customers' requirements are met. This transparency will allow business stakeholders to accelerate the go-to-market of competitive offerings, onboarding of new customers and monetization of businesses.

While AI and automation are key to achieving cost and operational efficiency, the service providers must start now to understand, implement, and invest in intent-based operations. For this reason, we at Ericsson believe that for the transition towards intent-based operations, service providers should consider implementation of seven key components:

1. Intent onboarding
2. Intent handling
3. Intent delivery
4. Autonomous and self-organizing platform
5. Knowledge base
6. AI agents
7. Intent management

Due to the expected scale, complexity, and criticality of future networks and use-cases, implementation of intent-based operations will be the way forward for service providers in terms of differentiation and helping them to monetize their investments.

**Glossary**

| | |
|---|---|
| AI | Artificial intelligence |
| ARPU | Average revenue per user |
| eMBB | Enhanced mobile broadband |
| FWA | Fixed wireless access |
| KPI | Key performance indicator |
| ML | Machine learning |
| MR | Machine reasoning |
| NOC | Network operations center |
| NPS | Net Promoter Score |
| SLA | Service-level agreement |
| URLLC | Ultra-reliable low-latency communication |

**References**

- Supercharging customer experience through AI and automation
- 5G for business: a 2030 market compass – Setting a direction for 5G-powered B2B opportunities
- Harnessing the 5g consumer potential – The consumer revenue opportunity uncovered
- Ericsson Mobility Report
- Cognitive technologies in network and business automation

# Contributors

**Jia You**
Jia has over 20 years of experience in the OSS industry and successfully incubated products in the AI/ML, Service Assurance and Analytics domains. He is currently leading innovation in the areas of 5G AI and automation for Ericsson Managed Services. He previously headed Product Management in Ericsson's customer experience management portfolio and held leadership roles in product management and marketing in the OSS domain. Jia graduated with a Master of Engineering from Columbia University in New York and a Bachelor of Engineering Science from the University of Toronto.

**Rajat Kumar Kochhar**
Rajat has over 16 years' experience in telecom, with roles in R&D, Managed Services, Business Development and Pre-Sales. He is a Senior Specialist at Ericsson for 5G test tools architecture and design. Prior to this, he had worked for Ericsson's Managed Services business unit in various customer-facing positions. Based in Linköping (Sweden), Rajat is a keen evangelist for emerging technologies and has written extensively on 5G, IoT, AI, CEM and Blockchain technologies. He holds a master's degree in Informatics and a bachelor's degree in Electronics, both from Delhi University, India.

**Agata Platek**
Agata is a Marketing Director at Ericsson. In her current role she works with strategic topics of AI and automation across Business Area Managed Services. She has over 10 years of experience with Ericsson working in various roles in marketing, communication and learning services. She is also the host of the Ericsson AI Operations podcast series.