Industry perspective: From automation to autonomy

Intent and knowledge as the foundation for decision-safe network operations

ERICSSON

In collaboration with

Telstra

This paper aims to advance industry discourse on Level-4+ autonomous networks by challenging a persistent assumption: that autonomy can be achieved simply by layering artificial intelligence (AI) onto existing automation paradigms. We argue that this approach is insufficient. Instead, a more robust foundation is required — one built on intent and a formalized knowledge plane. These elements are essential to enable decision-safe operations at machine speed. This paper outlines both the conceptual shift and a pragmatic path toward realizing this vision.

Foreword

For decades, Communication Service Providers (CSPs) have operated within an enduring trilemma: the need to balance scale, reliability, and efficiency. Historically, no operational model has successfully delivered all three simultaneously.
As digital infrastructure has become foundational to modern society—characterized by rising demand for capacity and expectations of continuous, uninterrupted service—CSPs have prioritized scale and reliability, often at the expense of efficiency.

Autonomy is not a new concept within telecommunications networks. The data and control planes—the layers responsible for forwarding traffic and determining routing—already operate with high degrees of autonomy. However, the management plane remains fundamentally different. It is here that networks are configured, maintained, and evolved, and where human involvement persists across nearly all lifecycle stages. While execution may be automated, planning, design, and decision-making remain human-driven.

As networks have grown in complexity, the demands placed on the management plane have exceeded the limits of human cognition. This has resulted in an environment where safe operation is increasingly difficult to maintain. Empirical evidence highlights the issue: approximately 80 percent of major outages stem from human-induced change. [1],[2].

The industry's apparent reluctance to modify live systems is therefore not conservatism, but a rational response to systemic risk. This leads to what we define as the fragility trap. To prevent failure, change is minimized, resulting in operational inertia, rising costs, and reduced efficiency. Breaking this cycle requires decoupling change from instability.

Current approaches, particularly the application of AI to traditional automation, fall short. While they enhance execution speed, they do not eliminate human dependency in design and decision-making. Scripts and workflows remain human-authored, with the potential of being critical points of failure. Intelligent automation is therefore an important transitional capability, but not the end state.

The fundamental shift lies in intent-based network autonomy. In this model, CSPs define desired outcomes, and the system autonomously plans, designs, validates, and executes changes. Trusted guardrails ensure that actions remain aligned with business objectives. This approach enables predictable, safe-by-default operations at scale. Engineers transition from operating within the loop of execution to governing the system itself—defining policies, setting intent, and overseeing outcomes.

By managing the lifecycle of the autonomous system rather than individual changes, CSPs can simultaneously achieve scale, reliability, and efficiency. This paper explores how this transition can be realised in a controlled and value-driven manner.

Mark Sanders,
Chief Architect,
Telstra

Mats Karlsson,
Vice President and Head of Solution Area Business and Operations Support Systems,
Business Area Cloud Software and Services,
Ericsson

What is broken, and why should we fix it now?

The fragility trap introduces ongoing costs, even in the absence of visible failure. To mitigate perceived operational risk, CSPs routinely over-provision hardware capacity. This creates a persistent capital expenditure burden while utilization remains suboptimal. In parallel, legacy configurations are often retained to avoid the risks associated with removal, leading to accumulated technical debt.

Operational governance mechanisms—such as approval boards and change reviews—further exacerbate inefficiencies. While intended to ensure safety, they introduce latency and consume valuable engineering capacity. Skilled professionals are diverted from innovation to administrative oversight. Additionally, restrictions on live-system access limit the ability of less experienced engineers to develop operational expertise, concentrating knowledge within a small group of individuals.

The transition to software-defined networks was intended to unlock agility. However, the surrounding operating model has not evolved accordingly. When every change is treated as high-risk, the update cadence of software-defined infrastructure mirrors that of legacy hardware systems. As a result, the promised agility remains unrealised.

From a financial perspective, the industry has effectively substituted agility with excess capacity. Over-provisioning serves as a hedge against the risk of change, reinforcing a cycle of high investment, low utilization, and constrained innovation.

Attempts to address these issues through enhanced analytics and automation are insufficient because they do not address the root cause. While execution may be accelerated, the upstream activities—interpreting intent, assessing current state, identifying gaps, and designing changes—remain manual. This is particularly critical given that many outages originate within the management plane rather than the underlying infrastructure.

Automation has successfully scaled execution, but it has not scaled cognitive capacity. The opportunity lies in transferring this cognitive burden to the system itself, allowing human judgment to focus on intent definition and oversight rather than operational design.

Automation and autonomy are not points on the same ladder

A fundamental misunderstanding persists within the industry: that automation and autonomy represent incremental stages along the same continuum. They address fundamentally different questions—specifically, who determines the sequence of actions required to achieve a desired outcome.

In automated systems, humans design the sequence of steps, which are then executed by machines. Scripts encode human reasoning, and any deviation from expected conditions requires human intervention. The machine acts as an executor, while the human remains the decision-maker.

In contrast, autonomous systems assume responsibility for both design and execution. Given a defined intent, the system determines the target state, evaluates the current environment, computes the necessary actions, and executes them safely. Crucially, this process occurs dynamically and can adapt in real time to changing conditions.

This shift changes the focus of the engineer. Rather than designing individual changes, engineers establish intent, define constraints, and govern system behaviour. They transition from being in the loop to operating on the loop, focusing on oversight and continuous improvement rather than direct intervention.

The urgency of this transformation is driven by three converging factors: increasing network complexity, rising customer expectations for rapid change, and the escalating cost of inefficiency. Incremental improvements to existing models will not suffice. A structural transformation of the management plane is required.

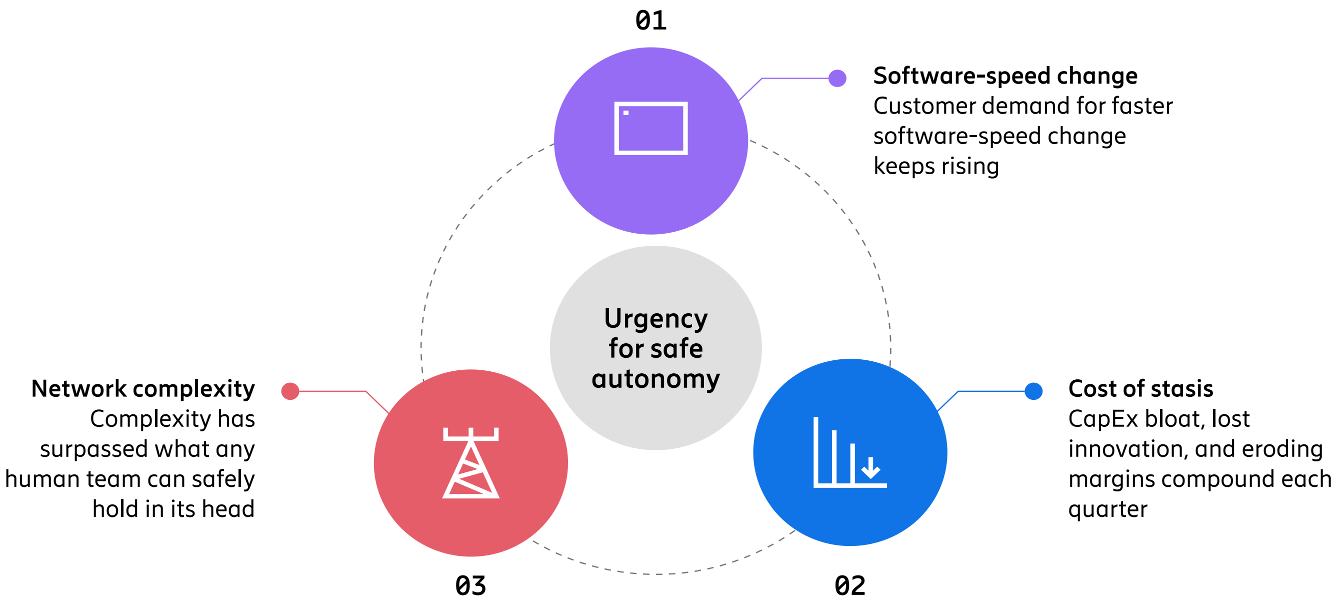

Figure 1. Factors leading to the urgency of safe autonomy

Intent and knowledge make decisions safe

Autonomous decision-making relies on two foundational elements: intent and knowledge. Intent defines the desired outcome, while knowledge provides the context required to achieve it.

The process of realizing intent can be understood as a sequence of stages: expressing intent, deriving the target state, determining the required actions, and executing those actions. For this process to be autonomous, intent must be represented in a formal, machine-interpretable manner, and the knowledge required for reasoning must be structured and accessible.

Foundation 1: Intent

Intent represents a formalized expression of desired outcomes. Unlike procedural scripts, intent is declarative and outcome oriented. It defines what must be achieved, rather than how to achieve it.

This abstraction enables flexibility and adaptability. As conditions change, the system can determine the most appropriate course of action while remaining aligned with the defined intent. Intent therefore provides both direction and governance, ensuring that decision-making remains consistent and controllable.

Intent guidelines for autonomous networks

Not all objectives are suitable for intent-driven autonomy. To ensure safe and effective implementation, intents must be clearly defined, unambiguous, and aligned with measurable outcomes. Poorly defined intents risk reverting to procedural automation, undermining the benefits of autonomy.

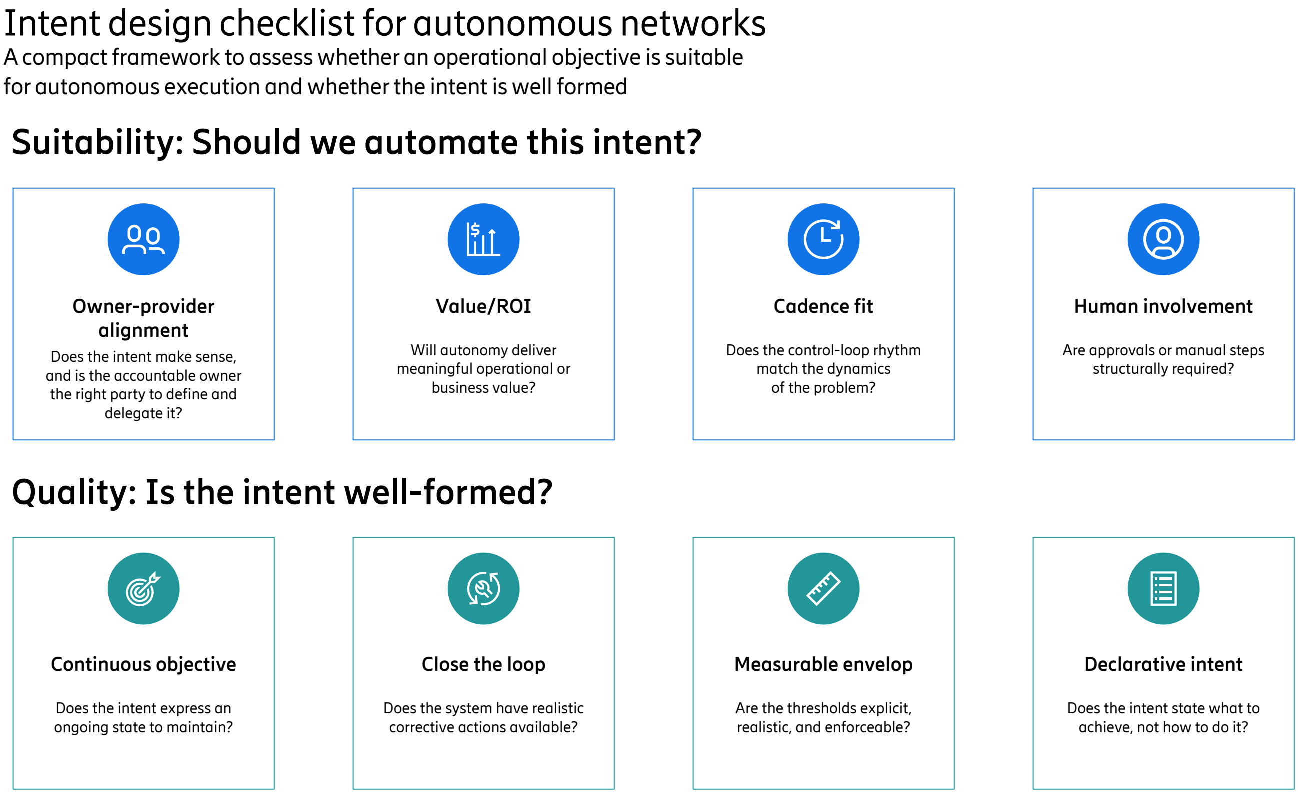

Figure 2. Intent design checklist for autonomous networks

Foundation 2: Knowledge plane
While intent provides direction, the knowledge plane provides the context required for safe decision-making. It encapsulates the information necessary for reasoning, including network ontology, semantics, dependencies, constraints, policies, and real-time state.

Unlike traditional documentation, the knowledge plane is dynamic and machine interpretable. It enables systems to validate decisions at runtime, ensuring that actions remain feasible and compliant with defined guardrails.

At a minimum, the knowledge plane must represent entities, relationships, constraints, and governance structures. Over time, it should evolve to incorporate learned insights and inferred knowledge, further enhancing system capability.

The implication is clear: achieving Level-4+ autonomy requires more than increased automation. It demands a structured knowledge plane capable of supporting safe, consistent, and explainable decision-making.

Realizing a continuous autonomous loop with intent and knowledge
Intent and knowledge together enable a continuous control loop in which decisions are not only executed rapidly but evaluated rigorously before action is taken.

This approach provides a scalable pathway to Level-4+ autonomy, enabling organizations to increase operational velocity without increasing risk. Speed becomes a natural consequence of competence rather than a substitute for it.

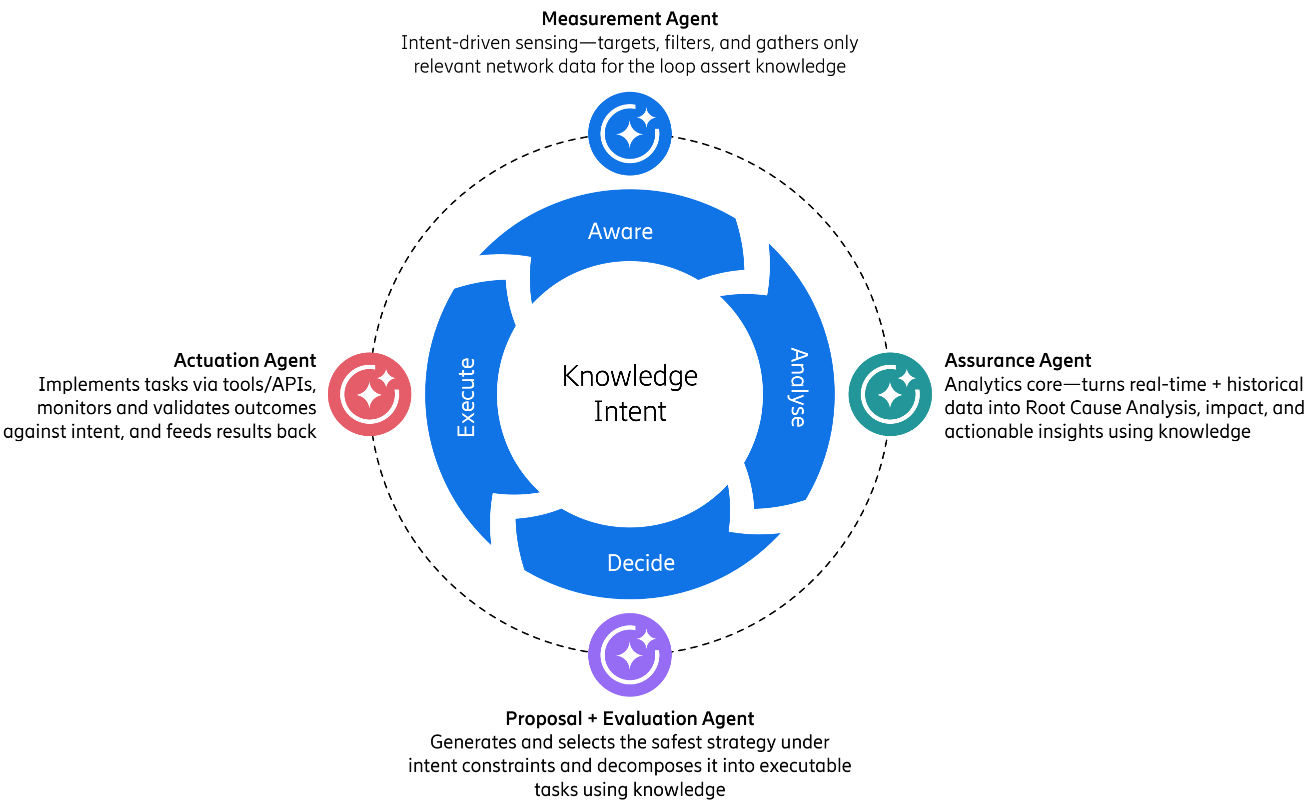

Figure 3. TMF IAADE control loop for autonomous networks

Agentic AI and trust in level 4+ autonomy

Agentic AI represents a significant accelerator for achieving advanced autonomy. However, its adoption introduces new challenges, particularly in relation to trust.

In this context, trust is operational rather than philosophical. Systems must be able to justify their decisions, demonstrate compliance with constraints, and provide traceable evidence for their actions. The combination of intent and the knowledge plane is essential to achieving this.

Intent defines the outcome boundaries, while the knowledge plane provides the structured context required for reasoning. Together, they constrain agent behavior, enable pre-action validation, and support explainability. This makes agent-driven autonomy both safer and more deployable at scale.

Autonomy that pays: The business case in measurable outcomes

To justify the investment and track progress, the CSP needs a business value measurement model that links Key Effectiveness Indicators (KEIs) to Key Business Value Indicators (KBIs) [3]. The model ties each outcome to measurable indicators and a clear value calculation method.

The benefits cluster into five business outcomes. Each one reversed a specific cost of the fragility trap.

- **Capital and energy efficiency**: Autonomous systems optimize resource utilization by right-sizing infrastructure to live demand. By dynamically deploying capacity to meet peaks and releasing it when demand falls, the network reduces the need for static over-provisioning, thereby lowering both capital and energy costs.
- **Innovation velocity**: By reducing the risk associated with change, autonomy enables faster release cycles and accelerates time-to-market for new services.
- **Operational health**: Continuous monitoring and proactive remediation reduce incident frequency and improve customer experience.
- **Governance and resilience**: Computational governance reduces reliance on manual processes while preserving and scaling institutional knowledge.
- **Commercial uplift**: Outcome-based service models enable differentiation and support higher-margin offerings. Instead of selling bandwidth as a commodity, autonomy allows CSPs to sell guaranteed outcomes—such as a latency target, an availability floor, or an experience level—moving commercial pricing from raw capacity to guaranteed results.

A practical path forward: Building knowledge case by case

A successful transition to autonomy does not require a comprehensive upfront ontology model. Instead, a scenario-driven approach is recommended.

Organizations should begin with targeted use cases where the cost of failure is high, and current processes rely on manual intervention. For each case, a minimum viable knowledge model should be defined, capturing only the elements necessary for safe decision-making.

This incremental approach enables rapid value delivery while supporting reuse and scalability. Over time, knowledge assets accumulate, accelerating subsequent implementations.

A federated operating model is essential for scale. Domain teams should manage local knowledge, while a central function ensures consistency and governance.

Call to action: Align and accelerate

The industry must recognize that true autonomy is not simply an extension of today's intelligent automation trajectory. It requires capabilities that allow machines to understand intent, reason over network knowledge, determine the required target state, and act safely without human intervention. Intent and the knowledge plane are foundational to this shift.

The practical path is clear and can begin immediately: build decision-safety scenario by scenario—starting with minimum viable knowledge slices and well-formed intents—and scale through reuse rather than one-off initiatives.

Accelerating maturity beyond early deployments will require strong industry collaboration. We believe there is an opportunity for TM Forum, together with its members, to consider dedicated collaboration workstreams to advance these foundational capabilities, align terminology and reference architectures, define common models and interfaces, and progress toward a common industry ontology that enables the knowledge plane to represent consistently, connect to, and reason over network knowledge across the ecosystem.

One concrete entry point is TM Forum's Moonshot Catalyst, Conflict management in intent-based networks – Phase II [4], where Ericsson and Telstra are collaborating with other members to explore how a knowledge plane can detect, explain, and resolve conflicts across intent layers. This provides a pragmatic environment to align on intent semantics, decision governance, and knowledge models—while demonstrating measurable operational value.

Further reading

- OSS/BSS: bridging business and operations - Ericsson
- OSS/BSS Solutions - Sell. Deliver. Get paid. - Ericsson
- AI-enabled service orchestration solutions - Ericsson
- AI-powered telecom data and analytics solutions - Ericsson
- IT with intent: the interconnected future of telco operations
- The future of telco IT operating models: the road to intent-based operations
- Telstra and Ericsson collaborate to shape the future of autonomous networks

Reference

1. https://www.cisco.com/c/en/us/solutions/collateral/executive-perspectives/technology-perspectives/automate-network-infrastructure-code-wp.html
2. https://www.ofcom.org.uk/siteassets/resources/documents/research-and-data/multi-sector/infrastructure-research/connected-nations-2024/connected-nations-uk-report-2024.pdf?v=386497
3. https://www.tmforum.org/resources/introductory-guide/ig1256a-autonomous-networks-business-value-measurement-v1-0-0/
4. https://www.tmforum.org/catalysts/projects/C26.0.930/conflict-management-in-intentbased-networks-phase-ii

Authors

Senthoor Punniamoorthy,
Principal Architect, Chief Architect Office,
Telstra

Senthoor is a principal architect and techno-strategist operating at the intersection of business strategy, emerging technology, and complex systems engineering. He specializes in synthesizing multidisciplinary thinking into long-term enterprise roadmaps and engineering credible paths to scale in high-ambiguity environments. Known for bridging industry and academia, Senthoor actively collaborates with industry leaders, research institutions, and global standards bodies to influence the industry's strategic direction early, ensuring long-term architectural rigor and execution momentum.

Vivek Singh Jain,
Strategic Product Manager, Data and AI,
Business Area Cloud Software and Services,
Ericsson

With over 20 years of experience in technology product management, Vivek specializes in shaping transformative product roadmaps and delivering solutions that replace manual processes, reduce cost, mitigate risk, and improve customer experience. He works closely with CXOs, product leaders, subject matter experts, and engineering teams to translate business strategy into scalable, high-performance products. He has a strong track record in driving go-to-market initiatives and supporting milestone deals through market insight and customer-centric thinking.

About Ericsson

Ericsson's high-performing networks provide connectivity for billions of people every day. For nearly 150 years, we have 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.

BCSS-26:016276 Uen