Architecting tomorrow

Towards cognitive and fully autonomous 6G networks

Strengthened by agentic capabilities, holistic orchestration, and ecosystem programmability

White paper

NOKIA

Executive summary

6G must move beyond fragmented, reactive automation to cognitive, agent-driven autonomy that enables networks to act on their own and act intelligently. Future services such as XR, digital twins, robotics, and autonomous mobility will demand all performance attributes simultaneously and at scale, creating vast complexity with dynamic, ephemeral slices across domains. Traditional automation cannot cope with this challenge.

The path forward rests on adaptive and cognitive autonomy that realizes Self-X [1] at scale. Autonomy makes the network act independently; cognition makes it act intelligently. This cognitive autonomy is underpinned by holistic orchestration and assurance, distributed-decentralized intelligence aligned with the autonomous network architecture and anchored in endogenic principles, agentic capabilities, intent-based networking, and open programmability. Together, these enablers allow networks to reason, learn, predict, and act proactively, delivering a true Zero-X [1] experience, enhancing operational efficiency, and maximizing business value.

The role of the network operator in 6G will transition from automating networks and services (aka. Automated Ops) to collaborating in a continuous dialogue with the network to define business and end user goals, validate outcomes, and sustain trust in cognitive and fully autonomous networks (aka. Autonomous Ops).

To realize this vision, harmonized standards are essential, embedding autonomy, AI, trust, security, and sustainability by design. With unified frameworks, 6G can achieve high to full autonomy, unlocking resilient, sustainable and value-generating digital ecosystems.

With its end-to-end portfolio built on service operations and autonomous domains (at the resource operations layer), Nokia has already made major strides on the path towards cognitive and fully autonomous networks [2].

Contents

Introduction

As services and ecosystems grow more demanding and complex, 6G is about creating networks that must thrive where 5G-Advanced reaches its limits. Future services, from immersive XR and holographic telepresence to large-scale digital twins, autonomous mobility, and mission-critical industrial robotics, will not just demand higher uplink, ultra-low latency, massive capacity, guaranteed reliability, seamless coverage, and energy efficiency. They will demand all of these at once, at scale, without interruption, and ubiquitously across every environment. In 6G, we can expect hundreds of specialized networks (i.e., network slices), because connectivity will need to be far more granular, context-specific and dynamic than today. Instead of a single healthcare slice, for example, there could be distinct slices for medical doctors with mixed reality glasses, for patient wearables, and for hospital IoT. Individual users may require tailored slices for their phone, AR glasses, watch, car, and home devices, while enterprises and cities will demand specialized vertical slices tailored to their operations and services. Crucially, these slices will not all exist permanently: they can be spun up or shut down on demand dynamically (rather than through static configuration), and autonomously, using fingerprinting and AI-enabled traffic characterization, creating a flexible mix of persistent and ephemeral slices optimized for each situation.

Network slicing in 5G/5G-Advanced has already brought significant value by enabling tailored virtual networks with differentiated SLAs (now referred to as service-level objectives and/or expectations, and abbreviated as SLO or SLE) and new business opportunities. But it has also introduced substantial complexity, each slice requires its own lifecycle management, spans multiple domains, and involves diverse stakeholders, making end-to-end assurance and scalability increasingly challenging. In 6G, this challenge will intensify as the number, diversity and simultaneity of slices expands and as slicing becomes far more dynamic with slices being created, adapted and terminated on demand. This will require intelligent, end-to-end orchestration and assurance embedded in intent-based closed loops that today's disconnected or monolithic automation schemes cannot deliver.

This growing complexity of services is compounded by the evolution of the underlying network complexity. Cloudification and virtualization, and the decomposition of the software stack into multiple layers of microservices, have brought significant value in terms of flexibility, scalability and innovation. Yet they have created fragmented foundations with each layer applying different automation principles and often controlled by different organizations adding substantial complexity to network design and operations.

For operators, this rising complexity translates into higher total cost of ownership (TCO), slower time-to-market, difficulties in assuring SLAs, SLOs and SLEs end-to-end, and limited ability to monetize new services at scale. Fragmented, reactive automation and manual interventions increase inefficiencies, errors, and risks—threatening both profitability and customer trust. What is required is a vision for autonomous networks that can establish an autonomous network operations strategy to address these challenges.

Key pillars for cognitive and fully autonomous 6G networks

Building on today's automation foundations, the next leap is toward adaptive and cognitive autonomy by design, strengthened by agentic capabilities. In essence, this refers to autonomy that can adapt dynamically, reason and learn, and act with a holistic, end-to-end view. The Nokia Autonomous Networks vision described below combines end-to-end service operations autonomy and domains autonomy, building on key pillars and design principles that together enable the path toward cognitive and fully autonomous networks. This vision advances our two-pronged ambition for operational efficiency and new monetization opportunities.

Achieving this level of autonomy requires intelligent, cooperative agents with access to invaluable data across domains, user equipment (UE) and external sources powered by technologies such as advanced AI, digital twins (DTs), and cloud-native architectures. Together, these provide the intelligence, foresight and scalability required for predictive and preventive assurance, holistic orchestration and open programmability, ensuring both technical excellence and sustainable business value creation.

The evolution of AI provides the foundation for this step change. Earlier generations of AI delivered narrow, rule-based automation, and today's AI enables data-driven analytics and prediction but often in isolated domains. The emergence of new AI technologies such as GenAI, CausalAI and Agentic AI opens the way for contextual reasoning, collaboration and proactive decision-making, making cognitive, autonomous networks a realistic goal for 6G. While AI, and agentic AI in particular, are clearly transformational and help accelerate the evolution toward cognitive and fully autonomous networks, it is important to weigh the power and computational cost of real-time AI and to have the capabilities to flexibly place the intelligence where and when its benefits justify the cost.

Complementing this, purpose-built digital twins will provide virtual replicas of networks, services and applications that enable simulation, prediction and proactive assurance. By combining AI with digital twins, agents gain the foresight and context required to manage complexity across domains and ensure continuous service quality.

This step-change rests on key pillars that together define how 6G networks will evolve toward full autonomy, as discussed in the following.

Adaptive and cognitive autonomy (Self-X at scale)

6G networks must move beyond static or reactive automation to fully adaptive and cognitive autonomy. This means self-fulfilling, self-assuring and self-optimizing (Self-X) capabilities that not only act dynamically but also reason, learn and predict to optimize resource utilization and quality of experience proactively. By combining adaptability with cognition, networks will be able to manage resources intelligently, anticipate and prevent disruptions, and continuously align with and satisfy service intents. Such capabilities also ensure a true Zero-X experience (zero touch, zero wait and zero trouble), consistently delivered across network domains and reflected in seamless user experiences across diverse services and use cases.

Holistic orchestration and assurance

Orchestration and assurance must evolve into a holistic, end-to-end framework that integrates technological domains such as radio access (RAN), fixed access (FAN), core, IP and optical transport, edge, and cloud and that spans multiple stakeholders. This requires a collaborative, multi-domain view of resources, services and policies, supported by knowledge graphs that provide semantic insight into disparate operational data and interconnected entities, enabling consistent outcomes across domains, avoiding conflicts and ensuring end-to-end service quality.

Agentic capabilities

The shift to 6G will be enabled by intelligent autonomous agents, each designed to operate with full autonomy for specific tasks, yet possessing the intelligence to cooperate with others when such collaboration creates shared value or broader benefits —for example, amplifying overall system performance.

These agents combine agentic autonomy with cognitive awareness—the ability to sense, think and act, drawing on rich and trusted data to support reasoning, learning and adapting. Crucially, they are not confined to pre-programmed tasks but can acquire continuously updated knowledge, respond intelligently to new situations, take on new tasks beyond their original scope, and evolve their strategies and actions to achieve greater effectiveness.

Agentic AI can be embedded within domain-level control loops (aka. inner loops, in-domain loops) to act locally in real time, and it can also operate as part of higher-level loops (aka. outer loops), where cooperative agents coordinate across domains, interpret intents, troubleshoot issues in rapidly changing environments, and generate automated resolution actions to ensure end-to-end outcomes. These agents leverage extensive, trusted data and, when suitable, large language models (LLMs) to analyze context and guide decisions, thereby reducing the need for human intervention in network operations.

AI for IT operations (AI-for-IT-Ops) will remain essential in 6G. In 5G/5G-Advanced, AIOps has been a critical step forward, applying AI to large volumes of operational and telemetry data to automate assurance through anomaly detection, outlier detection, root-cause analysis, predictive operations, and aspects of network optimization such as QoS, resource scheduling, and load balancing. In 6G, AIOps will continue to provide the data-driven operational fabric. Agentic capabilities will build on and complement AIOps by interpreting intents, contextualizing insights, prioritizing objectives, and resolving conflicts, and by orchestrating resources proactively across domains to generate and execute actions at scale, with minimal or no human intervention. These agentic capabilities will, for example, drive proactive improvements in energy efficiency, QoS, and resource scheduling, enabling more sustainable, optimized and value-generating networks.

While intent-based management (aka. intent-driven operations) is already being widely experimented with, it is often template-driven and domain-bound. As networks grow too complex to be managed through low-level policies, human-designed worfklows, and configurations, intent becomes essential. It allows operators, enterprises and applications to express what outcome is desired without prescribing how it should be achieved. In 6G, services and requirements will increasingly be expressed as dynamic intents—high-level goals that span domains and stakeholders. These intents may be provided through formal models, APIs or even natural-language prompts, making them accessible not only to operators but, also, to enterprises and applications. For example, a fire brigade commander could issue an instant request such as “immediately prioritize mission-critical communications and real-time video feeds from all responder devices in the affected area and also for the two drones that we have with us today”. The network would automatically (autonomously) establish a specialized, resilient slice on the spot, autonomously coordinating and assuring resources across domains to support diverse devices such as bodycams, drones with 360-degree cameras, and smartphones, to ensure seamless collaboration under emergency conditions.

The autonomous agents will be responsible for interpreting these intents, recommending feasible options, avoiding and resolving conflicts, prioritizing competing goals, and assuring fulfillment—with minimal or no human intervention. At the same time, they must maintain an audit trail of their actions and decisions, ensuring that outcomes can be reviewed and understood by humans whenever required. This is essential to make autonomous networks credible and trustworthy for operators, enterprises and regulators alike.

Open programmability

Finally, 6G must be open and programmable by design. Fit-for-purpose APIs will expose network services in a way that is developer- and enterprise-friendly, enabling third-party innovation and new ecosystems. This openness should further enable third-party intelligent controller apps to help specialize and customize the network, meeting the requirements of distinct vertical use cases from autonomous vehicles to mission-critical healthcare and industrial automation.

For operators, this openness transforms the network into a platform for co-innovation and co-creation enabling them to expand beyond connectivity and capture value across entire digital ecosystems. At the same time, by specializing the network to the specific needs of verticals, operators deliver differentiated value, becoming indispensable enablers of industry transformation by powering intelligent, adaptive and purpose-driven networks that accelerate innovation across sectors.

Nokia Autonomous Networks vision

Building on these pillars that define the path toward full autonomy, the Nokia Autonomous Networks vision (depicted in Figure 1) serves as Nokia's blueprint for the industry's transition toward Levels 4 and 5, as introduced by the TM Forum Autonomous Networks framework (IG1230, IG1218 and IG1251 series) [1]. The vision rests on a coherent set of common enablers that deliver a unified and superior customer experience, enhance operational efficiency, and unlock opportunities to monetize the network. This vision is not just theoretical; in collaboration with its technology partners, Nokia has already successfully showcased how Autonomous Network at maturity Level 4 can be realized [3]. The PoC showed the path to zero-wait, zero-touch, zero-trouble operations, through intent-feasibility analysis and proactive detection and resolution of issues. The PoC leveraged AI, agentic AI and digital twins for impact assessment and risk management, and it is a leading example of what the ambition level is in 6G.

Agentic autonomy paradigm

At the heart of this vision lies the agentic autonomy paradigm, which, as depicted in Figure 1, combines centralized or distributed AI-powered engines with intelligent agents to realize intent-driven, flexible, adaptive, and scalable automation across 6G networks and services. Its key capabilities include:

- **Autonomous engines** (centralized or distributed) that interpret high-level intent and orchestrate associated goals across domains, delegating them—with related intents—to agents for execution
- **Autonomous agents** (centralized or distributed) that make decisions and act autonomously to fulfil delegated goals, guided by intent, while continuously learning, adapting to context and environment changes including disturbances and collaborating across boundaries
- **Continuous feedback (control) loops** that enable adaptation and optimization across layers and domains
- **Hybrid intelligence models** that support both centralized and distributed intelligence, balancing global optimization with localized, real-time responsiveness.

Figure 1. Advancing intent-driven operations with the agentic autonomy paradigm

Service intent flows from business operations to the service operations, which acts as an autonomous engine—interpreting and fulfilling intents through end-to-end service orchestration, assurance, policy enforcement, and by delegating and coordinating intents across autonomous agents. The autonomous agents operate autonomously within their defined scope to satisfy delegated resource intents through, for example, resource orchestration and control, service fulfilment and assurance. The service operations autonomous engine reports the intent status to the business layer, providing insight not only into goal achievement but also into further business opportunities identified through network capacity utilization. In addition, proactive feasibility assessments will help minimize the risk of failures and improve responsiveness to dynamic service needs.

Beyond intent insight, the autonomous network will evolve from suggesting, ingesting and recommending intents to actively generating intent. In other words, reversing the flow from human-to-network to network-to-human, marking a shift toward networks that operate autonomously without human intervention, with the ability to assist humans when needed—creating a continuous dialogue with the network.

Each autonomous domain receives its intent, decomposes it, and instantiates the required services and resources leveraging, for example, AI models and control loops. Domain controllers, supported by autonomous agents, continuously and in real time monitor key performance indicators (KPIs) and SLAs to verify alignment with the intended state, autonomously performing self-healing and proactively optimizing performance. They report the current intent status back to the service operations layer. The autonomous agents may also act as engines, delegating intent to further agents to operate autonomously to satisfy the delegated intents within or across network domains, coordinating their actions to satisfy the delegated intents.

Key design principles

Realizing this paradigm requires adherence to a set of design principles, including:

1. **Decoupled operational layers** that ensure flexibility and scalable evolution by separating the service operation layer, the resource control layer, and the underlying resource/infrastructure layer. The service-operation layer focuses on achieving end-to-end objectives and assured outcomes, while the resource-control layer focuses on concrete, real-time actions, assuring the underlying resources and reporting status and assurance data back to the upper layer.
2. **Separation of concern** underpins single domain autonomy by giving each domain autonomy within its scope, with in-domain controllers keeping resource control close to the resources and abstracting complexity. Each domain maintains its own real-time state and inventory of resources, providing information for local control and assurance while exposing only the abstractions required by upper layers. This helps reduce complexity, increase modularity, and allows independent evolution of each domain.
3. **Model-driven open interfaces** expose network capabilities through standardized, developer- and enterprise-friendly service APIs that support multiple levels of abstraction, including intent-based interaction. These APIs provide a consistent framework to expose services at various abstraction levels, from low-level resources to high-level intent, streamlining integration, ensuring interoperability, and accelerating the delivery of automated services.
4. **Distributed (decentralized) intelligence**, anchored in endogenic principles, places cognition where data is available, at the right quality and performance requirements for demand decisions, thereby avoiding bottlenecks and enabling scalability, per endogenic principles.
5. **AI nativeness** at all layers and domains embeds intelligence end-to-end throughout the 6G system, supporting continuous learning, prediction and adaptation.
6. **Control-loop** combines sensing, thinking and acting capabilities within and across domains, potentially with agentic AI in the loop to strengthen Self-X capabilities (healing, optimization, etc.) and adaptive autonomy. They need to be dynamically composable, and their lifecycle must be manageable to adapt to changing network conditions and objectives.
7. **Cross-domain collaboration** enables cooperative agents and orchestration across RAN (including non-3GPP access), FAN, core, transport, edge, and cloud to assure end-to-end outcomes, while avoiding and resolving potential conflicts.
8. **Data-centricity** (as a critical-asset for automation, cognition and autonomy) and semantic abstraction use common models to unify how data is interpreted, shared and acted on across domains. This requires inherent access to varied, invaluable network data within and across domains along with UE and external data, to enrich the intelligence capabilities of each domain and of the 6G system as a whole. This data-centric approach is complemented by modern data architectures, such as data mesh, federated data and/or data fusion, that enable scalability, interoperability and shared intelligence across domains.
9. **Security-, privacy-, and trust-by-design** embed safeguards in every layer and domain to ensure autonomy is accountable and trustworthy.
10. **Resilience and adaptability** provide self-healing, fast recovery and graceful degradation under stress or attack.
11. **Sustainability awareness** drives energy efficiency and reduces environmental impact through optimized automation.

12. **Seamless integration of purpose-built DT**s ensures that virtual replicas of networks and services are embedded into operations to augment proactive assurance and support adaptive autonomy within and across domains, for example, through simulation, what-if analysis, service feasibility checks, and pre-validation of automations.
13. **Human-in-the-loop** (where needed) reserves trust and oversight in critical scenarios, while the overall ambition remains full autonomy with minimal or no human intervention, also known as 'human on the loop'.

These principles define how cognitive autonomy must be architected and governed. Beyond these foundations, autonomy can be further enriched by providing a trusted service environment where certified third-party intelligent controller apps can help customize network behavior for specific vertical use cases, unlocking the 6G network potential and fostering a vibrant ecosystem of innovation. Such a service environment can be hosted within each domain-controller environment or in the end-to-end service-operation layer, while maintaining privacy, security and accountability.

Conceptual viewpoint from Nokia Bell Labs research

Nokia Bell Labs has introduced the forward-looking concept of UNEXT (Unified Networking Experience Technology [4]) that complements the vision outlined in this paper. The approach is to build an autonomous, secure network software system, whose components may be embedded into all functional layers of the network.

UNEXT ensures a unified networking experience by providing a harmonized set of key architectural components and operational principles applicable to all network technology domains, hence, enabling a gradual composition of an autonomous network for any use case. It identifies key areas for unification and, in the context of autonomy, extends it across various dimensions of networking, such as:

1. Automated and unified symbiotic adaptation between applications and networks [5]
2. Autonomous runtime service composition and assurance using elementary services coming from multiple actors, including consumer and network operations services [6]
3. Automated creation and maintenance of a compute-network continuum and deployment of workloads across a multi-stakeholder environment [7]
4. Automated inter-operation using contracts in a low-trust multi-domain, multi-stakeholder environment [8]
5. Automated knowledge and data services across domains [9].

It further proposes a unified, multi-agent architecture approach that operates at the most granular level to advance towards Levels 4 and 5 of the TM Forum's network autonomy model [1]. This approach will be capable of addressing hidden challenges, such as those related to intent interpretation and conflict resolution and enable autonomous self-composition of the network for any use case.

The need for standards

To realize the journey toward cognitive and fully autonomous networks, the design principles and a future-ready, novel end-to-end architectural framework need to be underpinned by a coherent set of standards-based enablers.

Common capabilities and enablers

Standardization is a must to ensure interworking and interoperability in heterogeneous environments shaped by fragmented technologies, diverse vendors and technology suppliers, varied deployments, and multiple stakeholders. Beyond overcoming fragmentation, standardization also provides the enablers for consistency, accountability, explainability, and trust in autonomous operations, while enabling solutions to scale on a common foundation that fosters continuous innovation.

These standards provide the common framework, core principles, interfaces and models that enable an intent ecosystem—covering grammar and models, capability discovery, lifecycle management, explainability, traceability, negotiation, and reporting—so that intent can be consistently expressed, interpreted and assured across domains.

Furthermore, the standards define common information and data models, ontologies and model-driven APIs to support evolved management capabilities including advanced intelligent FCAPS (fault, configuration, accounting, performance and security)—paving the path from reactive to predictive and preventive networks—as well as predictive and preventive networks, service/slice lifecycle management and orchestration, and robust data products underpinning effective data management and services, and ensuring consistent access to data as a foundation for analytics and assurance.

Standards also enable data-driven decision-making through AI frameworks, the dynamic composition and observability of control loops and their interactions, and the seamless integration of purpose-built digital twins to provide simulation, validation and proactive assurance. To support AI capabilities, standards enable comprehensive AI model lifecycle management, including support for training, deployment, monitoring and retraining, while ensuring interoperability across platforms and embedding trust and explainability. Where agentic AI is introduced in the control loop as an implementation choice, the hosting environment must expose standard interfaces, lifecycle hooks and trust frameworks so agents can sense, decide and act in alignment with intent while remaining accountable and interoperable.

Together, these standardized common enablers underpin the capabilities shown in Figure 2, where foundational enablers at the base underpin the capabilities shown in the middle, all building toward cognitive and fully autonomous networks at the top.

Figure 2. Common enablers and capabilities underpin cognitive autonomy

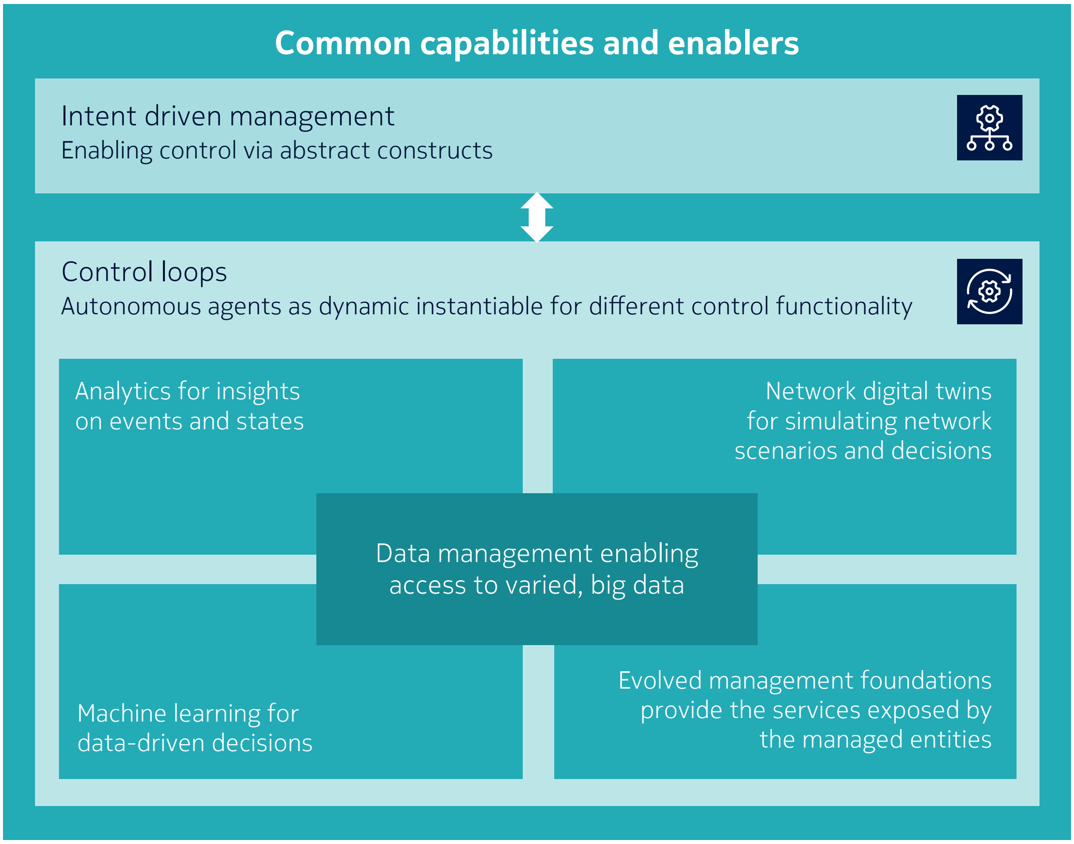

Beyond the capabilities described above, standards should further support the full lifecycle management of third-party intelligent controller apps, allow only authenticated intelligent controller apps to consume and expose management capabilities, and enable the portability of these applications across different platforms.

Today's standardization landscape

Given that cognitive control and autonomy cut across multiple network layers and technologies, the current industry landscape, illustrated in Figure 3, extends well beyond the scope of any single standards organization, with each working within its own domain of expertise.

Despite notable progress across many technology areas, the current standards landscape is a work in progress, representing varying levels of maturity, with alignment and coordination across organizations and domains needing to be strengthened. Key enablers—such as the outdated FCAPS services, control loops, intent frameworks, AI frameworks and services, data management and services, analytics services, and DT capabilities—are specified across different standards bodies, at varying levels of maturity, and too often in isolation, with no common approach and with differing timelines and decision processes that create gaps, overlaps and slow overall progress.

This fragmentation prevents a consistent framework for enabling a unified customer experience, coordination of intelligence across and within domains, and flexible placement of the intelligence so decisions can be made where they are most timely and effective. Besides, it demands additional effort for cross-domain interaction and for gaining access to data outside each domain. This, in turn, increases the risk of duplicated efforts, hinders efficient use of implementation resources, and delays the availability of these technologies. Moreover, native support for AI, security, privacy, trust, and sustainability is largely absent or treated as an afterthought, and DTs are typically added on top rather than seamlessly integrated.

In addition, standards already provide some support for the service environment for third-party intelligent controller apps within a specific technology domain, but this environment should ultimately be generalized to enable it across all domains.

Only through harmonization can these separate efforts converge into the seamless, cognitive and fully autonomous networks envisioned for 6G.

Figure 3. Industry landscape enabling cognitive network control and autonomy

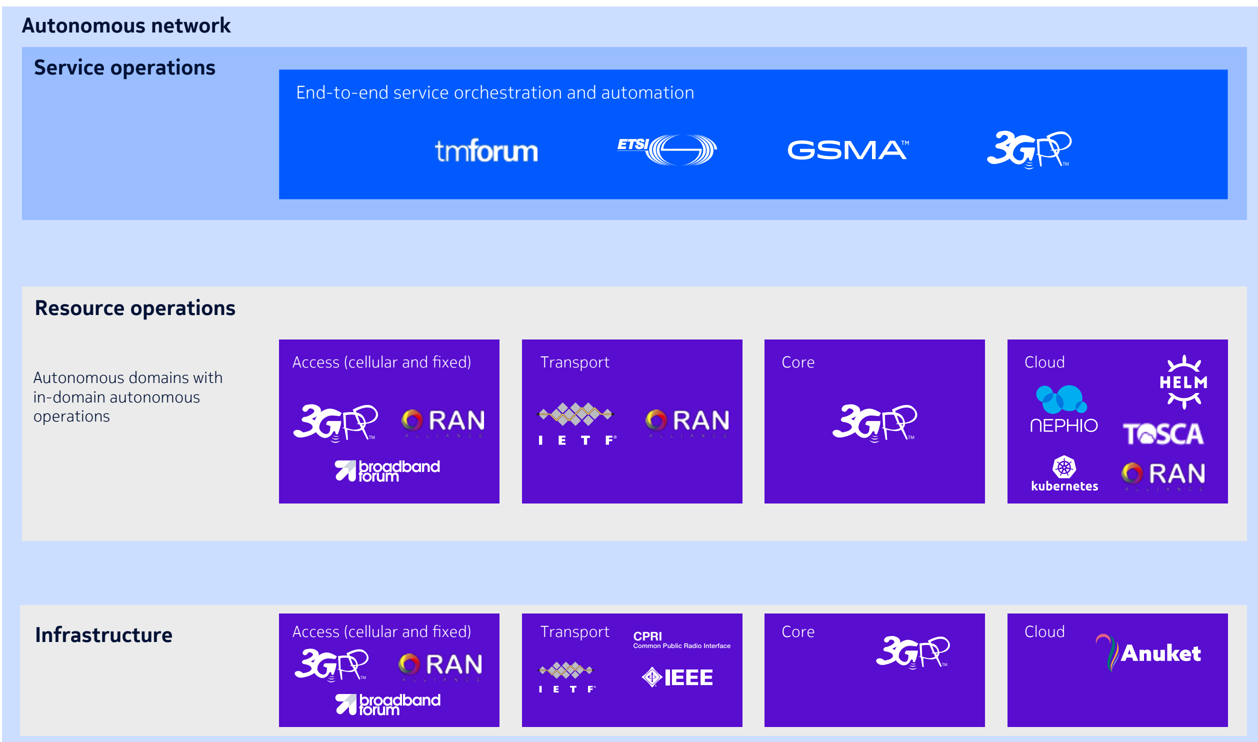

Driving harmonization to catalyze new and sustainable value

6G must be the catalyst for harmonizing and simplifying architecture, data models and enablers across and within standards organizations embedding AI, security, privacy, trust, sustainability and DTs by design. In this way, it will enable flexible multi-vendor deployment and deliver adaptive and cognitive autonomy that scale innovation, unlock growth and efficiency, and provide seamless user experience at lower cost. To fulfil the identified needs for standards highlighted earlier, harmonized 6G standards should provide a common foundation to avoid duplication, ensure end-to-end interoperability, and allow flexibility in implementation and deployment, focusing on the following key areas.

End-to-end architecture

6G standards should enable a simple, modular and flexible, future-proof, G-agnostic architecture for intent-driven, cognitive and fully autonomous operations. They should facilitate the use of agentic capabilities and balance backward compatibility with innovation as well as focus on integration points and exposed services/capabilities to support diverse implementations and deployments, all while reducing interoperability overhead.

Core capabilities by design

6G standards should embed AI, security, privacy, sustainability and DTs as intrinsic capabilities, designed into every layer and domain from the outset, with governance frameworks to ensure trust, accountability and sustainable operations. This requires common standardized enablers to make this native support real:

1. **Intent models and interfaces**—common grammar, information models and APIs so that intent can be expressed, mapped and assured consistently across domains and layers
2. **AI enablement**—common approaches for data access, model lifecycle management (training, deployment, monitoring, retirement) and trustworthy AI operation, providing the foundations to integrate paradigms such as agentic AI and LLMs and enable cross-provider interoperability across agents and models to support emerging AI techniques
3. **Data handling and exposure**—treats data as a core asset through a general framework for modeling, describing, correlating, abstracting and sharing it across domains, vendors and use cases, with consistent policies for privacy and governance, and providing a trusted foundation for AI, DTs and autonomous agents to reason, learn and adapt
4. **Control loops**—common means for the dynamic composition, interaction and conflict resolution of control loops spanning RAN, FAN, core, transport and cloud, allowing sensing, thinking and action to adapt to changing conditions and objectives
5. **Security, privacy and trust**—security and privacy by design, and accountability mechanisms that work coherently across domains
6. **DTs**—seamless integration of DTs as an operational component of operations, with access to APIs and real-time data, so that simulation, prediction, validation and proactive assurance become part of normal operations
7. **Foundational management services**—builds on advanced, intelligent FCAPS, enabling prediction, prevention, optimization and adaptive operations. Examples include at-source data processing for efficient and richer insight, proactive alerts and tracking of configuration-plan transactions, while aligning service/slice lifecycle management and cross-domain orchestration to support flexible placement of intelligence and end-to-end autonomy
8. **Native exposure and developer readiness**—make service exposure first-class across and beyond network domains, enabled by common model-driven APIs and tools (e.g., SDKs and sandboxes) that make network capabilities natively accessible, and application management capabilities that allow packaging/portability of third-party intelligent controller apps and agents across platforms. Moreover, they should enable secure access to enrichment information from outside the network, further enhancing services and innovation.

Together the harmonized reference architectural framework and common enablers create the foundation for cognitive autonomy, allowing the network to reason, learn and adapt across domains while supporting flexible placement of intelligence and end-to-end autonomy. Building on this foundation, a harmonized 6G standardization landscape can turn cognitive autonomy from vision into reality and unlock new value across the entire ecosystem.

Now is the time to act

To realize coherent, seamless end-to-end operations and true cognitive autonomy, the industry must act now to harmonize standards and make them an integral part of the 6G system design from day one.

Standards organizations and industry stakeholders should align specifications, roadmaps and timelines, establish coordination mechanisms across major organizations, and jointly develop interoperable frameworks for key enablers, such as intent, AI, data framework and services, control loops, and DTs. Understanding the roles of various organizations is crucial for effective coordination to fully leverage their complementary contributions. As a technology and standardization leader with a strong presence across many industry fora, Nokia is well positioned in the lead of this harmonization effort, helping advance the journey towards cognitive and fully autonomous networks.

By achieving this level of harmonization, holistic end-to-end approach and common enablers, 6G has the potential to realize the ambition of an autonomous network at maturity level of 4/5 (high to full autonomy) and transform networks to embrace Self-X and Zero-X capabilities, unlocking innovation with resilient, high-quality services at massive scale, real-time agility, and at a more sustainable, lower overall cost.

Enabling cognitive autonomy is not solely about technology; it fundamentally requires significant changes in both processes and human mindset. Building trust is an ongoing learning journey, and, as trust deepens, the need for direct supervision will naturally diminish. Crucially, the human role is not disappearing but evolving. The shifting role of humans involves injecting, tuning and monitoring intents, which will require cultivating new skills, a challenge for which agentic AI, operating “on the loop”, can provide invaluable assistance, helping human operators maintain a continuous dialogue with the network.

6G offers a once-in-a-decade opportunity to create more unified and capable enablers within cellular networks, paving the way for greater network autonomy. Today, operators often have to scale back their ambitions due to practical constraints and limitations in the enablers. Current efforts focus on identifying these challenges so that 6G can provide the right enablers to overcome the constraints and empower operators to build Level 4/5 autonomous networks. Although the 3GPP and autonomous networking evolution timelines are not directly linked, the vision remains clear: enable autonomy by design. Cognitive autonomy cannot be reached if seen merely as an afterthought. “It needs a thoughtful approach anchored in purpose, powered by data, designed with simplicity, and infused with AI—where it makes sense and with the right oversight” [4].

Operators will continue to advance their journey toward cognitive and fully autonomous networks, shifting their focus from managing network operations to autonomous controlled operations. This evolution will extract greater business value, drive new revenue streams, and enhance end-user experiences, leveraging the network as a strategic platform to deliver innovative services and meet evolving business and customer needs.

Abbreviations

AI	Artificial intelligence
AIOps	AI for IT operations
FCAPS	Fault, configuration, accounting, performance and security
KPI	Key performance indicator
DT	Digital twin
LLM	Large language model
RAN	Radio access network
SDK	Software-development kit
Self X	Self-fulfilling, self-assuring, self-optimizing
SLA	Service-level agreement
SLE	Service-level expectation
SLO	Service-level objective
TCO	Total cost of ownership
UE	User equipment
Zero X	Zero wait, zero touch, zero trouble

References

1. Autonomous Networks Business Requirements and Framework v3.0.0 (IG1218)", TM Forum web site, 28 June 2024, https://www.tmforum.org/resources/introductory-guide/autonomous-networks-business-requirements-and-framework-v3-0-0-ig1218/
2. Nokia, "Autonomous networks," Nokia web site, viewed 10 Oct 2025. Online: https://www.nokia.com/autonomous-networks/
3. Riggs, A., "Full network autonomy through intent management, enhanced observability, knowledge graphs, and AIOps-driven closed loops," TM Forum web site, 25 Jul 2025. Online: https://inform.tmforum.org/research-and-analysis/proofs-of-concept/full-network-autonomy-through-intent-management-enhanced-observability-knowledge-graphs-and-aiops-driven-closed-loops
4. Nokia, "UNEXT: unified networking experience," Nokia web site, viewed 22 Oct 2025. Online: https://www.nokia.com/bell-labs/research/unext/
5. Mukherjee, S., Francini, A. and Montag, M., "UNEXT: Network-application symbiosis," Nokia Bell Labs white paper, 2024. Online: https://onestore.nokia.com/asset/214399
6. Nokia, "6G orchestration and automation: a system software view," Nokia Bell Labs white paper, 2024. Online: https://onestore.nokia.com/asset/213955
7. Boussard, M., "Extended compute services for a unified networking experience," Nokia Bell Labs white paper, 2024. Online: https://onestore.nokia.com/asset/214403
8. Nokia, "A unified networking experience in decentralized environments," Nokia Bell Labs white paper, 2024. Online: https://www.nokia.com/asset/214086/
9. Conte, A., "The role of knowledge and data services in unified networking experience technology (UNEXT)," Nokia Bell Labs white paper, 2024. Online: https://www.nokia.com/asset/214174/
10. Ramachandran, D., "LinkedIn post," LinkedIn, Jun 2025. Online: https://www.linkedin.com/posts/deeparamachandran1_dtw25-dtwignite-activity-7345146265929732096-EBNi

About Nokia

At Nokia, we create technology that helps the world act together.

As a B2B technology innovation leader, we are pioneering networks that sense, think and act by leveraging our work across mobile, fixed and cloud networks. In addition, we create value with intellectual property and long-term research, led by the award-winning Nokia Bell Labs.

With truly open architectures that seamlessly integrate into any ecosystem, our high-performance networks create new opportunities for monetization and scale. Service providers, enterprises and partners worldwide trust Nokia to deliver secure, reliable and sustainable networks today – and work with us to create the digital services and applications of the future.

Nokia OYJ
Karakaari 7
02610 Espoo
Finland
Tel. +358 (0) 10 44 88 000

Document code: CID215139 (November)