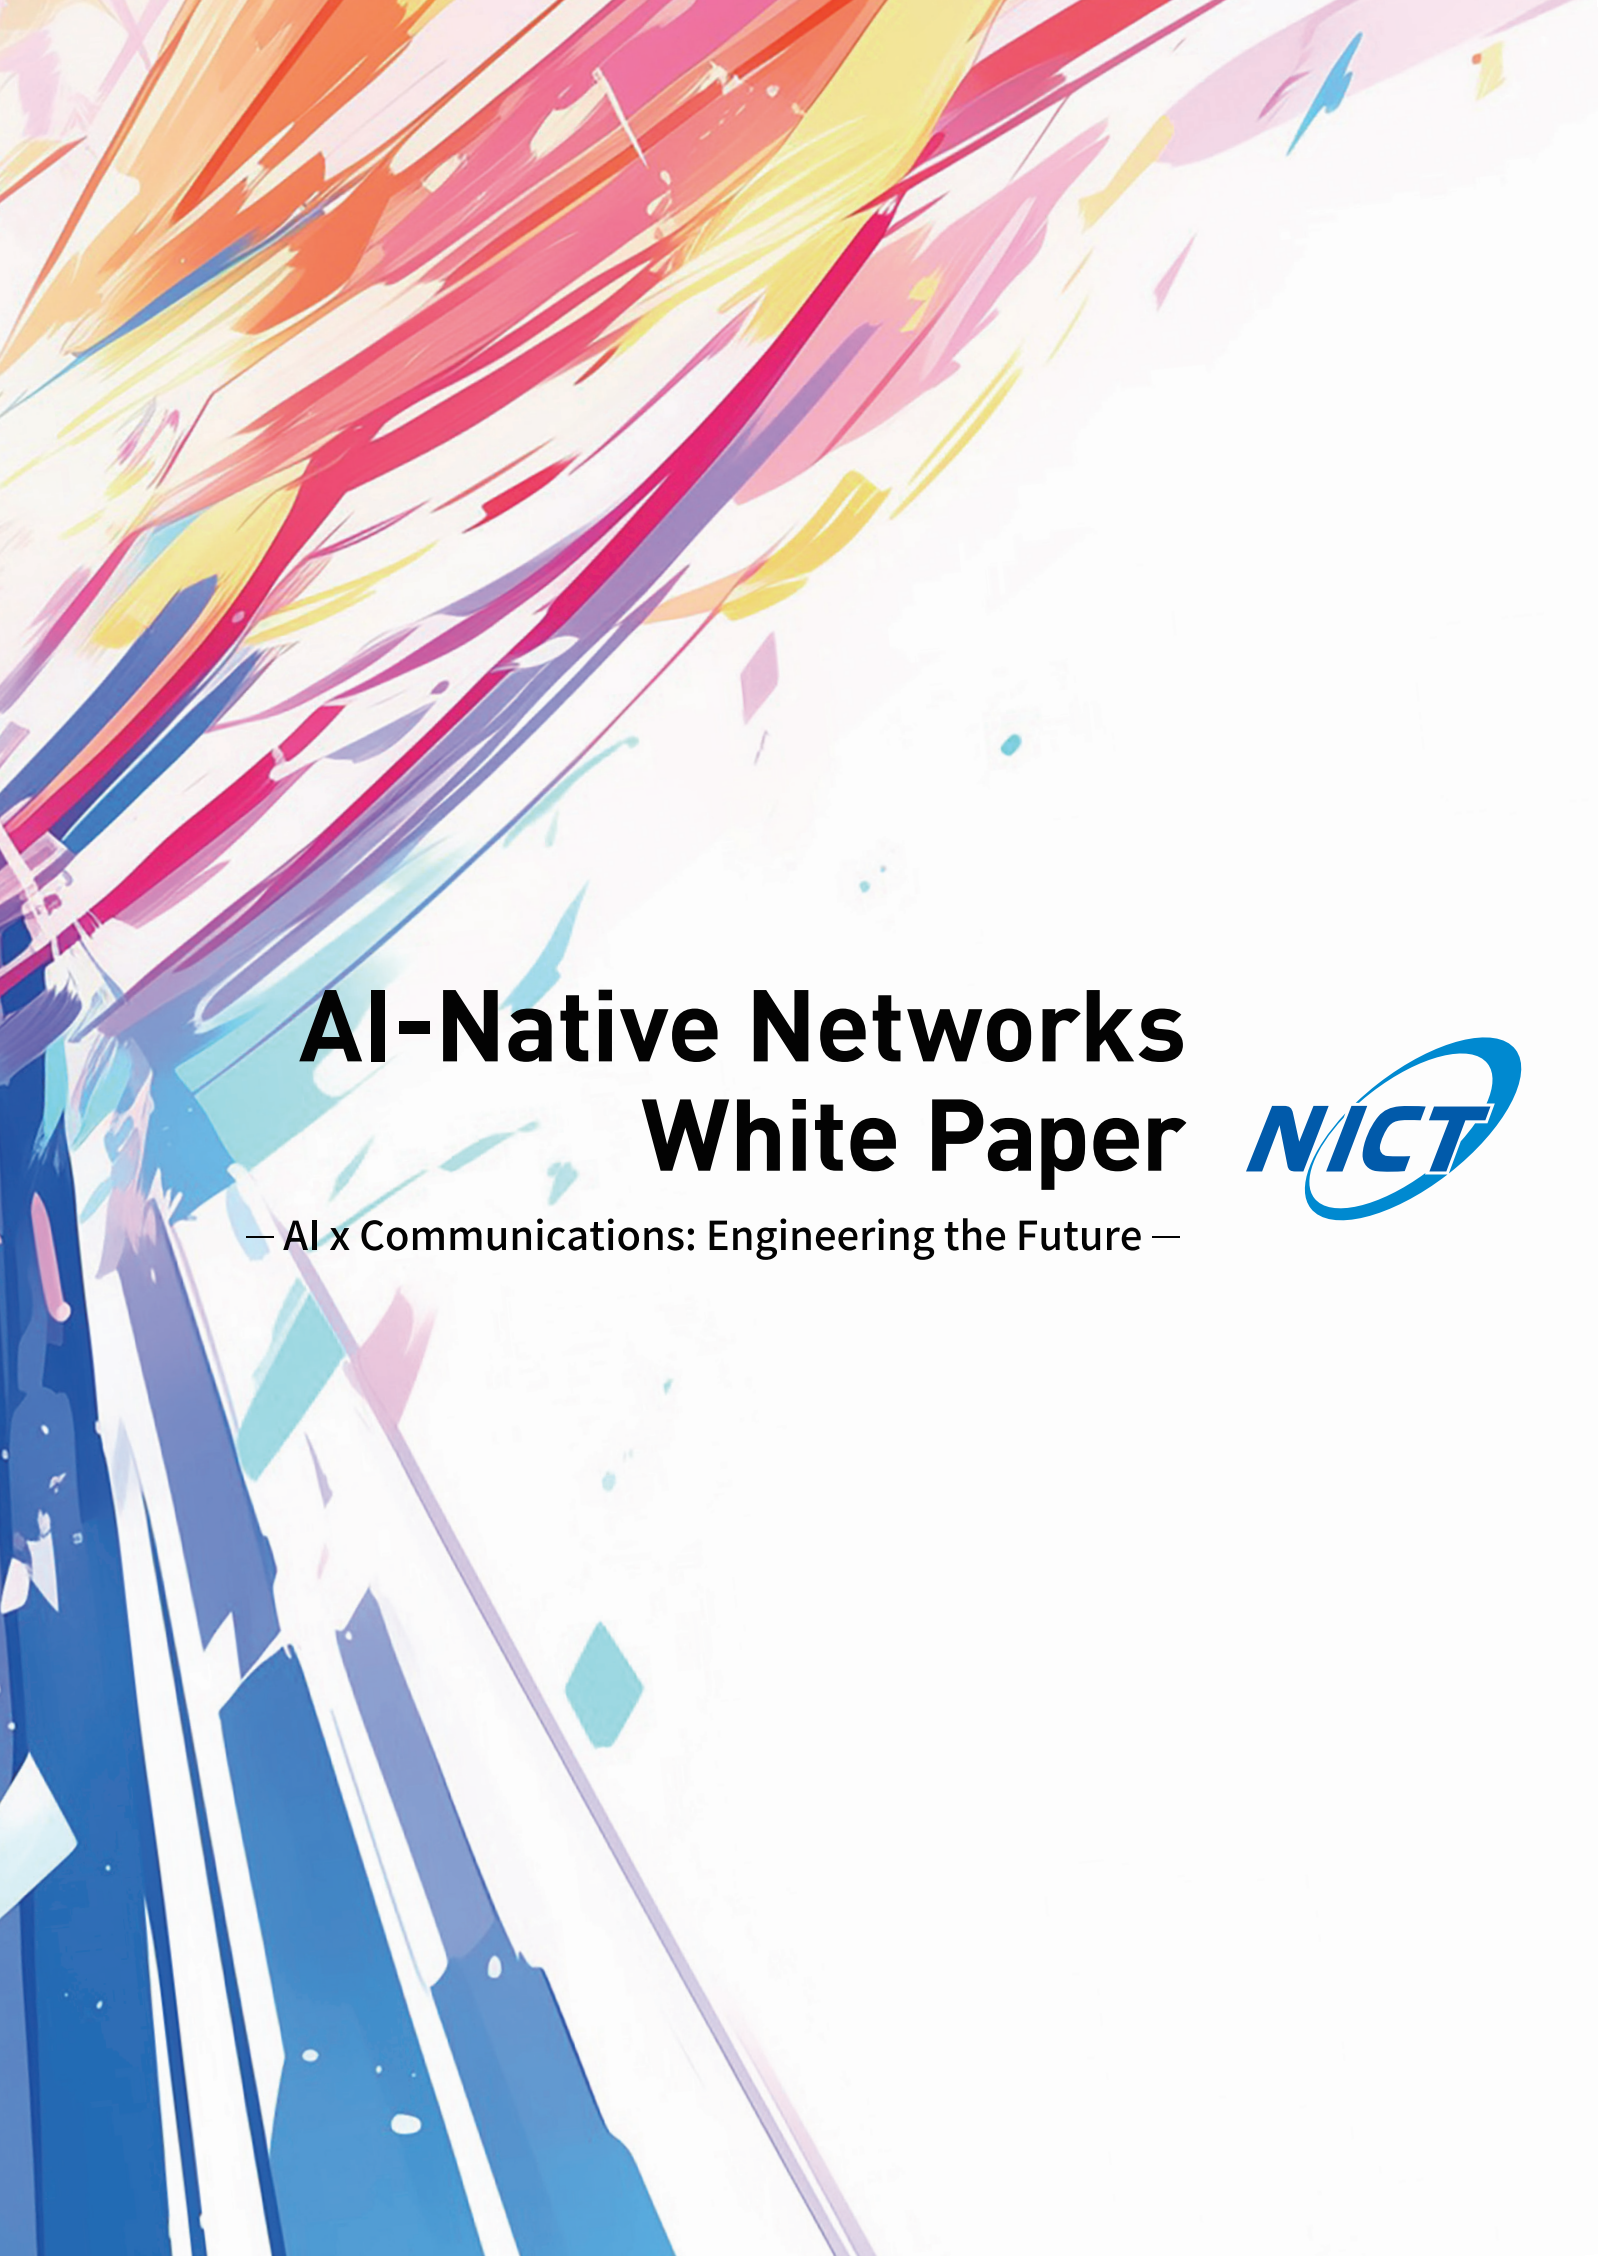

AI-Native Networks White Paper

— AI x Communications: Engineering the Future —

NICT

AI-Native Networks White Paper

— AI x Communications: Engineering the Future —

Executive Summary

The convergence of AI and communications is poised to transform the very role and value of communication networks. Partial deployment of AI within communication networks is already underway. The transition to what lies beyond – "AI-native networks," where AI and network lifecycles are fused and humans and AI dynamically determine the scope, location, and method of decision-making depending on the context – signifies not merely technological advancement but a transformation of societal infrastructure itself. Nevertheless, a framework for holistically explaining and verifying system-wide behavior remains underdeveloped.

Under this transition, three assumptions that have long underpinned communication networks as societal infrastructure are beginning to erode. First, due to interdependence among communication, computing, cloud, and power infrastructures, overall adequacy can no longer be explained through investment decisions made at the level of individual facilities (adequacy of capital investment/CAPEX). Second, because AI-driven control spans multiple infrastructures, the very premise of what should be evaluated and verified as a single unit is being called into question (premise for evaluation and verification). Third, because the division of roles between humans and AI may shift during operation, the boundary of responsibility under institutional frameworks cannot be determined in advance (boundary of responsibility). These are not hypothetical future concerns but challenges that are already becoming apparent in current implementation and operation.

This white paper makes this transformation visible and presents a foundation upon which stakeholders with differing perspectives can engage in dialogue based on shared assumptions.

Three Structural Indeterminacies – Where the Root of Change Lies

At the root of the three changing assumptions described above lie three structural indeterminacies that arise because AI's involvement in control decisions within communication networks renders the "location," "basis," and "agent" of decision-making indeterminable at design time (Chapter 2).

- Location and granularity of decision-making – Which controls to execute, where to execute them, and at what level of granularity are determined dynamically according to runtime conditions. Even for the same control decision, the balance among latency, accuracy, and cost varies depending on where it is executed: at the edge, in the cloud, or within the network.
- Basis and evidence trail of decision-making – The rationale for AI decisions becomes distributed across multiple layers, making tracing and explanation difficult. A structure emerges in which the question "why was this control action taken?" cannot be answered by any single entity.
- Decision-making agents – Because the division of roles among humans, rules, and AI may shift during operation, the answer to "who made this decision?" is no longer uniquely determinable. The locus of responsibility becomes structurally ambiguous.

These three structural indeterminacies are intertwined and simultaneously propagate across the domains of capital investment (CAPEX), evaluation and verification, and the boundary of responsibility.

Seven Questions Facing AI x Communications Societal Infrastructure

These eroding assumptions manifest in practice as the following seven questions.

- Design: To what extent should the scope of AI decisions and human involvement be treated as design parameters?
- Evaluation: By what criteria should the adequacy of AI-involved control be judged?
- Responsibility: Who bears accountability for the consequences of decisions, and to what extent?
- Cost: Who bears the costs of AI control deployment, operation, and verification, and how shall they be implemented?
- Trust: How can operators, users, and society trust AI decisions?
- Future sustainability: Which technical structures and institutional frameworks are sustainable?
- Definition: Is this fundamentally an issue of telecommunications, AI, or social institutions?

Why Conventional Verification Falls Short

In environments where AI is involved in control decisions, three assumptions that conventional evaluation and verification have implicitly relied upon – that the subject can be isolated, that conditions can be fixed, and that causality can be traced – no longer hold (Chapter 3). Control decisions are generated in a context-dependent manner, and providing identical conditions does not guarantee identical outcomes. Moreover, security threats specific to AI-native networks (training data poisoning, adversarial inputs, model inversion attacks, etc.) fall outside the scope of conventional network security frameworks. Verification confined to a single operator or a single technology cannot assess either the adequacy of AI decisions or their societal impact.

Proposal of This White Paper – Verification Collaboration Platform

It is difficult for individual entities to answer these questions on their own. This white paper proposes a "Verification Collaboration Platform" – a framework in which diverse entities bring their technologies together under shared conditions, conduct iterative comparison and verification, and incrementally develop evaluation metrics, institutional designs, standards, and other elements (Chapter 4). This platform goes beyond confirming technical adequacy. By sharing the framework of the location, basis, and agent of decision-making, verification items for new service deployment can be designed in advance, thereby limiting verification costs and the scope of trial-and-error. The platform is based on the assumption that telecommunications operators, AI providers, cloud providers, users, public institutions, and other entities each contribute their own conditions and expertise. It is fundamentally distinct from existing proprietary platforms in that it is configured as a neutral verification environment independent of any specific platform vendor.

- **From a design and implementation perspective** — A verification environment under comparable conditions becomes available. AI interoperability in multi-vendor environments can be demonstrated, enabling performance and cost comparisons across different control decision placement patterns (edge/cloud/in-network).
- **From an investment and management decision perspective** — The effectiveness of AI control deployment can be demonstrated incrementally. Comparative data between deployment and non-deployment scenarios accumulates, providing evidence to justify capital investment.
- **From an institutional design and standardization perspective** — Insights gained through verification are fed back into institutional design. In the EU, regulatory frameworks, including the AI Act, are already being developed; whether Japan can design its own verification frameworks and propose them internationally may influence the country's position in global rule-making.

Where to Begin — A Three-Phase Transition

Realizing the Verification Collaboration Platform is not about fully developing it all at once; rather, it proceeds through iterative cycles of phased demonstration and dialogue (Chapter 5).

- Phase 1 (present to near-term): AI control is partially introduced in limited use cases, and comparative data are accumulated through parallel operation with rule-based control. Concurrently, discussions among stakeholders on the requirements and configuration of the Verification Collaboration Platform begin, and consensus-building on shared conditions and evaluation metrics gets underway.
- Phase 2 (medium-term): Comparative verification of multiple entities and implementations is conducted using the Verification Collaboration Platform. Conditions for investment viability and patterns of breakdown in the boundary of responsibility are clarified based on empirical evidence, and feedback into institutional design begins at this stage.
- Phase 3 (medium- to long-term): The verified scope is gradually expanded, and the scope of application and level of autonomy of AI control are increased. Confirmation of safety and effectiveness through the Verification Collaboration Platform at each phase provides a rational basis for advancing to the next phase. This phase also aims to cultivate an ecosystem in which participation in the platform itself becomes a business opportunity.

This white paper does not present a complete solution. It calls on diverse stakeholders to take the structures and questions articulated here as a starting point and to incrementally form a shared framework through repeated cycles of verification and dialogue.

Table of Contents

Chapter 1: Why Start from "Decisions" — Rethinking Assumptions in AI-Communications Convergence

As the convergence of AI and communications advances,

there are assumptions that must be reexamined when discussing the future of networks, before choosing technologies or methods.

"Who decides what, where, and on what basis?" — Unless this question can be answered,

design, investment, and verification will all become misguided.

What is at issue is not AI making decisions autonomously outside the frameworks designed by humans.

Rather, the question is how far to delegate decision-making to AI within frameworks designed by humans, and how to maintain and update those frameworks.

Box: The Future of AI-Native Networks — When Two Lifecycles Converge

The "AI-native network" that this white paper addresses is not merely an extension of partial AI utilization, such as 3GPP NWDAF or O-RAN RIC. It envisions several possible futures for AI-native networks that come into view when the lifecycles of AI (training -> inference -> monitoring -> retraining) and communication networks (design -> deployment -> operation -> optimization -> decommissioning) converge (ITU-T FG-AINN, 2024).

For example, a user's AI inference request could be dynamically routed via wide-area networks to the globally optimal GPU cluster or edge node, with the execution locations for training, inference, and monitoring themselves being reallocated in real time according to network load and computing resource conditions. Furthermore, as the network continuously learns users' usage patterns, network configurations would be generated in reverse — not from specifications defined by designers but from actual usage. The boundary between design and operation dissolves, and the network becomes an infrastructure that continuously redesigns itself.

Moreover, the network itself could detect the emergence of unknown service categories from changes in traffic patterns and automatically generate corresponding APIs and bandwidth policies, thereby becoming the first to detect market changes. Beyond that, one can envision a world in which the AI of cloud providers trades computing resources, the AI of telecommunications operators trades bandwidth, and the AI of power providers trades electricity through real-time automated negotiation, while the network layer, acting as a third party, automatically records and verifies the history and evidence trail of decision-making. At that point, the network would no longer be a mere communication infrastructure but could serve as a trust infrastructure as well. Humans design the objectives, rules, and constraints, while decision-making and resource allocation within those boundaries are delegated to AI. When the two lifecycles converge, the very role of the network changes.

This white paper looks ahead to such a future and maps the changes that will unfold along the path toward it.

1.1 What is Happening to Communication Networks as Societal Infrastructure

In recent years, both domestic and international policy and research strategies have increasingly sought to reconceptualize AI-centered advanced technologies as well as information and communication infrastructure as an integrated foundation that spans research, social implementation, and operation. Behind this lies the fact that AI training, inference, and control are being embedded ever deeper into social systems, and that the interdependence with the information and communication, computing, and energy infrastructures that support them is rapidly intensifying.

Communication networks have long served as societal infrastructure, but their role has primarily been positioned as "infrastructure for providing stable transmission." Design and operation have been premised on the assumption that networks function under predefined requirements and conditions, and evaluation and verification have been conducted on a per-device or per-function basis, centering on metrics such as communication quality and availability. Although the possibility of dynamic optimization and adaptive control had been noted previously, it is fair to say that designs and institutional frameworks premised on such mechanisms being routinely operated, with control decisions being established distributively and dynamically at runtime, had not been adequately developed. The dynamic control and optimization that had been constrained under these limitations are now reemerging as realistic options.

The driving force behind this is the rapid advancement of AI technology. As the application of machine learning and generative AI progresses, communication has come to be positioned not only as infrastructure that supports AI training and inference, but also as a domain in which operations and control are themselves enhanced by AI.

Such changes are not confined to future visions; they have already begun, in part. For example, NWDAF (Network Data Analytics Function), standardized by 3GPP as an analytics function for 5G core networks, is being deployed as a mechanism that leverages AI for traffic prediction and anomaly detection[1]. The RIC (RAN Intelligent Controller) defined by the O-RAN Alliance is a framework that incorporates AI-based decisions into radio base station resource allocation and interference management, and multiple telecommunications operators and vendors are conducting demonstrations. Additionally, SON (Self-Organizing Network), introduced in the 4G era, has achieved automated base station parameter adjustment and fault recovery, and in recent years, its enhancement through the incorporation of AI/machine learning has been advancing.

At present, many of these frameworks remain at the stage of analysis, prediction, and supplementary optimization, and have not yet reached the stage where AI directly generates control decisions for immediate execution. However,

[1] For example, regarding architectures that assume AI processing distributed across terminals, edge, and cloud, discussions are being advanced within 3GPP's AI/ML-related studies (5G Advanced – Services and Architecture). In addition, the IETF is examining Computing-Aware Traffic Steering (CATS) as a framework that takes into account the state of computing resources when selecting communication paths. In these discussions, configurations are envisioned in which AI/ML applications present requirements (Intents) such as latency, bandwidth, computing resources, and energy efficiency, and the network dynamically considers and optimizes them.

the important point is that even with such partial deployment, the three structural indeterminacies that this white paper identifies – the shift from a deterministic control structure, in which the location and granularity of decision-making, the basis of decision-making, and the composition of decision-making agents can be predetermined at design time, to an adaptive and interactive structure – are already beginning to manifest, albeit to varying degrees.

For example, which control node uses the analysis results generated by NWDAF and how it differs with each deployment configuration, and the extent to which the decision-making basis of AI models used in RIC can be traced remains unresolved. This white paper takes the position that in order to advance such partial deployment safely and incrementally, it is necessary to organize the structure of control decisions and prepare verification frameworks in advance.

In this context, communication networks are coming to be recognized not merely as infrastructure for information transmission, but as a composite platform linked with data processing, computing resources, and service control. Coordination with cloud environments, data centers, edge computing infrastructure, and even execution infrastructure such as power and cooling facilities has become indispensable, and communication networks now function as part of societal infrastructure involving multiple technological elements and operational entities.

If these trends advance further, control and optimization in communication networks will transition from something determined entirely at design time to something premised on being continuously reconfigured according to runtime conditions. Control mechanisms that include AI will determine at runtime which control is performed, where, and when, while simultaneously taking into account multiple conditions such as communication quality, computing resource utilization, power constraints, service requirements, and surrounding environmental information, including human traffic and behavioral data. Moreover, the premises and policies underlying these control decisions will themselves continue to be updated during operation. Under such a structure, predicting and fixing network behavior at design time is no longer feasible. The reproducibility and ex ante verifiability that operators have long assumed will face growing pressure to be fundamentally reconsidered as design principles.

This growing uncertainty does not diminish the role of communication networks as societal infrastructure. Rather, on the premise that they will be operated while incorporating uncertainty, communication networks are increasingly expected to play the role of continuously linking control decisions and execution between the physical world and the digital world as a core foundation of societal structures represented by Society 5.0. In addition, for communication networks as societal infrastructure, perspectives beyond the traditional emphasis on connectivity and speed, including reliability, safety, explainability, and sustainability, are increasingly being recognized as evaluation criteria for societal impact. In particular, in systems involving AI, the possibility that unintended behavior or effects could surface as social issues has been noted, and situations where conventional technical performance evaluation alone is insufficient are becoming more common.

The convergence of AI and communications is no longer limited to the sophistication of network configurations or technological elements. As the lifecycles of AI and communication networks begin to overlap, questions that cannot be captured within conventional frameworks are emerging, namely, how network behavior is determined and what societal impact may result.

Building on this awareness, the following section organizes the changes that the convergence of AI and communications brings to communication networks.

1.2 Starting from Decision-Making — What Has Changed

When envisioning the future of networks in which AI and communications are converged, the question "who decides what, where, and how?" arises before any discussion of AI, network placement, configuration, or methods.

In environments where AI has begun to participate in communication control and operations, the overall behavior of the network is increasingly shaped not so much by which devices are connected and how but by **what choices are made, when, and based on what information**.

In conventional communication networks, most decisions about "what to do" were made in advance by humans at the design or operational planning stage. For example, policies such as under what conditions to switch routes or at what load level to intensify control were defined in advance as rules or configuration values, and at runtime, devices faithfully applied them. Because of this structure, there was almost no need to explicitly consider the selections made by devices at each layer during runtime as "decisions," and control decisions and execution were treated as temporally and functionally separate by design.

On the other hand, in environments where AI participates in network control and operations, this assumption does not necessarily hold. AI operates within the framework of objectives, constraints, and priorities established by humans at design time. However, at runtime, control content and policies are updated each time based on network state, traffic conditions, computing resource and power constraints, and past behavior. As a result, control decisions by AI manifest not as the mere selection of configuration values but as the selection of behavior itself – "which control to apply, to which target, and when."

Another important point is that such AI control decisions not only take network state as input, but are also constrained and shaped by the configuration of the network and computing infrastructure itself. AI processing, such as training, inference, and model updating, is strongly dependent on conditions including latency, bandwidth, computing resources, and power constraints, and these are not merely exogenous constraints but serve as the preconditions that determine what control decisions AI can generate. In this sense, the relationship between AI and communications must be understood not as a unidirectional relationship in which AI controls the network, but as an interdependent relationship that includes reverse causality, in which the network shapes the decision space of AI.

For example, even if control that prioritizes latency minimization is selected at a given point in time, when the impact on computing resources or other services becomes apparent, the control policy may be switched, and choices may be made to reduce processing accuracy or control frequency. Predefining all such control decisions as rules is difficult, and they are increasingly assumed to be made dynamically according to runtime conditions.

AI operates through the stages of training, inference, and output. Training is the stage of acquiring patterns from past data and corresponds to preparation for decision-making. Inference is the stage of applying acquired patterns to the current situation to generate output, and this output constitutes the control decision itself. However, the manner in which inference output functions as a control decision is not uniform. In the case of optimization / control AI, inference output directly becomes the control decision and is executed immediately. In the case of operational decision-support AI, inference output is merely a recommendation that is used by humans to make the final decision. In composite systems, control decisions are formed as the result of interactions among the inference outputs of multiple AIs, so a decision is not completed on the basis of the output of any single AI platform alone. This distinction is essential for

considering the attribution of decision responsibility and verification methods, and is discussed in detail from Chapter 2 onward.

In this white paper, AI is treated not as a single entity but is distinguished into three roles: an entity that utilizes communication resources, a decision-making agent that determines control policies, and a mechanism that executes decisions as control actions. When the main text refers to "AI as a decision-making agent," it denotes the logical agent involved in determining control policies. AI that utilizes communication resources and AI that translates decisions into execution can be found in conventional environments as well. However, what characterizes AI-native transformation is the incorporation of AI as a decision-making agent that generates and updates control policies based on learning into the cyclical structure.

Furthermore, in light of these changes, this white paper treats "decision-making" not as something attributed to a specific technology or entity, but as a functional process that directs network behavior. Decision-making as used here encompasses human decision-making, rule-based automated control, and AI-based inference and optimization processing, referring broadly to all acts of selecting or updating policies and actions. What is at issue is not the relative merits of individual control methods, but the structure: under what framework decisions are generated and applied, and how that framework itself should be designed and updated.

By taking "decision-making" as the starting point for examining structure in this way, it becomes possible to organize, across technological fields, why perspectives on verification and design that differ from conventional approaches are needed in networks where AI and communications are converged.

From Chapter 2 onward, taking this shift in the positioning of "decision-making" as a premise, the challenges that manifest in actual operational environments and the problems facing evaluation and verification frameworks will be discussed concretely.

Column 1: Why Did Decisions That Were Correct Within Each Domain Cease to Align?

Even before the convergence of AI and communications advanced, voices had intermittently been heard in operational settings saying, “everything should be working as designed, yet something feels off overall.”

A telecommunications operator says:

> “Network control is operating according to specifications and quality metrics are being met.”

An AI system manager says:

> “Model accuracy and stability have been confirmed in the test environment.”

A power and facilities manager also says:

> “Supply margins and redundancy are not a problem by design.”

They are each rational within their respective domains, and one cannot definitively say that any of them is clearly wrong.

However, when AI begins to connect these domains across boundaries as a decision-making agent, situations arise in which the “correctness” of each domain no longer aligns as a whole.

Moreover, even when such misalignment is noticed, it is unclear which entity should intervene. The communications side considers that “the AI’s decision may be the cause”; the AI side considers that “communication or power constraints may be the cause”; and the power side determines that “this is outside our jurisdiction.” Previously, problems were contained within a specific domain, so the personnel responsible for that domain could identify the cause and address it. However, in environments where AI generates and executes decisions across domains, the locus of the problem itself spans multiple entities, and it is not predetermined who should escalate and who should halt operations.

Such incidents tend to be handled individually as “minor malfunctions” or “unanticipated events.” However, when decisions are generated and updated at runtime and span multiple infrastructures, the problem is no longer a mere technical issue.

The inability to clearly answer the following questions after such incidents becomes a significant issue when considering accountability and investment decisions for societal infrastructure.

- Who made the decision
- Which decision had the impact
- Under which institutional or evaluation framework should verification be conducted

What this white paper discusses is not the merits or demerits of any specific technology, but rather what kind of societal challenges will arise if societal infrastructure continues to be operated in a state where explanation, verification, and the boundary of responsibility remain ambiguous.

Box: Organizing Technologies Through the Lens of Decision-making

In recent years, a diverse array of technologies related to AI has emerged, including federated learning, distributed inference, cooperative control, autonomous operations, and agent-based systems. These appear to belong to different domains — learning, inference, control, and operations — but when organized from the perspective of "decision-making," they can be understood as sharing a common structural transformation. Specifically, it becomes possible to compare across technological fields where decisions are made, at what granularity they operate, and how they are updated.

Examples of Organizing Technologies Through the Lens of Decision-making

Technology Domain	Technology Category	What Is Happening Structurally	Change in Decision-making
Federated Learning	Learning Method	Decision criteria (models) are updated at multiple locations and integrated via communications	Decentralization of decision updating
Distributed Inference	Inference Placement Optimization	Where decisions are made is not fixed; it varies with communication state and load	Fluidity of decision-making location
Cooperative Control / Distributed Control	Control Method	Each entity makes local decisions, and overall behavior emerges as the result of interactions	Multiplicity of decision-making agents
Autonomous Operations (AIOps, etc.)	Operational Automation	Execution results influence subsequent decisions, and decision criteria are continuously updated during operation	Continuity of decision-making over time
Agent-Based Systems	Implementation Form	Observations, decisions, actions, and feedback are cycled while interacting	Cyclicalization and autonomization of decision-making

1.3 Five Domains That Can No Longer Be Determined at Design Time

From the perspective of "who decides what action to take, where, and how," the "undefined domains" that newly become apparent through AI-native transformation are organized below.

In conventional communication networks, many preconditions were implicitly shared. Decisions were primarily made by humans at the design or operational planning stage, and at runtime, devices applied predetermined rules – this structure was the norm. Consequently, there was little need to explicitly define the decision-making agent, location, and timing, or the locus of responsibility for decision outcomes.

However, in environments where AI participates in network control and operations, these assumptions become difficult to sustain. As control decisions are generated and updated by AI at runtime and directly affect network behavior, matters previously taken as self-evident must be reexamined. Specifically, at least the following domains tend to remain undefined or ambiguous.

The first domain is the **locus of the decision-making agent**.

When AI is involved in decision-making, the decisions are not necessarily attributable to a specific device or a single system. Even if a supervisory control entity (master) is established, when control decisions are formed as the result of

interactions among processes distributed across multiple locations such as edge, network, and cloud, it is not necessarily apparent which entity should be considered to have "made" the decision.

The second domain is **treatment of the basis, history, and evidence trail of decision-making**.

Decision-making processes involving AI are generated based on multiple factors, including training results, runtime input data, resource constraints, and past decision histories. Consequently, the basis for decision-making is not consolidated in a single design document or log but is increasingly distributed across multiple layers and systems.

The third domain is the **boundary between control decisions and execution**.

When control decisions are generated at runtime, they are no longer processes separated from execution in terms of time or function. What counts as a "decision" and where "execution" begins are difficult to define with fixed boundaries.

The fourth domain is the **unit of evaluation and verification**.

When control decisions change dynamically and interact with constraints in the execution environment, it becomes difficult to isolate and evaluate a single device or control logic. What scope should constitute the evaluation target, and under what conditions behavior should be deemed adequate, cannot be fully organized within conventional frameworks.

The fifth domain is **how to handle cases where the adequacy or reliability of decisions is compromised**.

When control decisions are generated at runtime and formed through interactions among multiple elements and entities, they may be unintentionally distorted or subject to external interference. How to detect this, which entity should assume responsibility for it, and to what extent, is not obvious in advance.

These undefined domains are not problems of individual technical implementation per se, but structural problems arising from the transformation in the nature of decisions. Accordingly, they cannot be resolved within a single technological field or a single verification environment, and must be considered under the premise that multiple entities, layers, and constraints are involved.

In Chapter 2, the undefined domains surveyed in this section are analyzed in depth as three structural indeterminacies: "the location and granularity of decision-making," "the basis, history, and evidence trail of decision-making," and "the locus of the decision-making agent."

1.4 Problem Framing and Structure of This White Paper

This section organizes the problem framing and discussion structure of this white paper. As outlined in the preceding section (Section 1.3), AI-native transformation is creating situations in which the locus, basis, and boundary of responsibility for decision-making cannot be determined at the design stage. These structural indeterminacies do not manifest as a single issue but arise simultaneously across multiple layers – edge, communication infrastructure, cloud, and operations – and propagate into all three domains on which this white paper focuses: the adequacy of capital investment (CAPEX), system evaluation and verification, and clarification of the boundary of responsibility. Consequently, the challenges cannot be resolved locally through the maturation of individual technologies or

implementation-level ingenuity, and must be organized across multiple technological fields, including communications, AI, cloud, evaluation, and institutional design.

The problems addressed by this white paper are not specific to any particular AI technology. Methods such as online learning, which update at runtime, make complete prior verification fundamentally difficult; even with pre-trained or rule-based methods, since decision outcomes vary with input conditions, defining all behavior at design time is challenging. This white paper focuses not on the characteristics of individual methods but on the challenges that commonly arise across methods when AI is incorporated into the decision structure of societal infrastructure. The discussion in this white paper is generalized in a manner not dependent on any specific company or product, while drawing on publicly available technical documents and standardization discussions.

Based on the above, this white paper conceives of AI-native networks not as a single technology or product but as societal infrastructure that is realized through the coordination of multiple physical elements – edge, communication infrastructure, cloud, power infrastructure, and management and control systems. On that basis, it organizes how the relationships among these elements change when AI participates as a decision-making agent, and what structural questions newly arise.

Figure 1-1 is a reference diagram showing the physical components of the AI-native network addressed by this white paper; it does not depict any specific implementation or method, but aims to share the positional relationships among the physical elements that form the premise for discussion. Table 1-1 organizes the major components shown in Figure 1-1 and their roles, and Table 1-2 organizes the correspondence between the issues discussed in each chapter and the overall structure. When reading subsequent chapters, please refer to Table 1-2 as an overarching guide.

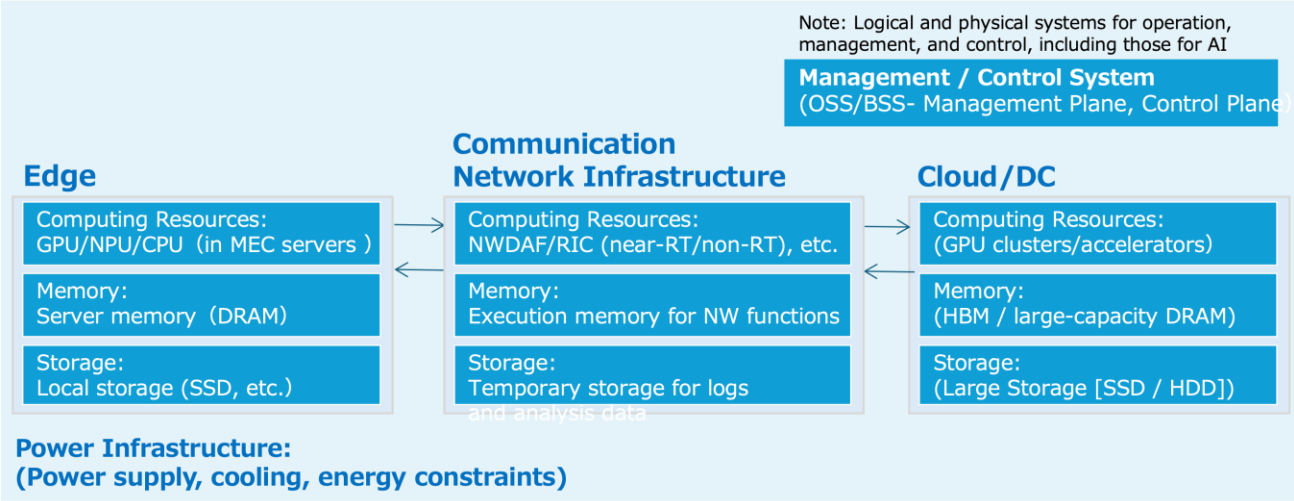

Figure 1-1: Example Physical Configuration of an AI-Native Network

Table 1-1: Classification of Physical and Functional Elements Comprising an AI-Native Network (Example)

Component	Primary Role	Included Functions / Examples
Communication Users/ Applications	Generation of service requests	Human usage, AI applications (inference, analysis, etc.), generative AI services (dialogue, content generation, etc.)
Edge	Low-latency processing/ local control	MEC servers (including computing, memory, and storage), near-RT RIC, edge AI inference, sensor processing

Communication Network	Connectivity/ transport/ control	Radio, core NW, traffic control, in-network computing (NWDAF, etc.)
Cloud/DC	Aggregated computing/ training	AI training, large-scale inference, management systems
Computing Resources	Provision of processing capacity	CPU/GPU/accelerators, memory, storage
Power/Energy	Execution constraint conditions	Power consumption, power control
Operations/ Management/Control Systems	Overall control/policy enforcement	Orchestration, operations
AI as a Decision-making Agent	Integration of decisions	Decision-making considering objectives and constraints

Table 1-2: Mapping of Questions Surrounding AI-Native Transformation to the Chapter Structure of This Whitepaper

Perspective/ Chapter	Chapter 2: Three Structural Indeterminacies	Chapter 3: Limitation of Operation and Verification	Chapter 4: Verification Collaboration Platform
Chapter Theme	**Decision-making location, basis, and agent are not predetermined**	**Structural indeterminacies become apparent in both operation and verification**	**Verification frameworks and the roles of entities**
	Sec. 2.1 Location and Granularity of Decision-making	**Sec. 3.1** Distribution of Locations and Operational Issues	**Sec. 4.1** "Correct Operation" Alone is Not Enough
	Runtime Variability of Location and Granularity	Operational Impact of Distributed Decision-making	Verification Perspectives for Three Structural Indeterminacies
	Sec. 2.2 Location of Basis and History	**Sec. 3.2** Difficulty in Tracing the Basis	**Sec. 4.2** Who Bears the Structural Indeterminacies
	Basis Distributed Across Multiple Layers	Loss of Reproducibility and Explainability	Operators, Vendors, Research Institutions, etc.
	Sec. 2.3 Overlap of Humans, Rules, and AI	**Sec. 3.3** Combined Effects and the Collapse of Verification Assumptions	**Sec. 4.3** Functional Requirements for the Verification Coordination Platform
	Agent Not Attributable to a Single Entity	Limitations of Isolation, Reproducibility, and Traceability	Four Functional Domains and the Value of Participation
Chapter Conclusions	Decision-making cannot be fixed at design time	Post-hoc verification is insufficient	Structural clarification enables verification
To the Next Chapter	---> What Happens? (Chapter 3)	---> How Should We Conceptualize it? (Chapter 4)	---> Where Should We Start? (Chapter 5)

This white paper does not present a completed solution or specific implementation methods for AI-native networks. Rather, its purpose is to make visible the structural challenges that cannot be avoided in future societal infrastructure, and to present a foundation for discussing the stance and frameworks with which these challenges should be addressed. It is hoped that the problem framing and structure organized in this chapter will serve as a starting point for constructive discussion and verification among diverse stakeholders, including industry, academia, and government.

Regarding frameworks in which AI participates in the control and operation of communication networks, ITU-R M.2160 (IMT-2030 Framework) positions AI as a design principle for 6G, and ITU-T FG-AINN is developing definitions and reference architectures for "AI-native networking." Furthermore, ETSI ENI is advancing AI-assisted network automation, ETSI ZSM is developing a reference architecture for zero-touch management, 3GPP is specifying AI functions through NWDAF (TS 23.288) and AI/ML for Air Interface (TR 38.843), and the O-RAN Alliance is promoting an open AI control platform through multi-layer RIC architecture.

This white paper does not address the individual interfaces or functional specifications handled by these standardization activities, but rather organizes in a cross-cutting manner the changes that arise when AI participates as a decision-making agent – the structural indeterminacies of the location, basis, and agent of decision-making – and positions itself as an attempt to make visible the structure of challenges commonly faced by the aforementioned standardization activities.

Chapter 2: What Cannot Be Determined in Advance — The Three Structural Indeterminacies

– Why the location, basis, and agent of control decisions cannot be fixed at design time –

In conventional communication networks, the location, basis, and agent of control decisions can be fixed at design time.
AI-native transformation causes all three of these premises to change simultaneously:
where decisions are made is unknown, the basis on which they rest is untraceable, and who made them is not identifiable.
This chapter elucidates how these "three structural indeterminacies" propagate into capital investment (CAPEX), evaluation and verification, and the boundary of responsibility.

The "structural indeterminacy" addressed in this white paper differs from uncertainty in the sense of probabilistically varying behavior; it refers to a state in which the very premises, such as the location, agent, and basis of decision-making, can no longer be uniquely defined at design time or evaluation time. This is not a matter of the possibility that the system may err, but rather of the structural change whereby "where, by whom, and on what basis decisions are being made" can no longer be ascertained in advance. When this white paper states that "AI generates decisions," this does not mean that AI autonomously makes decisions of its own accord. It is shorthand referring to the entire process by which the inference output of AI is adopted and executed as a control decision within a framework designed by humans.

In conventional communication networks, decisions were made primarily by humans or by fixed rules, and therefore the location, boundary of responsibility, and basis of decision-making were explicitly organized at the design stage, with operations and evaluation built upon those premises.

However, as AI begins to generate context-dependent control decisions at runtime and becomes embedded as part of the control loop, these premises are no longer self-evident. Control decisions become distributed and capable of dynamically shifting, not only across physical locations such as edge, communication network, and cloud/data center, but also across functional layers such as resource control, service control, and operations management.

In such an environment, humans, rules, and AI all participate as decision-making agents, with their relative weights unfixed and decision bases and histories generated and updated in a distributed manner. As a result, it becomes fundamentally difficult to predetermine at design time "where a control decision is made," "to whom the decision should be attributed," and "on what preconditions it is based." [2] Even in conventional networks, not everything can be prescribed in advance; however, there is a qualitative difference between cases in which the objects of prescription are finite and enumerable and those in which they themselves are subject to variation at runtime.

[2] For example, in an environment where a Near-RT RIC is performing real-time traffic steering, if communication quality degrades, it is difficult to determine whether the cause lies in the RIC's decision, in resource shortages on the core network side, or in a learning model updated on the cloud side. This is a typical example in which the location and basis of decision-making are distributed across multiple layers.

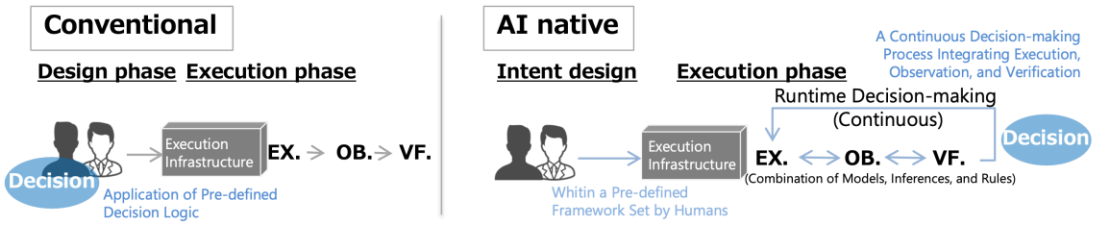

Perspective	Design with Decisions Fixed at Design Time	Design with Decisions Adaptively Determined at Runtime
Decision-making Agent	Humans/predefined rules	Within a pre-defined framework set by humans, humans, rules, and AI collaboratively decide
What Is Determined at Design Time	All Decision Contents below	Intent of decision-making: objectives, constraints, and priority ordering Example: establishing a state that satisfies both latency requirements and processing continuity
Decision Content	Concrete operations and configuration values Example: setting bandwidth to XX Mbps	Policy for Generating and Selecting Actions (Based on Objectives and Constraints Defined at Design Time) Example: Selecting Actions While Prioritizing Latency Minimization Within Resource Constraints
Decision on Processing Placement	Fixed in advance (device/layer) Example: this control is always performed by this device	Selected at runtime (edge/network/cloud) Example: first by the edge, next by the cloud
Form of Decision Logic	if–then branches → enumerable	Optimization- and inference-based selection → exploration of combinatorial solution spaces
Conditions for Decision-making	Enumerable states and thresholds	Combinations of states, constraints, and histories
Temporal Characteristics of Decisions	Fixed cycle	Adaptive to context (including event-driven behavior)
Location of the Basis of Decision-making	Centralized in design specifications and configuration values	States, constraints, and histories exist in a distributed manner
Evaluation Unit for Decisions	Single KPI, single function → pointwise evaluation is feasible	Multiple KPIs and overall system behavior consistency → need to evaluate the coherence of the cycle

Note: The "policy for generating and selecting actions" referenced in this white paper is distinct from orchestrator-level execution controls (such as workload placement and resource allocation). It refers to the decision-generation layer in which intent is translated into concrete control actions based on runtime state and constraints.

Figure 2-1: Difference between decisions established at design time and decisions that are continuously formed at runtime (Left: Conventional communication network / Right: AI-native network)

This figure organizes how the decision-making agent, decision content, temporality of decisions, and evaluation unit change between conventional communication networks and AI-native networks. Sections 2.1 through 2.3 of this chapter describe, in order, how structural indeterminacy arises with respect to each perspective shown in this figure.

2.1 Where and at What Granularity Are Control Decisions Made?

As illustrated in Figure 2-1, in environments where AI-native transformation is advancing, it becomes difficult to predetermine at design time the location at which control decisions are made or their temporal and spatial granularity.

In conventional communication networks, control decisions are placed at specific layers and control points, and their granularity is treated as a premise defined in advance. However, as AI becomes involved in decision-making, control

decisions are no longer defined as simple input-output correspondences but are generated as inferences that simultaneously consider numerous pieces of state information and constraint conditions.

The factors considered include not only radio conditions and traffic, but also the execution location of the decision (edge/network/cloud), the granularity of control, and constraints such as latency, computing resources, power, and operational policies.

As a result, the decision space expands combinatorially, making it practically infeasible to enumerate and verify all behaviors at design time.

In AI-native environments, the fact that **the location and granularity of decision-making are determined at runtime is itself becoming a precondition.**

This issue of the location of decision-making also relates to the choice of where within the network to perform AI training and inference. Table 2-1 organizes the deployment patterns for training and inference that are currently envisioned.[3]

Table 2-1: Deployment Patterns for AI Training and Inference (Reference)

Location	Training	Primary Training Data	Inference
Cloud / OAM/SMO	Load prediction, energy-saving optimization, slice load prediction, etc.	Traffic statistics, KPI history, network configuration information, etc.	Load prediction, energy-saving decisions, slice resource allocation, etc. (global control based on trained models)
Edge / Near-RT RIC	Online reinforcement learning-based fine-tuning (resource allocation, etc.)	Real-time RAN metrics, CQI, RSRP, etc.	Traffic steering, resource allocation, DU parameter optimization
Base station (gNB)	Localized training for energy saving and load balancing	Cell load, neighboring cell information, UE measurement reports, etc.	Energy saving (cell ON/OFF), load balancing, mobility optimization
Terminal (UE)	Federated learning client (terminal mobility prediction, gradient sharing for video recognition, etc.)	Terminal local data (location, sensors, video, etc.)	Terminal mobility prediction, video/audio recognition, V2X control
NWDAF (Core NW)	Federated learning (anomaly detection, quality prediction, etc.)	NF load, UE events, slice utilization status, etc.	Anomaly detection, quality prediction, UE location prediction
Router / Switch	No track record	—	(Limited) Encrypted traffic analysis, anomaly detection, etc. (training performed in the vendor cloud)

2.2 Where Do the Basis and History of Control Decisions Reside?

As illustrated in Figure 2-1, in AI-native networks, the assumption that prevailed in conventional communication networks – that the basis and history of decision-making can be organized after the fact in a single location and format – becomes difficult to sustain. Control decisions are no longer generated just once based on predefined rules; instead, they are dynamically generated and updated based on state information observed at runtime, past decision

[3] This table was compiled with reference to 3GPP TR 37.817, O-RAN specifications, 3GPP TS 23.288 (Rel-18/19), 3GPP TR 23.700-80, and related documents.

results, and constraint conditions such as computing resources and power. To organize this change, it is necessary to view the relationship between AI and communications from two perspectives.

AI for Network (network control and optimization by AI) and Network for AI (network provisioning as an AI execution platform) are the two representative axes. In this white paper, from the perspective of where the basis of control decisions resides, these two are reframed as follows.

- AI for Network: the side responsible for executing decision results as concrete control actions
- Network for AI: the execution environment that provides the state and constraint conditions enabling decisions and their execution

While the relationship – AI controls the network / the network supports AI – remains unchanged in both cases, from the AI for Network perspective, control decisions generated by AI are applied to the network as traffic control and resource allocation. The objects of such control include the securing and allocation of communication, computing, and power resources that support the training and inference of AI itself. This constitutes the Network for AI aspect, and the "AI" receiving these resources includes AI that functions as communication users and applications. Furthermore, the changes in communication quality, traffic, computing resource utilization, and power consumption that arise as a result of control become inputs to AI decisions once again.

In this structure where the basis and history of decision-making are distributed in a circular manner, the question inevitably arises as to who, or what, ultimately makes the decision; that is, the decision-making agent itself comes into question.

This circular structure is shown in Figure 2-2. AI generates control decisions using the network state as input, and the results of those decisions alter the network state, which in turn becomes input to AI once again. The state of the execution environment that arises as a result of this control is fed back to the decision side and becomes the precondition for the next decision. In this way, a structure is formed in which control decisions, execution, and constraint conditions are linked in a circular manner.

In terms of implementation, control decisions are not necessarily generated by a single entity at a single location. In some cases, decision-making and execution are completed within a single AI; in other cases, multiple AIs render different decisions in parallel; and furthermore, the objective function or optimization target itself may be updated during execution. In other words, "decisions" are, in practice, distributed, fluid, and overlapping.

As such control decisions come to constitute a circular structure, the basis and history of decision-making can no longer be contained in a single log. As exchanges, such as the collection of state information, transmission of decision results, and feedback of execution results, are repeated, the basis and history become dispersed across these processes. At this point, changes also arise in the traffic structure within the network. In addition to traffic that was conventionally generated primarily for the provision of communication services, traffic to support this circular decision process begins to occur in a layered manner.[4] As AI, cloud, and mobile networks become more tightly coupled in the future, this trend will intensify.

[4] Recent reports indicate that traffic within and between data centers is growing rapidly, with AI/machine learning workloads identified as one of the primary drivers (Cisco, 2023; Google, 2019). Furthermore, the proliferation of generative AI and AI assistants has raised the possibility of increased uplink-dependent communications and bursty, difficult-to-predict traffic characteristics in mobile networks (Nokia, 2024; NVIDIA, 2025). In addition,

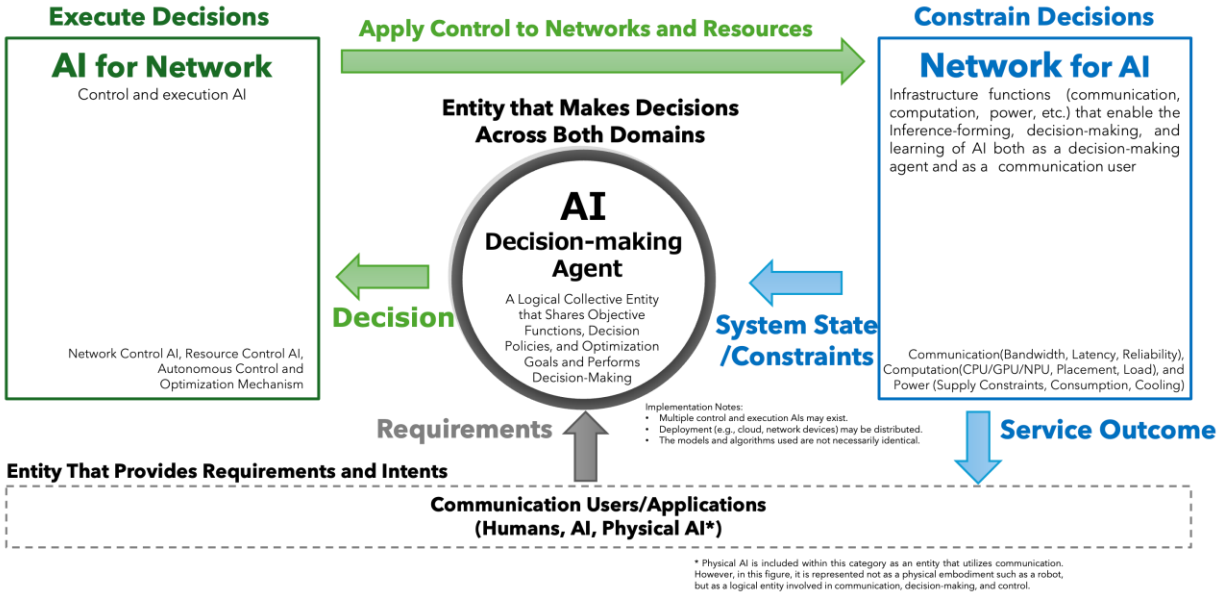

Figure 2-2: Conceptual diagram illustrating the structural factors behind the distributed nature of decision-making basis, history, and related information

This figure conceptually organizes three relationships in AI-native networks: AI as the decision-making agent, the side responsible for executing decision results as control actions (AI for Network), and the execution environment that provides the state and constraint conditions enabling decisions and their execution (Network for AI). This structure represents both AI for Network and Network for AI as part of the same decision-to-execution flow.

This white paper organizes its discussion around the axis of decision-making not to ignore this complexity, but rather to systematically capture it. Posing the questions "where are decisions generated, on what basis, and who assumes responsibility for them" makes it possible to provide a common organizing axis for the overall picture of distributed and fluid control.

The positioning of AI as a decision-making agent overlaps with the concept of Agentic AI, which has been advancing rapidly in recent years – namely, AI that autonomously executes tasks to fulfill a given intent.[5]

industry reports have observed a trend toward higher uplink ratios relative to the traditionally downlink-centric mobile traffic structure, driven by the influence of AI services (Ericsson, 2025). These trends indicate that the fundamental design and control assumptions of communication networks are undergoing transformation.

[5] TM Forum IG1230 (Autonomous Networks Technical Architecture) presents a structure in which Intelligent Agents controlling each layer make autonomous decisions based on Intents. Moreover, technologies are advancing in which LLM-based Agentic AI autonomously invokes tools via the Model Context Protocol (MCP); and Agent-to-Agent for Telecom (A2A-T), specialized for telecommunications management, has also been proposed within the same forum.

Column 2: No Problems by Design — Why Changes Still Arise in Operations

In the history of communication networks, “decision-making” had long been excluded as a central topic of discussion. This was because networks were assumed to operate in accordance with specifications and rules defined at design time, and it was presupposed that runtime behavior could be predicted deterministically.

The background to this change lies in the growing complexity of what networks handle and the increasing sophistication of control. The diversification of traffic, the variability of radio environments, and the escalation of service requirements have made it infeasible to define all situations as rules in advance. To address these challenges, approaches that use AI to select optimal control actions at runtime have begun to be introduced.

As a result, decision-making that was previously confined to the design stage has come to occur dynamically during operations. Decision-making has ceased to be a fixed premise and instead functions as a context-dependent element that can no longer be ignored when explaining, tracking, and verifying network behavior. The reason “decision-making” is problematic today is not the introduction of AI per se, but rather the transition to a structure in which decisions are continuously generated within the network.

2.3 Who is the Decision-Making Agent? — The Overlap of Humans, Rules, and AI

Even when AI is involved in the decision-making process, the criteria and permissible ranges of decisions are prescribed by frameworks defined in advance by humans. However, it becomes difficult to attribute individual control decisions generated at runtime under those frameworks to a specific single agent. The background to this lies in a structure whereby control decisions are generated not by a single entity in isolation, but through overlapping layers of human-defined policies and operational rules, rule-based control, and inference processing by AI.[6] Moreover, the configuration of involvement changes depending on the situation, time of day, and execution conditions.

In addition, policies established by humans are not directly reflected in terminal-level control decisions. High-level policies are translated at the operations layer into concrete rules and thresholds, and then interpreted at the execution layer as AI inference parameters and constraint conditions. Through this multi-stage translation process, the intent of the original policy may be reduced or altered, and verifying whether the final control decision is consistent with the original policy itself becomes a challenge. Such multi-stage policy translation is already becoming a design premise in, for example, 3GPP intent-based management and the multi-layer RIC architecture of O-RAN.[7]

As organized in Section 1.2, the manner in which AI inference output functions as a control decision is not uniform. In some cases, the output is executed directly, in others, it remains a proposal for humans, and in yet others, the outputs of multiple AIs can also interact to form a decision. Furthermore, AI does not assume a single role within the network; it may be implemented in a distributed manner as multiple mechanisms with different roles, such as AI as a communication user, AI that determines control policies (decision-making agent), and AI in the control and execution system that translates decisions into action. For this reason, the very act of asking "who made the decision" needs to be redefined not as the identification of a single agent in the conventional sense, but as a description of the composition of roles.

Moreover, under such composite decision structures, when a decision produces an unintended result, it is no longer clear to which entity the responsibility should be attributed. Distinguishing whether the issue lies in the framework defined by humans, in the inference output of the AI, or in the constraints of the execution environment itself becomes difficult, and the locus of responsibility cannot be organized without considering the decision structure as a whole. This point is directly connected to the discussion in Chapter 4 on the Verification Collaboration Platform and the division of roles by entity.

---

[6] As an implementation example of AI's involvement in control decisions for Radio Access Networks (RAN), studies and demonstrations in the context of AI-RAN (Artificial Intelligence Radio Access Network) have been advancing in recent years. AI-RAN is a concept that encompasses AI-for-RAN, which enhances radio control through AI; AI-and-RAN, enabling the coexistence of RAN and AI workload on computing resources; and AI-on-RAN, which utilizes RAN as an execution platform for AI applications. As an example, demonstrations of radio resource optimization using the RAN Intelligent Controller (RIC) within the O-RAN architecture have been reported (O-RAN Alliance, 2021–2024; AI-RAN Alliance, 2024).

[7] In mobile networks from the fifth generation onward, standardization efforts are progressing to expose network functions as APIs for dynamic utilization by applications and control logic. Within 3GPP, API exposure through the Network Exposure Function (NEF) and Policy Control Function (PCF) is specified, built upon the Service-Based Architecture (SBA) of the 5G core (3GPP TS 23.501, TS 23.502). In addition, ETSI is advancing the conceptual organization of Intent-based Networking Management, where discussions are underway on the direction of delegating high-level requirements (Intents) in network operations to control logic such as AI. These developments are positioned as premises for design and discussion regarding the partial transition of network control from manual configuration to automation and autonomy.

The changes in the decision structure described above are organized in Table 2-2, showing the correspondence between the location at which decisions are generated and the three structural indeterminacies.

Table 2-2: Correspondence Between Decision Generation Location and the Three Structural Indeterminacies

Agent / location of decision-making	Representative decision items	Indeterminacy of location	Indeterminacy of basis	Indeterminacy of agent
Human (designer/operator)	Configuration policies, policies, operational rules	Limited to design and operational stages	Traceable through documents and logs	Clear as a single agent
Rule-based control	Route selection, bandwidth control, failover	Control points fixed in advance	Conditions and results correspond	Attributable to the rule designer
AI (control / support)	Resource allocation, parameter tuning, operational proposals	Dynamically determined at runtime	Dependent on inference internals, and difficult to trace	Involvement of humans and AI becomes multi-layered
NW: AI-request-driven	Low-latency path assurance, bandwidth reservation, computing resource control	NW side adapts to AI requirements	AI requirements and NW constraints intersect	Requesting agent and executing agent are separated
Composite system (AI x NW x infrastructure)	Computing resource allocation, control frequency adjustment, control under energy constraints	Distributed across multiple layers	Decision basis is multi-dimensional	Impossible to identify a single agent alone

Note: This table is intended to organize the structure in which decision-making agents coexist, and captures changes in communication networks from a different perspective than maturity-level frameworks for autonomous networks (e.g., TM Forum Autonomous Network Levels).[8]

[8] See, for example, various recommendations regarding the TM Forum Autonomous Networks Manifesto.

Column 3: How is “Handan” Translated into English? — And Why it Feels Incongruous

The Japanese term “handan” (decision / decision-making/ judgment) used in this white paper does not correspond to a single English word. It overlaps with decision (the act of choosing a policy from multiple options), inference (the process by which an AI model generates output from input), and control (the execution of system operations based on a chosen policy), yet it does not fit into any of them. In conventional networks, these were separated into the distinct stages of design, inference, and execution, and “handan” referred to a one-time act of decision-making by a specific entity at a specific point in time (decision understood as a point).

However, in AI-native networks, observation, inference, execution, and feedback become a cyclically repeating process (a decision-making process in the form of a loop), and the meaning of “handan” itself has changed. This change is the cause of the sense of incongruity produced by the word “handan.” When readers assume the conventional “decision understood as a point” while reading, they find descriptions such as “the location of decision-making changes dynamically” or “the basis of decision-making is distributed” to be unnatural. The reason this white paper deliberately continues to use “handan” (decision / decision-making) is to make this circular structure visible and to treat it as a subject of discussion.

2.4 Practical Consequences of the Three Structural Indeterminacies

This chapter organizes the three structural indeterminacies that arise when AI becomes involved as a decision-making agent as: the location and granularity of decision-making; the basis, history, and evidence trail of decision-making; and the locus of the decision-making agent.

Because the location and granularity of decision-making fluctuate at runtime, the appropriateness of capital investment cannot be determined in advance. Because the basis of decision-making is distributed across multiple layers, the premises underlying conventional standalone evaluation and reproducibility verification break down. Because decision-making agents are configured in a composite and context-dependent manner, the locus of responsibility when failures or quality degradation occur cannot be uniquely identified.

In other words, the three structural indeterminacies organized in this chapter explain why these three practical challenges that cannot be avoided in the operation of societal infrastructure arise, from the common root of changes in decision structure. The key point is that these three challenges do not occur independently, but rather manifest simultaneously as the structural indeterminacies of the location, basis, and agent of decision-making become intertwined with one another.[9]

This organization based on the axis of decision-making is not in conflict with the entity-based organization (developer, provider, user) adopted by existing AI governance guidelines; rather, it serves a complementary role. While the entity-based organization indicates "who should do what," this white paper's decision-axis organization clarifies "what is decided and how." By combining the two, it becomes possible to discuss more precisely the locus of responsibility in cases where the same entity is involved in multiple stages of decision-making.

Chapter 3 organizes how these structural indeterminacies manifest as problems in evaluation and verification within real-world societal infrastructure.

[9] TM Forum IG1339 identifies fault recovery and optimization as high-value scenarios for automation level 4, and it is anticipated that autonomous operation will progress starting from cases where the definition of a single KPI is relatively straightforward, such as recovery time and cost. On the other hand, scenarios related to Planning remain at an undeveloped stage.

Chapter 3: What Happens in Real-World Operations, and Why Does Conventional Verification Fall Short?

– Examining Operational Fraying and the Limits of Evaluation and Verification in an Integrated Manner –

When control decisions become distributed across terminals, networks, and cloud,
what happens at the operational level, and why can conventional evaluation and verification no longer cover them adequately?
This is not due to operations becoming more difficult, but rather
to the fact that operational fraying and the limits of verification both arise from the same structural origin.

This chapter examines how the three structural indeterminacies identified in Chapter 2 are beginning to manifest as operational challenges in real-world societal infrastructure, and discusses in an integrated manner why conventional evaluation and verification frameworks are ceasing to function adequately. The issues addressed here extend beyond the implementation of individual technologies or operational workarounds; they emerge as problems that directly affect business decisions themselves, including capital investment planning, resource allocation, and service quality evaluation for telecommunications operators.

The following sections address, in order, the operational distortions produced by the distribution of the location and granularity of decision-making (Section 3.1), the structure by which the distribution of the basis of decision-making renders explanation and tracing infeasible (Section 3.2), and the structure by which constraints on the execution environment distort decision outcomes (Section 3.3). On this basis, the limits of conventional evaluation and verification frameworks are examined (Section 3.4).

3.1 What Happens in Practice When the Location and Granularity of Decision-Making Become Distributed?

As discussed in Section 2.1, in environments where AI is involved, the location and granularity of decision-making are determined at runtime. This discussion is directly linked to capital investment decisions regarding which layer – edge, network, or cloud – should receive investment to achieve the desired effect when introducing AI-based control. In real-world operations, humans, rule-based controls, and AI coexist as decision-making agents across multiple layers such as edge, network, and cloud, and structural distortions emerge as the location, granularity, and timing of decision-making fall out of alignment. It should be noted that this distribution extends not only to physical placement but also to protocol layers (from the physical layer to the application layer); however, this white paper focuses its discussion primarily on the axis of physical placement.

The time scale at which control decisions are made is also an important factor. In operational domains requiring millisecond-level responses, human intervention is impractical, and automated execution by AI becomes a prerequisite. Meanwhile, at scales of seconds, minutes, or hours, coordination between humans and AI or human intervention is anticipated. The same applies to spatial granularity: cell-level control, area-level control, and network-wide resource

allocation each involve different information requirements and optimization objectives. In this manner, decision-making agents are differentiated according to the time axis and spatial granularity, and operational design must take these differences as a given.

Table 3-1 organizes the constituent elements of the operational distortions caused by control decisions distributed across multiple layers and locations. When independent control decisions are made simultaneously at different layers for the same service, inconsistencies in control intent and effect can arise. Here is a concrete example: a telecommunications carrier's edge AI performs bandwidth control in response to a traffic surge, while a cloud operator's AI simultaneously offloads inference processing to the edge due to GPU cluster power constraints. As neither party accounts for the other's actions, load concentrates on the edge node. This example demonstrates that the capital investment decision of whether to augment edge capacity depends on the cloud-side's AI policies and that investment decisions cannot be grounded exclusively in single-layer evaluation. While inter-operator SLAs can define constraints, static agreements cannot address all situations in an environment where each entity's decision logic continues to change based on learning and state. Furthermore, because AI decisions depend on runtime state information and internal states, ostensibly identical conditions may produce different decision outcomes at different points in time.

Table 3-1: Constituent Elements of Operational Distortions Caused by Decisions Distributed Across Multiple Layers and Locations

Perspective	Primary distortion	Specific elements involved
Decision-making location	Distribution of decision-making location	Ambiguity of decision-making agents Unclear authority boundaries Stratification of control loops
Decision-making granularity	Differences in decision-making granularity	Differences in objective functions Divergence between local optima and global optima Discrepancies between service-level and infrastructure-level decisions
Timing of decision-making	Differences in the timing of decision-making	Misalignment of control cycles Perception gaps due to temporal differences in state recognition Feedback delay
Decision outcome	Non-uniqueness of decisions	Dependence on internal state Dependence on learning state updates Dependence on execution environment and time

3.2 Why Does Tracing the Basis of Decision-Making Become Infeasible?

As discussed in Section 3.1, in AI-native networks, control decisions are generated in a distributed manner across multiple layers and locations. While this structure enhances the flexibility and adaptability of control, it also makes it extremely difficult to explain and trace the basis of decision-making after the fact. This is directly connected to the structuring of the boundary of responsibility – specifically, whether it is possible to identify whose decision was the cause of a failure when one occurs.

Table 3-2 organizes the distributed structure of the basis of decision-making. In concrete operational scenarios, it becomes difficult to immediately answer the question "Why was this control action taken?" during failure response.

For example, if the load concentration on an edge node described in Section 3.1 materializes as a failure, it is necessary to determine whether the cause lies in the telecommunication carrier's bandwidth control, the cloud operator's offload decision, or in the coincidence of timing between the two. However, the decision logs of each AI are recorded in different formats and at different levels of granularity by each operator, and retroactively cross-referencing the input states at the time of each decision is itself far from straightforward. When short-cycle control decisions are combined with learning states and policies that are updated over longer cycles, the post-hoc disentanglement of causal relationships becomes even more difficult. Moreover, when an AI's learning state is updated during operation, ensuring reproducibility under identical conditions also becomes challenging.

What is required is not only real-time monitoring but also the design of audit trails that enable the basis of decision-making to be explained after the fact.

Table 3-2: Distributed Structure of the Basis of Decision-Making in Networks Where AI Is Involved in Decision-Making

Element of decision-making basis	Primary location / point of generation	Characteristics	Challenges in explaining and tracing
State information	Edge, network equipment, cloud	Dynamically acquired at runtime	Acquisition timing and granularity are not aligned
Learning results	Internal to the AI model	May be updated during operation	Difficult to reproduce past states
Past decision history	Distributed logs, control history	Multiple decisions are chained	Correspondence with decision units is unclear
Execution environment conditions	Computational resources, latency, power constraints	Varies depending on conditions	Decision premises change after the fact
Policies and constraints	Human, operational rules, configuration	Updated over long cycles	Prone to divergence from runtime decisions

Note: This table organizes the structure in which the basis of decision-making is not consolidated in a single location but exists in a distributed manner across multiple elements. The simultaneous actions of these elements make it difficult to explain decisions and trace them after the fact.

The key point here is that this is not an ethical or abstract discussion about AI decisions being "inexplicable." In real-world operations, the inability to explain or trace decisions manifests as practical problems, such as delays in the failure response, difficulties in building consensus among stakeholders, and the inability to share verification results. The very structure in which the basis of decision-making is distributed imposes a new burden on operational and verification processes.

3.3 When Correct Decisions Still Produce Distorted Outcomes — Combined Effects of the Three Structural Indeterminacies

Even when the control decisions of AI for Network are rational, the resource constraints of the infrastructure on which they are executed can distort the outcomes. Sections 3.1 and 3.2 addressed how the decision logic of AI for Network (the side where AI controls the network) becomes interrelated across layers and operators. The current section shifts the perspective to examine the structure in which AI for Network and Network for AI (the execution environment that enables decisions to be executed) compete for computing resources on the same physical infrastructure.

To illustrate this problem in accessible terms, consider the example of an edge node with limited computing resources. Suppose a telecommunications carrier deploys an MEC server near a base station to advance AI-based network control, and operates scheduling optimization via RIC xApps – a configuration already being deployed in current O-RAN environments. Even if this achieves the expected control quality in isolation, when user-facing AI inference services (such as video analytics and generative AI) are co-located on the same MEC server, GPU/CPU queues become congested, directly contributing to response delays in network control. Although sophisticated operational techniques such as static resource isolation and dynamic resource orchestration exist, in environments where AI inference load fluctuates dynamically, the optimal allocation itself changes at runtime, leaving a trade-off in which excessive margins reduce investment efficiency while insufficient margins result in contention. As a result, capital investment made to improve quality through AI control may fail to deliver the expected benefits due to resource contention on the same infrastructure.

Resource contention of this kind can, in principle, be addressed through an operator's internal capacity design, but because AI inference load changes dynamically due to model updates, request fluctuations, and offloading from other operators, design-time estimates alone cannot fully address all scenarios. Furthermore, the three structural indeterminacies act in a compound fashion on this problem. On the same MEC server, the extent of computing resource consumption by AI inference for network control and user-facing AI inference fluctuates at runtime (structural indeterminacy of location); it is difficult to determine whether quality degradation is attributable to the xApp itself, resource contention, or offloading from the cloud side (structural indeterminacy of basis); and it is not uniquely determinable whether responsibility for degradation should be attributed to the network control side, the service side, or the infrastructure management side (structural indeterminacy of agent).

Tables 3-3a and 3-3b organize the typical patterns in which decision outcomes become distorted due to execution environment constraints and resource contention, even when the AI's decisions themselves are not erroneous.

Table 3-3a: Impact of Structural Divergence Between Decision and Execution on Evaluation and Verification

What happens	Why this is problematic in AI-native environments	How it should be verified
Time lag between decision and execution	Delay fluctuates due to computational load of AI inference	Record decision time and execution time separately
Contention for resources	AI inference itself consumes GPU / CPU	Evaluate by isolating the impact of resource contention
Interference between controls	Learning states of multiple AIs differ	Observe interactions between controls

Decisions on halting under anomalous conditions	AI's decision space is broad and conditions cannot be exhaustively enumerated	Pre-design halt and intervention rules for unforeseen conditions
Oscillation in feedback	Response characteristics change with retraining	Include stability and convergence as evaluation metrics

Table 3-3b: Examples of Phenomena That Emerge When Communication, Computation, Power, and Operational Constraints Act Simultaneously

Constraint element	What happens	How it should be verified
Computing resources	Concentration of AI inference processing causes network control to lag, and load volume fluctuates with model updates	Evaluate the correctness of decisions and execution delay separately
Communication	The timing of AI decisions itself changes dynamically, and misalignment with control cycles causes control instability	Verify, including interaction with communication cycles
Power and facilities	During peak periods, power and cooling constraints prevent some processing from executing, and the power demand of AI inference fluctuates with learning state	Evaluate on the assumption of fluctuating resource conditions
Operational procedures	Speed differences between AI automated control and approval processes cause decision delays or suppression	Verify technical rationality and operational feasibility separately
Compound (overall)	Multiple constraints become strained simultaneously, and the dynamic changes of AI decisions make prediction difficult	Evaluate based not only on a single KPI but also on headroom and degradation trends

3.4 Summary — Why Conventional Verification Falls Short and What Must Be Questioned

This chapter organized the structural challenges in AI-native networks from three perspectives: decision distortions and their impact on investment decisions (Section 3.1), the difficulty of tracing the basis of decision-making (Section 3.2), and computing resource contention (Section 3.3). What these challenges share is that the premises of conventional verification – isolation, reproducibility, and traceability – all cease to hold.

Conventional network verification has relied on three premises: the evaluation target can be isolated; conditions can be fixed and results reproduced; and the causality of decisions can be traced. As discussed in Section 2.2, AI for Network and Network for AI are cyclically interlinked, and with AI embedded as a decision-making agent within this cycle, the decision logic itself changes through learning with each iteration. Under this structure, all three premises become difficult to maintain. Because the generation, execution, and environment of decisions collectively form behavior, the evaluation target cannot be isolated as a standalone entity (in the example from Section 3.1, quality degradation cannot be clearly attributed to edge AI or cloud AI). Because internal states and learning histories differ with each execution, identical conditions cannot be reproduced. Because the basis of decision-making is distributed across multiple layers and entities, causal relationships cannot be uniquely traced (Section 3.2).

The same structural problem arises in the domain of security. In environments where AI dynamically generates and updates control decisions, distinguishing between "normal adaptive behavior" and "anomalous behavior caused by

an attack" becomes difficult, breaking down the premises of conventional detection methods based on fixed patterns.[10] This problem connects to all three structural indeterminacies: the distribution of decision-making locations expands the attack surface (location); the difficulty of tracing the basis of decision-making delays intrusion detection (basis); and the multi-layered nature of decision-making agents renders the locus of defense responsibility ambiguous (agent).

Security implications extend beyond conventional threat models. Because decision positions are distributed and dynamic, adversaries can target the weakest node in a shifting topology — a challenge that static perimeter-based defenses cannot address. The distributed nature of decision basis creates opportunities for data poisoning and adversarial input attacks that are difficult to detect when no single entity holds the complete decision context. The composite nature of decision subjects means that supply chain attacks — compromising one AI component among many — can propagate effects through the entire decision cycle without triggering conventional anomaly detection. These threats demand verification approaches that treat security not as a bolt-on assessment but as an integral dimension of the Verification Collaboration Platform, enabling red-team exercises across operator boundaries, shared threat intelligence regarding AI-specific attack vectors, and continuous security validation that evolves alongside AI model updates.

Data governance and privacy considerations further compound these challenges. AI-native network control requires continuous collection of network state information, traffic patterns, and potentially user behavioral data as inputs to decision-making. In jurisdictions with comprehensive data protection frameworks — such as the EU's General Data Protection Regulation (GDPR) or emerging privacy legislation in other regions — the question of how such data is collected, processed, and retained within distributed AI decision loops introduces additional regulatory constraints. The Verification Collaboration Platform should therefore incorporate data governance as a cross-cutting requirement, ensuring that verification scenarios account for privacy-preserving approaches such as federated learning, differential privacy, and data minimization principles.

The collapse of these verification premises leads to three practical consequences. First, the validity of capital investment (CAPEX) can no longer be evaluated in advance. The effectiveness of AI control depends on interactions across layers, yet with verification premises collapsed, means to compare not only "AI control introduction versus non-introduction scenarios" but also configuration patterns regarding "in which layer to deploy AI control and how" under identical conditions are lost, making it impossible to construct a basis for investment decisions. Second, the design of evaluation and verification itself becomes difficult. Without the ability to determine in advance what constitutes the evaluation target and what scope the verification should cover, there is no basis for test item comprehensiveness. Third, the boundary of responsibility becomes difficult to establish. Identifying "which entity's decision caused the issue" during a failure requires the ability to trace the causality of decisions, yet the collapse of the three premises is precisely what makes such tracing infeasible. These points are revisited in Chapter 4 as design requirements for the Verification Collaboration Platform. Given this methodological limitation, Chapter 4

[10] Research on security, reliability, and self-healing in AI-involved network control includes work on the organization of Security/Privacy/Trust in 6G Open RAN environments (Porambage et al., 2025), intrusion detection and cooperative defense using machine learning (Li et al., 2021; Ataa et al., 2024), and self-healing design in cyber-physical systems (Johnphill et al., 2023). All of these indicate the necessity of considering AI control not as an individual function but as a structure encompassing the entire network.

addresses how to reconstitute evaluation and verification as a framework and under what division of roles to advance AI-native networks.

Column 4: Where Should Investment Be Directed to Yield Results? — Significance of the Verification Collaboration Platform as Illustrated by Router Path Selection

In AI-native networks, determining "where to invest to yield results" requires a verification framework different from conventional approaches. This can be illustrated through the familiar example of router path selection.

In conventional routers, the rules for path selection are defined at design time. Paths are selected based on explicit criteria such as cost and AS-path length, following protocols including OSPF (Open Shortest Path First) and BGP (Border Gateway Protocol). Evaluation and verification of a router only need to confirm whether the router delivers the specified path selection and forwarding performance under given traffic conditions. The evaluation target is self-contained within the individual router, and the basis of decision-making is accessible through configuration files and routing tables.

In an AI-native environment, the situation differs even for the same router path selection. For example, suppose a router reroutes traffic from path A to path B with the objective of minimizing overall service latency. This rerouting may not have been the router's own decision but the result of edge-side AI predicting a traffic surge, cloud-side resource management AI detecting computational load imbalance, and network control AI inferring that "under current conditions, path B would better suppress latency" and issuing a path change instruction. Moreover, the inference itself may have reflected the fact that the processing capacity of a specific data center was temporarily reduced due to power constraints.

Even if the router is isolated as an evaluation and verification target, the appropriateness of the path change behavior cannot be assessed. The router's forwarding performance is consistent with specifications, and there is no problem with its operation. However, in an AI-native environment, a single behavior such as path selection is established as a result of interactions among decisions and environmental conditions across multiple layers, including network control AI, cloud-side resource management, and power constraints. Consequently, it becomes difficult to disentangle the interactions and explain "why this path was selected" and "under what preconditions that selection holds."

The difficulty of explaining this extends beyond post-hoc analysis following failure incidents. What is more important is that it directly constrains routine decision-making related to design and investment. If the interactions cannot be disentangled to clearly identify the basis of this path selection — a network control decision, an effort to avoid shortages of computing resources, or a reflection of power constraints — then it is impossible to rationally determine "where to focus next." The appropriate target for improvement — whether it is to augment the network, add computing resources, or revise the power infrastructure — can only be determined once it becomes possible to compare which element governs which behavior in the current configuration.

Furthermore, even when achieving the same service quality, cost structure and risk profile vary significantly among configurations for absorbing variability: through network control; through computing resource redundancy; or through the accuracy of AI decision-making. Unless these configurations can be compared in advance, investment decisions inevitably rely on the bottom-up aggregation of individual equipment.

Enabling such comparison requires a verification environment that can reconstruct the chain of decisions and compare behaviors under varying configurations and conditions. However, given that the chain of decisions spans the network control of telecommunications carriers, the resource management of cloud operators, and the supply constraints of power utilities, such an environment cannot be built by a single entity.

The Verification Collaboration Platform is needed not only for the post-hoc response when failures occur. It is needed to enable routine decisions, such as "where to invest to yield results" and "which configuration is superior overall," to

be made through comparison by the relevant stakeholders under common conditions, rather than through the confined internal deliberations of individual operators.

Chapter 4: How to Verify and Who Bears Responsibility — The Concept of the Verification Collaboration Platform

— Designing a Forum Where Industry, Academia, and Government Bring Their Resources Together to Compare, Compete, and Verify —

The analysis through Chapter 3 points to a single conclusion.
Verification of networks in which AI is involved cannot be completed through individual verification efforts alone.
Then, who should bear what responsibilities, and in what forum should comparison and verification be conducted?
This chapter presents the concept of the Verification Collaboration Platform – a framework in which industry, academia, and government bring their resources together and verify such networks under common conditions through competition.

This chapter presents the concept of the Verification Collaboration Platform designed to address the three practical consequences identified in Section 3.4 – the validity of capital investment (CAPEX), the design of evaluation and verification, and the boundary of responsibility. It first organizes the significance of expanding the perspective from technical verification to a collaborative verification process (Section 4.1); it then clarifies the roles and scope of responsibility for each entity involved (Section 4.2); and it further presents the functional structure and design philosophy of the Verification Collaboration Platform that supports these elements (Section 4.3).

4.1 Why is Verifying "Correct Operation" Alone Insufficient?

Conventional network verification asks whether something "operates correctly from a technical standpoint." However, as discussed in Section 3.4, in AI-native networks, the three verification premises (isolation, reproducibility, and traceability) break down, giving rise to multiple questions that cannot be answered by that approach alone. To begin with, questions on investment effectiveness, the boundary of responsibility, and configuration comparison have been asked in the past as grounds for decision-making, but in an environment where the three verification premises have collapsed, the very means to answer these questions have been lost.

Furthermore, as AI becomes involved in decision-making, additional questions are encountered routinely in the field of technology selection and operational design – questions that lack clear answers and have been left unanswered. These include: "How should the execution space of AI be designed and who should grant it what degree of freedom?"; "Who should switch control, and at what stage, under unforeseen conditions?[11]"; "To what extent must the basis of decision-making be traced and recorded to fulfill accountability?"; "What should be established as common conditions to compare different implementations fairly?"; "How should verification be repeated each

[11] The switching mechanism referred to in the main text corresponds to the concept generally known as the "AI-Kill-Switch." This is not intended to completely halt AI control; rather, it is organized as a means of temporarily switching the decision-making and control agent for the purpose of ensuring operational safety and verifiability. It should be noted that while technical fallback mechanisms (3GPP), operational-level definitions (TM Forum), and legal human oversight obligations (EU AI Act) exist, no unified standard has been established regarding "when and by whom the switch is to be made."

time learning is updated?" ; and "How should normal adaptive behavior of AI be distinguished from anomalous behavior caused by an attack?" While some of these issues have been discussed within their respective organizations, they are by nature impossible to resolve solely within a specific industry or standardization framework, and remain as implicit assumptions.

There are three reasons for this. First, because AI decisions continue to change through learning updates, specifications agreed upon at a given point in time are not necessarily valid after the next update. Verification requires a mechanism that enables continuous, rather than a one-time, conformance confirmation. Second, premises such as halt criteria, the level of accountability, and comparison conditions cannot function without practical agreement existing among the relevant entities, rather than as definitions in specification documents. Third, because the chain of decision-making crosses operator boundaries, confirming overall behavior requires the actual environments of multiple entities to be brought together. This necessitates not only agreement on specifications, but also a forum for comparative verification in real environments. Table 4-1 presents examples of the differences that the presence or absence of the Verification Collaboration Platform may bring to various aspects of decision-making.

On this basis, this white paper presents the following as the direction that such consensus-building processes should aim for: to establish a framework in which multiple entities bearing different roles and positions can bring together their technologies and environments under common conditions and continue comparison and verification toward the realization of AI-native networks.

The Verification Collaboration Platform shares certain characteristics with regulatory sandboxes (as deployed in EU fintech and AI regulation) and open testing environments (such as TIP and ONAP in telecommunications), but differs fundamentally in scope. It is designed not for single-technology conformance testing but for cross-infrastructure behavioral verification involving multiple autonomous decision-making entities. Unlike vendor-specific testbeds, the platform operates under conditions agreed upon by participating stakeholders, ensuring neutrality in both condition-setting and result attribution.

Collaborative verification does not require the disclosure of competitive domains. The object of comparison is not the internal logic of each entity's AI, but the results obtained under common conditions. By establishing common conditions and bringing together behavioral tendencies and acceptable ranges under those conditions, verification can be conducted without disclosing internal logic. Conversely, without collaborative verification, introduction decisions must be made with comparison conditions remaining unaligned, while the locus of responsibility in the event of a failure is left to post-hoc negotiation. The costs of such delays in decision-making and rework are already being incurred, though they may remain invisible. The Verification Collaboration Platform is not something that imposes a new burden; rather, it is an attempt to structure costs that are already being borne. The starting point of this framework is precisely the gathering of relevant entities to discuss how the conditions and scope of collaborative verification should be established.

Table 4-1: Practical Functions That the Verification Collaboration Platform May Bring to Decision-making (Examples)

Decision-making context	Current state	With the Verification Collaboration Platform

Investment decisions	The layer where effects materialize is difficult to discern, and evaluation remains limited to individual equipment units	Different configurations can be compared under identical conditions, providing a basis for determining how much to invest in which configuration
Design choices	Comparison conditions are not aligned, making it difficult to evaluate alternatives	Essential differences can be isolated, enabling design decisions that avoid dependence on a specific vendor
SLA and the boundary of responsibility	Identifying the cause of quality fluctuations is time-consuming, and responses are predominantly reactive	Conditions for fraying can be identified in advance, enabling SLA terms and the boundary of responsibility to be set with supporting evidence
Sharing of failures	The preconditions of failures are not easily shared, and experience remains confined to individual companies	Preconditions that did not hold can be shared, preventing recurrence of the same failures across the industry as a whole
Verification of new deployments	Criteria for what to observe and compare are difficult to establish	Verification items can be designed in advance, enabling risks and benefits of deployment to be evaluated prior to introduction

4.2 Who Brings What to the Table? — Conditions and Expertise by Entity Type

In AI-native networks, AI-based control, rule-based control, and human decision-making coexist, and the decision-making agent shifts dynamically depending on the situation. For this reason, rather than assigning responsibility to a specific entity in a fixed manner, it becomes important to ensure that "what happens under which conditions" can be verified in advance. It should be noted that the following categorization of entities is based on functionality, and a single organization may assume multiple roles. Table 4-2 provides examples of the conditions and expertise that each entity may bring to this verification. While Table 4-2 is organized primarily around technical entities, the ultimate recipients of verification results are the users of telecommunications services. The comparative data and evaluation metrics generated by the Verification Collaboration Platform contribute to ensuring transparency for users, serving both as reference material for selecting services and as evidence supporting the accountability of operators.

Table 4-2: Conditions and Expertise Brought by Each Entity to the Verification Collaboration Platform (Examples)

Entity	Conditions and expertise contributed
Technology and development	
Technology entity (AI algorithm and model development)	• Presents behavioral conditions and tolerance ranges to visualize the impact of its AI control on other system components • Provides verification data on update impacts to grasp performance changes after learning and updates in advance
AI platform operator	• Presents configuration conditions of the inference platform and model update history to enable tracing of the basis of decision-making in AI control • Provides API specifications and compatibility conditions to grasp the impact of AI platform switching and migration in advance • Presents load conditions such as traffic characteristics, latency tolerance, and bandwidth demand to reflect the communication requirements of AI workloads in network design
Terminal and device manufacturer	• Presents terminal-side constraint conditions (processing capacity, latency characteristics) to prevent inconsistencies between edge decisions and network decisions

	• Provides terminal behavior verification data to enable isolation of terminal-originated failures
Operations and provision	
Telecommunications operator	• Presents verification conditions for configuration and control policies to ensure interoperability in multi-vendor environments • Accumulates comparative evaluation data on SLA compliance to provide a basis for AI deployment investment decisions • Accumulates decision logs and response histories for configuration changes and failures to expedite root cause identification in the event of incidents • Verifies the effectiveness of fallback procedures to minimize service impact in the event of control switching
Cloud / data center operator	• Presents execution environment constraint conditions to enable the isolation of failures that originate from its own systems from those that do not • Provides advance assessment based on behavioral data to grasp impacts under compound load conditions • Verifies the validity of detection conditions and threshold settings to reduce false positives in anomaly detection
Power utility	• Presents supply fluctuation conditions and constraints to proactively grasp the impact of power variability on AI control • Provides verification data on fluctuation impacts to enable isolation of power-originated failures
Usage and evaluation	
Users and enterprise customers	• Specifies usage conditions and required performance levels to grasp the impact of AI control on service quality in advance • Verifies conformity with business requirements to reduce business impact from unexpected behaviors
Institutions and safety	
Security specialist organization	• Presents threat scenarios and evaluation criteria to enhance response capabilities against AI-specific threats • Verifies the effectiveness of identification methods to prevent misidentification between normal adaptation and attack responses
Public institutions, etc.	• Presents the framework of institutional requirements and verification criteria to enable verification results to be reflected in institutional design • Designs feedback pathways from verification results to institutional frameworks so that they can keep pace with technological progress

4.3 What Must the Verification Collaboration Platform Be Capable Of?

The preceding section organized the conditions and expertise that each entity can contribute. When these are grouped by the nature of their functions, four domains emerge, as discussed below.

Composition of the Four Functional Domains

The four functional domains that constitute the Verification Collaboration Platform are as follows: the AI Control Decision Domain; the Observation, Evaluation, and Verification Domain; the Execution and Resource Infrastructure Domain; and the Interoperability and Competition Design Domain. These four domains operate in mutual coordination. The conditions and expertise that each entity brings, as organized in the preceding section, are

classified into these domains and shape their respective functions. The domains are not intended to function independently but are premised on mutual coordination.[12]

1. **AI Control Decision Domain:** This domain is responsible for the design, learning, and inference of AI models that perform network control. It requires an assurance of decision transparency, reproducibility, and explainability, and outputs decision results in a form that enables coordination with other domains. The impact of network quality (latency, bandwidth) on AI inference results is also included within the scope of verification for this domain.
2. **Observation, Evaluation, and Verification Domain:** This domain verifies the behavior of AI control models in virtual environments or partial physical equipment, and visualizes and observes their performance. It benchmarks KPIs such as throughput, latency, and stability, and evaluates the validity of decision outcomes. The impact of changes in network and infrastructure state on AI behavior is also included within the scope of measurement for this domain.
3. **Execution and Resource Infrastructure Domain:** This domain encompasses the infrastructure on which decision results are actually executed, including telecommunications networks, clouds, data centers, and power and cooling facilities. In AI-native environments, AI itself is an entity that consumes computing resources, and resource contention can occur, making this domain in and of itself an important verification target. Enabling comparative verification under common conditions by interconnecting the verification environments already held by each entity is also an important role of this domain.
4. **Interoperability and Competition Design Domain:** This domain plays the role of connecting the aforementioned domains and defining minimum common specifications and rules while enabling the coexistence of competition and co-creation. Rules regarding resource allocation between AI infrastructure and network infrastructure are also addressed within this domain. Standardization bodies, public institutions, and industry-academia-government collaborative consortia are envisioned as its stewards, and it is expected to function as a forum for pre-standardization implementation verification and competition. While verification environments provided by specific platform operators are effective within their respective ecosystems, they depend on the provider for condition-setting and attribution of results, and are therefore unsuitable for neutral comparative verification across multiple entities. This domain is fundamentally different from existing specific platforms in that setting verification conditions and management of results are conducted based on agreement from participating entities.

[12] In international discussions centered on ETSI and ITU, the importance of evaluation and verification methodologies that treat everything, from the physical layer to control and operations, as an integrated whole has been noted. Verification approaches that link measurement, simulation, and emulation are attracting attention as practical methods for ensuring the reliability of networks in which AI is involved in decision-making (ETSI, ITU). In addition, telecommunications operators and platform vendors are advancing the construction of digital twins that reflect actual operational environments. For example, NTT DOCOMO is conducting behavioral verification at the RAN scale using a digital twin targeting 5G networks, and NVIDIA Omniverse provides a digital twin environment that integrates AI control and communications in the manufacturing and robotics domains. In wireless networks, because physical and spatial factors such as building placement, terrain, and mobile objects significantly affect communication quality, high-fidelity digital twins at the urban scale are considered important for evaluating the appropriateness of AI control (NTT DOCOMO; NVIDIA).

Cyclical Structure of the Four Domains (Figure 4-1)

For the cyclical structure to function, designing a "condition-defined execution space" that determines how much freedom to grant AI is a prerequisite. This design also involves perspectives related to Ethical, Legal, and Social Issues (ELSI) – including responsibility, accountability, and governance. The functional requirements borne by the above four domains, when organized by domain, are as follows (Table 4-3).

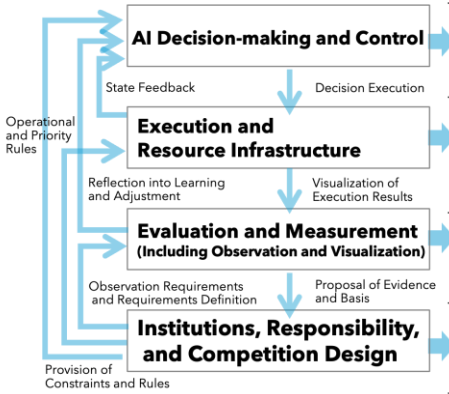

	Primary Question	Primary Time Scale	Primary Location	Primary Granularity
AI Decision-making and Control	**What does AI base its decisions on, and to what extent can we autonomously entrust them?**	Milliseconds to seconds (inference and control cycles)	Edge/within the network/cloud	Flow, session, user
Execution and Resource Infrastructure	**Are decisions feasible under current constraints of computation, communication, and power?**	Milliseconds to minutes (traffic fluctuations and resource contention)	RAN, core network, data centers, power facilities	Physical units such as CPU, bandwidth, and power
Evaluation and Measurement (Including Observation and Visualization)	**To what extent can decisions and behaviors be traced and explained? What is considered success or failure? Are comparisons possible?**	Seconds to hours (log aggregation and analysis)	Distributed logs, monitoring platforms	KPIs, statistics, events
Institutions, Responsibility, and Competition Design	**Who assumes responsibility, and to what extent?**	Days to years (rule updates and accountability decisions)	Organization, contract, and legal/regulatory layers	Policies, rules, and scopes of responsibility

Figure 4-1: Cyclical Structure of Decision-Making, Execution, Evaluation, and Institutional Design in AI-Native Networks

Table 4-3: Functional Domains and Functional Requirements of the Verification Collaboration Platform

Functional domain	Functional requirements
AI control decision	• The ability to trace the basis and process of decision-making and record them in an explainable form • The ability to detect behavioral changes following learning and updates and to re-verify
Observation, evaluation, and verification	• The ability to compare and evaluate decisions of different AI models under common conditions • The ability to measure and compare KPIs of multiple entities against a common standard
Execution and resource infrastructure	• The ability to execute verification under conditions approximating real environments (including communication, computing, and power constraints)
Interoperability and competition design	• The ability to share and compare verification results without disclosing the internals of competitive domains • The ability to agree on verification conditions and rules among relevant entities and to update them on an ongoing basis

4.4 Summary — Structural Clarification Enables Verification and Motivates Participation

This chapter organized the challenges of verification in AI-native networks in three stages. It first demonstrated the reasons for the breakdown of the premises of conventional verification and the necessity of collaborative verification in response (Section 4.1), and then it organized the conditions and expertise that each entity brings (Section 4.2) and presented the four functional domains, the cyclical structure, and the functional requirements that bind them together (Section 4.3). Through this structural clarification, it becomes possible to explore a path toward conducting verification, which has conventionally been confined within individual entities, under common conditions across multiple entities.

Participation in the Verification Collaboration Platform holds practical significance for each entity. The ability to share verification costs, the availability of a common platform for tracing the basis of decision-making and delineating responsibility in the event of failures, and the ability to bring proposals backed by verification track records to international standardization forums – each of these represents a practical advantage that is difficult to obtain independently. Furthermore, if an economic cycle emerges through the provision and utilization of data and APIs required for verification, participation in the platform itself serves as a business opportunity, leading to the formation of a self-sustaining ecosystem.

Verification of AI-native networks is no longer of a scale that any single entity can cover alone. The Verification Collaboration Platform is a forum in which each entity brings its own strengths, and through comparison and verification under common conditions, collectively builds the reliability of technology and the basis for business decisions.

Chapter 5: Where to Begin — A Phased Approach to Demonstration

— Beginning not with a finished form, but with iterative cycles of verification and dialogue —

Chapter 4 organized the structure and functional requirements of the Verification Collaboration Platform. However, this platform is not something to be constructed in its complete form all at once. It must be advanced incrementally as a process in which the underlying assumptions themselves are updated through repeated cycles of verification and dialogue.

Global approaches to the convergence of AI and telecommunications vary widely, including regulation-led, market-driven, and state-led models. Japan's approach is based on a soft-law foundation that enables such convergence to advance incrementally through iterative coordination between institutional design and implementation. This approach can prove advantageous in initiatives premised on consensus-building among multiple entities, such as the Verification Collaboration Platform.

Specifically, a transition through the following three phases is envisioned.

Phase 1 (Short-term): Partially introducing AI control in limited use cases and ensuring the observability of control decisions. Comparative data between AI and non-AI decisions are accumulated through parallel operation with existing rule-based control. Concurrently, stakeholders initiate discussions on the requirements and architecture of the Verification Collaboration Platform, building consensus on common conditions and evaluation metrics. Phase 1 success criteria might include, for example, demonstrating measurable differences in control quality between AI-driven and rule-based approaches across at least three distinct use cases, with sufficient statistical confidence to inform initial investment decisions.

Phase 2 (Medium-term): Comparative verification among multiple entities and across diverse implementations using the Verification Collaboration Platform. The conditions under which investment effects can be realized, patterns of the boundary of responsibility, and the validity of evaluation metrics are concretized on the basis of empirical demonstration. Feedback to institutional design also commences at this phase.

Phase 3 (Medium- to Long-term): The verified scope is gradually expanded, and the applicable areas and degree of autonomy of AI control are elevated. The confirmation of safety and effectiveness at each phase through the Verification Collaboration Platform constitutes the rational basis for proceeding to the next phase. Furthermore, at this phase, the aim is to foster an ecosystem in which an economic cycle matures through the provision and utilization of data and APIs necessary for verification such that participation in the platform itself serves as a business opportunity.

This white paper was prepared with the objective of organizing structural issues accompanying the transition to AI-native networks and articulating them in a form that can be shared among stakeholders. It is hoped that this organization of issues will serve as a common language leading to ensuing concrete discussions, demonstrations, and collaborations.

Appendix A: Definitions of Terms and Concepts

This appendix organizes the definitions of key terms as they are used in this white paper.

Term	Definition
AI	Refers to software mechanisms that employ machine learning and inference technologies to participate in the control, operation, and configuration of telecommunications networks and computing resources or in decisions and selections pertaining thereto. It is not necessarily limited to entities that autonomously continue learning; it encompasses mechanisms that are combined with pre-designed models and rules to make decisions under specific objectives and constraints.
AI (Decision-making Agent)	Refers to a logical entity that generates decisions based on policies, optimization objectives, and constraints pertaining to the control of telecommunications networks. It does not denote a single physical entity or algorithm, but is positioned as a conceptual entity that may be implemented in a distributed manner across multiple AI mechanisms. It corresponds to "AI (Decision-making Agent)" at the center of Figure 2-2.
AI as a Telecommunications User	Refers to AI positioned on the side that utilizes resources such as telecommunications, computation, latency, and reliability. It operates as a part of applications or services and presents its processing requirements and performance demands as requests to the telecommunications infrastructure. It corresponds to the bottom part of Figure 2-2.
Control and Execution AI	Refers to AI mechanisms responsible for reflecting the decisions generated by AI as a decision-making agent into the system as concrete controls. These mechanisms operate in proximity to network equipment and infrastructure and are frequently implemented as individual functions. They correspond to the left part of Figure 2-2.
AI Decision Items	Refers to the specific objects of decision-making that AI undertakes in the control and operation of networks and computing resources. These include route selection, bandwidth allocation, radio resource assignment, and determination of whether configuration changes are necessary, as well as selections pertaining to computing resources such as processing placement, inference frequency, and execution timing.
AI Involvement in Decision-making	Refers to a state in which AI assumes part of the decision-making process in the design, operation, and control of telecommunications networks, and the outcomes thereof affect the behavior of networks and societal infrastructure. This white paper encompasses not only cases in which AI directly performs control, but also forms in which: (1) humans or other systems execute actions based on AI decisions; (2) the network side is controlled in accordance with AI requirements; (3) AI decisions and network states are mutually dependent; and (4) decision outcomes propagate to social and public infrastructure. These are collectively regarded as the common structure of the "migration of decision-making agents." The state in which the degree of AI involvement in decision-making deepens such that the dynamic generation and updating of control decisions at runtime becomes a premise of network operation is what this white paper terms an "AI-native network."
AI for Network	Refers to an approach for controlling, operating, and optimizing network infrastructure, including telecommunications, through the use of AI.
Network for AI	Refers to the constraints, responsibilities, and design premises that the entirety of infrastructure, including telecommunications, computing resources, data distribution, execution platforms, and energy constraints, must assume in order to establish a societal infrastructure in which AI participates as a decision-making agent. It does not refer to network optimization solely aimed at maximizing AI performance.
AI-native Network	Refers to a structure, or the design and form of implementation thereof, in which AI is incorporated not as a retrofit optimization measure but as a premise of network operation. It encompasses both the aspect of AI controlling the network (AI for Network) and of infrastructure enabling AI inference

	and learning (Network for AI). Control, decision-making, and optimization are performed dynamically under the collaboration of humans and AI, and depending on the maturity of operations and the applicable domain, AI involvement may transition incrementally from an auxiliary role to a leading role. Whereas in conventional networks the logic of control could be determined at the time of design, in AI-native networks, control decisions are dynamically generated and updated at runtime, making it impossible to prescribe "what is operating where and by whose decision." This structural difference compels a reexamination of conventional frameworks such as the assessment of capital investment validity, system verification, and the boundary of responsibility in the event of failures.
Control Decision (in this White Paper)	Refers to a structural process that emerges dynamically from rules established by humans, inferences generated by AI at runtime, the physical constraints of infrastructure, and the cyclical interactions among them. The concept referred to as "decision" in this white paper differs from static control instructions by conventional orchestrators or individual inferences by standalone AI agents. Consequently, it is structurally difficult to attribute the "correctness" of a decision to any single entity or algorithm.
Decision Structure	Refers to a foundational concept for comprehending the nature of decisions that are established across multiple layers, including AI, telecommunications, operations, and institutional frameworks. "Decision structure" in this white paper refers to the totality of the relationships that define where decisions are generated, which entities are involved, at what granularity and on what time axis they are updated, and how these are linked to execution, evaluation, and the boundary of responsibility to form overall system behavior in AI-native networks.
Three Structural Indeterminacies	Refers to a structural situation in AI-native networks in which the location of control decisions (where decisions are generated), the basis (on what grounds decisions are made), and the agent (who assumes the decision-making role) cannot be determined at design time. This white paper positions these three structural indeterminacies as the foundation of the discussion in Chapter 2.
Sou / Layer	Refers to the distinction whereby "layer" is used to denote technical-functional divisions or hierarchical structures within an architecture, while "sou" is used to denote conceptual and analytical perspectives such as institutional frameworks, responsibility, and decisions.
Autonomous Network	Refers to a network that minimizes manual configuration and operation and autonomously performs state recognition, decision-making, and execution through AI and automation technologies. It is considered to evolve incrementally according to the level of autonomy.
API (Application Programming Interface)	Refers to an interface for connecting different systems and functions. In this white paper, emphasis is placed not on fixing individual specifications but on a design that enables interconnection through a bring-your-own approach.
Boundary of Responsibility	Refers to the boundary defining which entity bears responsibility and to what extent when failures or erroneous decisions occur. In AI-native networks, because decisions are dynamically generated across multiple entities and layers, it becomes difficult to determine this boundary in advance through contracts or design-time agreements, as was conventionally done.
Cyclical Structure	Refers to a structure in which the four functional domains of the Verification Collaboration Platform (AI control decisions; observation, evaluation, and verification; execution and resource infrastructure; and interoperability and competition design) are interconnected in a flow of control decision generation, execution, observation, evaluation, and the next decision, and continue to cycle on an ongoing basis. This corresponds to Figure 4-1.
Verification Collaboration Platform	Refers to a platform through which multiple entities bring together their respective conditions, expertise, and verification environments to compare and verify the behavior of AI-native networks under common conditions. It is not merely a platform for technical verification, but possesses the character of an institutional platform encompassing the formation of bases for investment decisions, consensus on the boundary of responsibility, and feedback to institutional design. It is premised on neutrality that does not depend on any specific platform provider.

Acknowledgments

In the course of preparing this white paper, NICT Open Summit 2025 was held over two days on October 29 and 30, 2025. This event brought together diverse participants, including researchers within NICT as well as stakeholders from industry engaged in cutting-edge initiatives in the fields of information and communications technology, AI, and data utilization, and vigorous discussions were held from the perspectives of implementation and societal deployment.

This white paper was prepared against the backdrop of the many insights and issues raised through lectures, panel discussions, and exchanges of views among participants at the event.

Although this white paper does not specifically reference or cite individual presentations or viewpoints, the discussions exchanged throughout the event provided significant insight toward its structure and framing of issues.

The authors express their deepest gratitude to all those who provided valuable expertise through this event.

Authors

Role	Organization/Division
General Supervision	AI Research and Development Promotion Unit
Planning	AI Research and Development Promotion Unit
Planning Cooperation	Beyond 5G Unit
Planning Cooperation	Strategy Planning Office, Strategic Planning Department
Editing	AI Research and Development Promotion Unit
Compilation	AI Research and Development Promotion Unit
Notation and Format Review (Proofreading)	Press Office, Public Relations Department

List of Internal Technical Review Contributors (Listed by Family Name in Japanese Alphabetical Order)

Role	Name	Affiliated Research Institute, etc.
Technical Advice	Hisashi Ibaraki	Advisor
Technical Advice	Hideyuki Tokuda	Executive Advisor
Technical Advice	Kentaro Ishizu	Beyond 5G Research and Development Promotion Unit
Technical Input	Go Kato	Advanced ICT Research Institute, Koganei Frontier Research Center
Technical Advice	Hikaru Kawasaki	Network Research Institute, Wireless Networks Research Center
Technical Advice	Toshiyuki Kamiya	Universal Communication Research Institute
Technical Input	Yutaka Kidawara	Executive Officer
Technical Advice	Hideyuki Tokuda	Executive Advisor
Technical Advice	Hiroaki Harai	Network Research Institute
Technical Advice	Tao Ban	Cybersecurity Research Institute, AI Security Research Center
Technical Advice	Takahiro Hirayama	Network Research Institute
Technical Advice	Iwao Hosako	Beyond 5G Research and Development Promotion Unit
Technical Advice	Motoaki Yasui	Vice President

* The technical review was conducted during the preparation of this white paper for the purpose of obtaining technical opinions and advice. Final responsibility for the content rests with the AI Research and Development Promotion Unit.

Colophon

Title:
AI-Native Networks White Paper — AI × Communications: Engineering the Future—
Edition:
Version 1.0
ISBN:
978-4-904020-46-3
Publication Date:
18 May 2026

Publisher and Contact:
National Institute of Information and Communications Technology: NICT
AI Research and Development Promotion Unit
4-2-1 Nukui-Kitamachi, Koganei-shi, Tokyo 184-8795, Japan
E-mail: AI-prom-pub@ml.nict.go.jp
https://www2.nict.go.jp/aipromo/

Errata and Supplementary Information:
https://www2.nict.go.jp/aipromo/whitepaper/appendix.html
The content of this document is based on information available at the time of publication. Corrections of errors identified after the publication of this white paper and supplementary information reflecting updates to policies and guidelines are provided as needed at the above URL.

Feedback and Issue Proposals:
This white paper has been published as a foundation for future discussions.
Feedback on the content and proposals for issues to be discussed are accepted through the methods described on the above website.
* Please note that individual responses or public disclosure of submissions are not guaranteed.

ISBN:978-4-904020-48-7

National Institute of
Information and Communications Technology

National Institute of Information and Communications Technology
AI Research and Development Promotion Unit
E-mail:AI-promo-pub@ml.nict.go.jp

Published:Version1.0/May 2026