Algora Volume: 2 Issue: 2, 2025 pp. 124-190

Journal of AI, Data Science and Engineering, Peer Reviewed, Open Access Journal

Interoperability of AI Agents in Multi-Vendor Telecommunications Networks: An Open Framework for Autonomous Network Coordination

[1]Mohammad Mustafa
[2]Mohammad Firoz

[1] *Jawaharlal Nehru Technological University; Contact Center Company (ccc by stc), Riyadh, Saudi Arabia.*
ORCID: 0009-0002-7535-3321

[2] *Jawaharlal Nehru Technological University; Contact Center Company (ccc by stc), Riyadh, Saudi Arabia.*
ORCID: 0009-0006-4529-9263

Abstract

Autonomous network management increasingly requires AI agents from different vendors and operational domains to exchange intents, reconcile competing actions, and respect network policies. This article proposes the Open Framework for Autonomous Network Coordination (OFANC), a five-layer architecture combining agent discovery, authenticated communication, a common semantic interoperability profile, bounded negotiation, policy checks, and domain-specific execution. The paper defines a constrained coordination problem and an implementation-oriented protocol, including capability validation, conflict handling, and escalation. Its quantitative component is a reproducible stochastic illustration, not a deployment or packet-level network experiment. A seeded Python generator produces 168,000 task records under seven assumed scenarios, four parameterized architecture profiles, and five agent-population settings. In this illustrative model, the OFANC profile yields 83.58% interoperability and 79.53% intent fulfillment, compared with 75.89% and 63.85% for the generic messaging profile. These differences follow from assumed baseline probabilities and cannot establish empirical superiority. The OFANC profile also incurs higher signaling overhead (0.579 KB per task) and longer modeled recovery time (51.04 s), and its unauthorized-request blocking rate (88.36%) falls short of production security expectations. The contribution is therefore an explicit architecture, message specification, reproducible model, and validation protocol, rather than demonstrated multi-vendor network performance. Independent implementations, instrumented network trials, and security conformance tests are required to establish operational effectiveness.

Keywords: AI Agent Interoperability; Multi-Vendor Telecommunications Networks; Autonomous Network Coordination; Multi-Agent Systems; Intent-Based Networking; 5G and 6G Networks; Zero-Touch Network Management

Introduction

Evolution of Intelligent Telecommunications Networks

Telecommunications networks have undergone substantial changes with the introduction of fifth-generation (5G) technology, network virtualization, software-defined networking, and intelligent resource management. These developments have increased the flexibility of network operations while introducing new challenges in managing distributed infrastructure. Modern telecommunications networks increasingly consist of interconnected radio access, core, transport, and cloud-based systems that must operate together to deliver reliable services. The growing diversity of applications, particularly those requiring low latency, high bandwidth, and consistent connectivity, has also placed greater demands on network management systems (Giordani et al., 2020; Barakabitze et al., 2020).

Traditional network management relies heavily on predefined rules, centralized control, and manual intervention. Although these approaches remain useful, they are less effective when network conditions change rapidly or when service delivery depends on several independently managed domains. Consequently, artificial intelligence (AI) and machine learning have gained attention as tools for improving network monitoring, fault prediction, resource allocation, and operational decision-making. Coronado et al. (2022) identified automated closed-loop management as an important requirement for future telecommunications networks, particularly as operators move toward increasingly autonomous operations.

However, increased automation does not necessarily guarantee cooperation between intelligent systems. As network functions become more distributed and are supplied by different vendors, the ability of autonomous management components to communicate and coordinate their decisions becomes increasingly important.

Emergence of AI Agents in Network Management

AI agents represent a further development in intelligent network management. Unlike conventional automation systems that execute predefined instructions, autonomous agents can interpret operational objectives, assess available information, select actions, and respond to changes within their assigned responsibilities. Such agents may support fault diagnosis, traffic optimization, network slicing, energy management, and service assurance.

Bonati et al. (2021) demonstrated the importance of programmable intelligence in Open Radio Access Network (O-RAN) architectures, where machine learning applications can support data-driven network control. More recently, Nezami et al. (2025) examined agent-based architectures for radio access networks and emphasized the need for standardized agent capabilities, tool interfaces, coordination mechanisms, and validation procedures.

In a telecommunications environment, different agents may be responsible for separate network functions. A radio access agent may optimize spectrum utilization, while a transport agent manages traffic routing and a core network agent monitors service performance. Although each agent can make decisions within its domain, effective network-wide management requires these decisions to be compatible with shared service objectives and operational restrictions.

Interoperability Challenges in Multi-Vendor Networks

Multi-vendor telecommunications networks combine equipment, software, controllers, and management platforms from different manufacturers. While this approach offers flexibility and reduces dependence on individual suppliers, it introduces interoperability challenges arising from differences in communication interfaces, data structures, management policies, and implementation methods.

O-RAN has addressed part of this problem by promoting open interfaces and programmable control functions. Nevertheless, standardized infrastructure interfaces do not automatically ensure that independently developed AI agents understand one another's objectives or produce compatible decisions. Polese et al. (2023b) highlighted the architectural opportunities and technical challenges associated with open and intelligent radio access networks, including the integration of heterogeneous control applications.

Conflicting decisions present a particular concern. For example, an agent seeking to reduce network energy consumption may propose changes that conflict with another agent's objective of maintaining throughput or service coverage. Wadud et al. (2024) examined such conflicts among O-RAN applications and demonstrated the importance of considering service requirements when resolving competing control actions.

These challenges indicate that autonomous network coordination requires more than connectivity between network components. It also requires consistent interpretation of operational intentions, controlled information exchange, and mechanisms for reconciling incompatible decisions.

Research Problem and Motivation

Existing telecommunications management solutions provide important capabilities for automation, network analytics, and programmable control. However, coordinating independently developed AI agents across multiple network domains remains difficult when their communication models, decision-making procedures, and operational constraints differ.

Consider a network congestion event affecting radio access, transport, and core network resources. An agent operating within the radio access domain may recommend traffic redistribution, while a transport agent may select an alternative routing strategy. If these actions are implemented independently, the resulting configuration may fail to achieve the intended service improvement or create additional network instability.

Although existing research has addressed autonomous management and application-level conflict mitigation, there remains a need to investigate how heterogeneous agents can exchange meaningful operational information, negotiate shared objectives, and execute coordinated actions under common policy and security constraints (Coronado et al., 2022; Wadud et al., 2024; Nezami et al., 2025).

This study addresses that problem by proposing an open framework for interoperable AI-agent coordination in multi-vendor telecommunications environments.

Research Aim and Objectives

The aim of this study is to develop and evaluate an open, standards-aligned framework that enables heterogeneous AI agents to communicate, negotiate operational objectives, and coordinate autonomous network management activities across multiple telecommunications domains.

The specific objectives are to:

1. Examine existing approaches to autonomous network management and identify interoperability requirements for heterogeneous AI agents.
2. Develop a common communication and semantic interoperability model for exchanging network management intents and agent capabilities.
3. Design coordination mechanisms that support distributed decision-making, policy compliance, and conflict resolution.

4. Establish integration mechanisms between the proposed framework and relevant telecommunications management interfaces.
5. Evaluate the framework against selected baseline approaches using interoperability, latency, scalability, reliability, and network service performance metrics.

Research Questions and Hypotheses

Four research questions guide the investigation.

RQ1: How effectively can independently developed AI agents exchange and interpret network management intents across heterogeneous vendor environments?

H1: The proposed framework will achieve a higher cross-vendor task completion rate than a generic agent communication approach without telecom-specific semantic coordination.

RQ2: To what extent can coordinated AI agents improve network management outcomes across multiple telecommunications domains?

H2: The proposed framework will reduce service-level agreement violations and fault recovery delays compared with independently operating agents.

RQ3: How does increasing vendor heterogeneity and agent population affect interoperability performance?

H3: The proposed framework will exhibit a smaller degradation in coordination performance as agent and vendor populations increase than direct vendor-specific integration approaches.

RQ4: Can policy-aware agent coordination reduce conflicting and unauthorized network management actions?

H4: Explicit policy validation and conflict resolution will reduce the occurrence of unauthorized or incompatible network control actions compared with coordination approaches lacking these mechanisms.

Research Contributions

The study proposes a framework for addressing communication, semantic interpretation, coordination, and operational control challenges within a common architecture. Its intended contribution is not the replacement of existing telecommunications standards, but the definition of mechanisms through which independently developed agents can use them consistently.

The principal contributions are a layered interoperability architecture, a common semantic profile for network management intents, a policy-aware coordination mechanism, and a reproducible evaluation methodology. The proposed approach also incorporates agent discovery, capability identification, conflict resolution, and decision traceability.

Through comparative evaluation, the study seeks to establish the conditions under which these mechanisms improve cross-vendor coordination and identify the performance costs associated with introducing additional interoperability functions.

Organization of the Article

The remainder of the article is organized as follows. Section 2 reviews relevant studies and interoperability standards. Section 3 presents the system model and problem formulation, while Section 4 describes the proposed framework. Section 5 explains the research methodology and experimental design. Section 6 presents the experimental findings, followed by their interpretation in Section 7. Section 8 addresses research limitations and threats to validity, and Section 9 concludes the article.

Background and Related Work

Autonomous and Self-Organizing Telecommunications Networks

Autonomous telecommunications management has developed from conventional network automation and self-organizing networks toward systems capable of continuously monitoring operational conditions and adjusting network behavior. Early self-organizing approaches concentrated on functions such as automatic configuration, performance optimization, and fault recovery. More recent approaches extend these capabilities through closed-loop management, where network observations are used to evaluate performance and determine corrective actions.

Benzaid and Taleb (2020) examined AI-driven zero-touch network management and identified the importance of intelligent decision-making in reducing dependence on manual operations. Similarly, Coronado et al. (2022) reviewed network automation approaches for 5G and 6G, highlighting the role of integrated management functions in achieving autonomous network operations.

El Rajab et al. (2024) further examined zero-touch networks, identifying applications in traffic monitoring, optimization, energy management, and security. However, the effectiveness of

autonomous management depends on the ability to coordinate actions across interacting network functions. When separate management systems operate under different policies or control models, local improvements may not produce desirable network-wide outcomes.

Multi-Agent Artificial Intelligence for Network Operations

Multi-agent systems provide an approach to managing complex environments through multiple autonomous components with different responsibilities. In telecommunications networks, agents may specialize in monitoring, fault identification, resource optimization, policy enforcement, or service recovery. Their decisions may be coordinated through centralized orchestration, hierarchical control, or distributed negotiation.

Distributed coordination offers flexibility because agents can operate close to the network functions they manage. However, this arrangement introduces challenges involving inconsistent information, conflicting objectives, communication delays, and shared resource constraints. The problem becomes more demanding when agents are implemented by different vendors and do not share the same internal representations.

Nezami et al. (2025) identified multi-agent orchestration, capability discovery, lifecycle management, and agent-specific validation as important areas for standardization in intelligent radio access networks. Their work also emphasized the need to account for telecommunications requirements when applying general-purpose AI-agent architectures.

Unlike conventional software agents, network control agents must operate within defined authorization boundaries and timing constraints. Consequently, autonomous coordination must account for both the quality of decisions and the safety of their execution.

Standardized Network Management and Interoperability

Several existing telecommunications standards and protocols provide functions relevant to autonomous network coordination. These technologies address different parts of the management process, including network access, analytics, service intent, and communication between intelligent components. Understanding their respective responsibilities is necessary to avoid duplicating established mechanisms in the proposed framework.

O-RAN Architecture and Open Interfaces

The O-RAN architecture promotes disaggregated radio access networks through open interfaces and programmable management functions. Its principal components include the Service Management and Orchestration (SMO) framework, non-real-time RAN Intelligent Controller (non-RT RIC), and near-real-time RAN Intelligent Controller (near-RT RIC).

Key interfaces include A1 for policy and related interactions between the non-RT RIC and near-RT RIC, E2 for communication between the near-RT RIC and supported RAN nodes, O1 for operations and management, and O2 for cloud infrastructure management and orchestration. These interfaces support integration across different network components and provide opportunities for intelligent control applications (Polese et al., 2023b).

Bonati et al. (2021) examined how O-RAN supports machine learning-based control, while Lacava et al. (2025) investigated real-time AI-based control through distributed applications. Although these developments improve the programmability of radio networks, the coordination of independently developed AI agents across radio, transport, and core domains requires additional attention to semantic compatibility, policy consistency, and shared operational objectives.

ETSI Zero-Touch Network and Service Management

The European Telecommunications Standards Institute (ETSI) developed the Zero-Touch Network and Service Management (ZSM) framework to support automated management and orchestration across heterogeneous network domains. Its architectural approach emphasizes end-to-end automation, management service integration, and closed-loop operations.

Benzaid and Taleb (2020) described ZSM as an important foundation for intelligent network operations and discussed challenges associated with introducing AI into automated management systems. Coronado et al. (2022) further examined how network automation solutions relate to the broader development of zero-touch management.

ZSM provides a relevant architectural foundation for the proposed research because it addresses cross-domain management rather than restricting automation to individual network functions. However, a framework for independently developed AI agents must also specify how capabilities are described, intentions are interpreted, and disagreements are resolved across different implementations.

3GPP Service-Based Architecture and Network Analytics

The Third Generation Partnership Project (3GPP) supports modular communication between 5G core network functions through its Service-Based Architecture (SBA). An important component is the Network Data Analytics Function (NWDAF), which provides analytics that authorized network functions can use to support operational decisions.

The functions and services of NWDAF are specified in 3GPP TS 23.288 and TS 29.520. These specifications establish mechanisms for requesting and receiving analytics concerning network conditions, service performance, and other operational characteristics.

In autonomous management, network analytics can provide information required by agents involved in congestion control, resource allocation, and fault management. Nevertheless, the availability of common analytics does not establish a mechanism for coordinating independent agent decisions. A separate interoperability function is therefore needed to determine how analytics are interpreted, shared, and translated into authorized cross-domain actions.

TM Forum Intent Management

Intent-based networking allows operators to describe desired service outcomes without specifying every underlying configuration action. This approach separates operational requirements from the technical procedures used to satisfy them.

Leivadeas and Falkner (2023) reviewed intent-based networking and identified intent expression, translation, resolution, activation, and assurance as central components of autonomous management. These functions provide a foundation for translating service requirements into network behavior.

TM Forum TMF921 supports the exchange, negotiation, and reporting of intents between an intent owner and an intent handler. In addition, RFC 9315 provides conceptual definitions that distinguish intent-based management from conventional policy and configuration mechanisms.

For heterogeneous AI agents, these approaches offer a basis for expressing common objectives. However, agents must still determine whether a requested objective is feasible, identify the relevant participating functions, and reconcile competing proposals before initiating network changes.

Agent-to-Agent Communication Protocols

Agent-to-agent communication protocols address interoperability between intelligent software components developed using different tools, programming languages, and internal architectures.

The Agent2Agent (A2A) protocol, introduced in 2025, provides mechanisms for capability discovery, structured message exchange, and collaborative task management.

These functions are relevant to telecommunications because independently developed agents require a common method for identifying available capabilities and exchanging requests. However, general-purpose communication protocols do not, by themselves, define network resource semantics, service-level constraints, or the authorization procedures required for operational changes.

Nezami et al. (2025) emphasized the importance of agent-specific interfaces, validation requirements, and standardized orchestration in radio access networks. Their analysis supports the need for telecommunications-specific extensions that allow general communication mechanisms to operate within established network management and safety boundaries.

Existing Research on Multi-Vendor Network Coordination

Previous studies have established important foundations for intelligent network management, particularly within open radio access networks. Bonati et al. (2023) introduced OpenRAN Gym, an experimental environment designed to support data collection, machine learning development, and testing of O-RAN control solutions. Its contribution lies in providing a practical platform for evaluating intelligent network applications under controlled conditions.

Polese et al. (2023a) presented ColO-RAN, demonstrating how machine learning-based applications can support closed-loop RAN control on programmable experimental platforms. Their work provides evidence that intelligent control functions can be implemented and evaluated using open network architectures.

Lacava et al. (2024) investigated programmable traffic steering using O-RAN, illustrating how intelligent applications can influence network resource management. These studies primarily examine specific control functions and their effects on radio access network performance.

Conflict management has received more direct attention from Wadud et al. (2024), who proposed a QoS-aware mechanism for mitigating conflicts among xApps. Their findings show why independently generated control decisions must be assessed against shared performance requirements.

Despite these contributions, the literature reviewed here provides less direct evidence concerning semantic interoperability and coordinated decision-making among independently developed

agents operating across multiple telecommunications domains. This distinction forms the basis of the present investigation.

Comparative Analysis of Relevant Standards and Technologies

The reviewed standards address complementary aspects of network automation and intelligent system integration. Table 1 summarizes their functions and relevance to autonomous agent interoperability.

Table 1. Existing standards and their relevance to autonomous agent interoperability

Standard or technology	Principal functionality	Relevance to the proposed framework
O-RAN Alliance architecture	Open radio access interfaces, programmable RAN control, and management integration	Provides standardized interfaces for radio access agents and control systems
ETSI ZSM	End-to-end network and service management automation	Provides an architectural basis for cross-domain coordination and closed-loop management
3GPP NWDAF	Standardized network analytics services within the 5G core	Supplies analytics that can inform agent decisions
TM Forum TMF921	Intent expression, negotiation, and reporting	Supports communication of service objectives and management requirements
Agent2Agent (A2A) protocol	Agent discovery, communication, and collaborative task exchange	Provides a general communication foundation for heterogeneous AI agents
IRTF RFC 9315	Intent-based networking concepts and definitions	Supports consistent interpretation of intent fulfillment and assurance

Table 1 shows that existing technologies provide important building blocks but operate at different architectural levels. O-RAN and 3GPP primarily define network-specific interfaces and services, while ETSI ZSM addresses broader management architecture. TM Forum focuses on intent management, and A2A provides general communication mechanisms for intelligent agents. RFC 9315 establishes a shared conceptual foundation for intent-based networking.

The proposed framework seeks to connect these capabilities through a common semantic model and policy-aware coordination mechanism. Its value must therefore be assessed by determining

whether this integration provides measurable benefits beyond the individual standards and existing orchestration approaches.

Identified Research Gaps

The reviewed literature indicates several areas requiring further investigation.

First, studies on intelligent network control have concentrated largely on individual network domains or particular application types. Although O-RAN supports programmable and multi-vendor environments, evidence concerning independently developed agents coordinating across radio access, core, and transport networks remains limited within the selected studies (Bonati et al., 2023; Polese et al., 2023a).

Second, technical interface compatibility does not guarantee semantic interoperability. Agents may exchange valid messages while interpreting network objectives, resource descriptions, or operational constraints differently. Existing work on intent-based networking provides useful foundations, but these must be linked to agent-level capability descriptions and executable coordination procedures (Leivadeas & Falkner, 2023; Nezami et al., 2025).

Third, conflicting decisions remain a concern when autonomous components influence shared resources. Although Wadud et al. (2024) addressed conflict mitigation among O-RAN applications, further investigation is needed into conflicts involving different network domains, administrative policies, and levels of control authority.

Fourth, security mechanisms must extend beyond protecting network interfaces. Autonomous coordination requires the verification of agent identities, permissions, proposed actions, and execution outcomes. Research by Groen et al. (2025) and Atalay et al. (2025) demonstrates the importance of securing programmable network components, but cross-domain agent authorization and decision accountability present additional design challenges.

Finally, reproducible experimental comparisons are needed to establish how agent interoperability affects service performance, coordination latency, scalability, and operational resilience. Existing testbeds provide useful starting points, although additional workloads and conformance tests are necessary to evaluate heterogeneous agent interaction directly.

These gaps motivate a framework that addresses communication compatibility, semantic interpretation, conflict resolution, and policy enforcement within a single experimental design.

Research Positioning

This study builds on established telecommunications standards rather than proposing an alternative network management infrastructure. The intended contribution is a coordination layer that allows independently developed AI agents to discover relevant capabilities, exchange consistently structured intents, negotiate actions, and verify compliance with network policies.

The framework will combine general agent communication mechanisms with telecommunications-specific semantic descriptions, standardized network interfaces, and explicit decision validation procedures. Its design will distinguish between communicating an intention, authorizing a proposed action, and confirming that the resulting network state satisfies the requested objective.

Unlike studies focused primarily on individual RAN control applications, the proposed research will assess coordination across multiple network domains and heterogeneous agent implementations. Comparative experiments will examine whether the proposed mechanisms improve interoperability and operational outcomes, while also measuring their communication and processing overhead.

The resulting evaluation will provide evidence of the framework's practical strengths, limitations, and suitability for future autonomous telecommunications networks.

System Model, Interoperability Requirements, and Problem Formulation

Multi-Vendor Telecommunications Network Model

The proposed system considers a heterogeneous telecommunications environment comprising multiple network domains, management platforms, and autonomous AI agents developed by different vendors. The network is organized into four functional domains: the Radio Access Network (RAN), 5G Core Network, transport and backhaul network, and service orchestration and assurance domain. Each domain contains network resources and management functions responsible for specific operational tasks, including traffic management, resource allocation, fault detection, and service performance monitoring.

The adoption of open network architectures has improved the integration of equipment and software from different vendors. However, differences in management interfaces, data representations, and control procedures continue to complicate coordinated network operations.

Polese et al. (2023b) identified interoperability, intelligent control, and integration among heterogeneous components as important considerations in O-RAN deployments. Similarly, Coronado et al. (2022) emphasized the importance of coordinated management functions for achieving end-to-end automation in 5G and emerging 6G networks.

To represent the network, let the set of operational domains be defined as:

$$\mathcal{D} = \{D_{\text{RAN}}, D_{\text{Core}}, D_{\text{Transport}}, D_{\text{Service}}\}$$

Each domain D_j contains a set of network resources R_j, management functions, and associated AI agents. The proposed framework introduces a common coordination layer through which agents can exchange operational information, interpret shared objectives, and negotiate management actions without requiring identical internal implementations.

For example, during a congestion event, a RAN agent may recommend traffic redistribution, while a transport agent identifies alternative routing paths and a core network agent evaluates service-level implications. The coordination layer assesses these proposals against shared constraints before permitting authorized network changes. This arrangement is intended to preserve local management autonomy while improving the consistency of decisions affecting multiple domains.

The system supports communication through standardized network interfaces and vendor-specific adapters where necessary. O-RAN interfaces provide access to supported radio network management functions, while relevant 3GPP services supply core network analytics and management information. These interfaces remain responsible for executing network operations within their respective control boundaries.

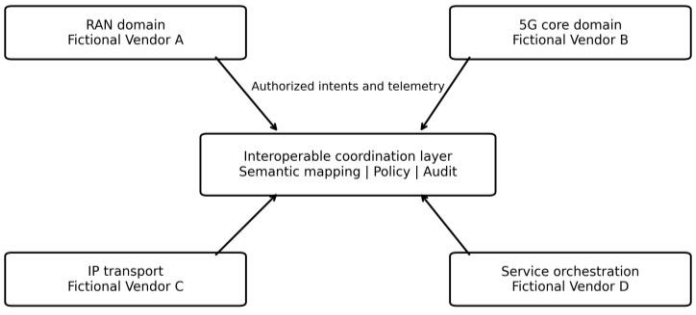


Figure 1. Proposed multi-domain agent coordination model with fictional vendor assignments (conceptual, not deployed).

Agent and Network Resource Models

The system consists of autonomous agents with different operational responsibilities, capabilities, and access privileges. Each agent observes network conditions within its authorized scope and may generate recommendations, respond to coordination requests, or initiate approved management actions. Agents are not assumed to possess complete knowledge of the network, since observations may be restricted by vendor interfaces, administrative boundaries, and communication delays.

The set of participating agents is expressed as:

$$\mathcal{A} = \{a_1, a_2, \dots, a_n\}$$

where $a_i$represents an individual AI agent and \(n\) denotes the total number of participating agents.

Each agent is characterized by:

$$a_i = (D_i, C_i, O_i, U_i, P_i)$$

Where D_i is its assigned network domain C_i represents supported capabilities, O_i denotes accessible observations, U_i defines permissible management actions, and P_i specifies applicable operational policies.

Three categories of agents are considered. Monitoring agents collect telemetry and identify changes in network conditions. Advisory agents analyze available information and recommend corrective actions. Control agents execute approved configurations through authorized network interfaces. This distinction separates decision support from actual control authority and reduces the possibility of unrestricted network modifications.

Network resources are represented by:

$$\mathcal{R} = \{r_1, r_2, \dots, r_m\}$$

where m is the number of managed resources, including spectrum, bandwidth, processing capacity, storage, and virtualized network functions. Each resource has an available capacity, an operational state, and applicable allocation constraints.

The system state at time $t$is represented by x_t, which combines relevant observations from participating domains. Because network measurements may be incomplete or delayed, agents must consider the validity and freshness of exchanged information before selecting coordinated actions. Nezami et al. (2025) emphasized the importance of clearly defined agent capabilities, standardized interaction mechanisms, and coordinated management procedures in agent-based radio networks. These principles inform the separation of agent capabilities, resource ownership, and control privileges adopted in the present model.

Functional and Non-Functional Interoperability Requirements

Interoperability in multi-vendor telecommunications networks involves more than establishing communication between autonomous agents. Effective coordination requires agents to identify compatible functions, interpret exchanged information consistently, comply with operational policies, and respond appropriately when network conditions change.

The proposed framework therefore considers both functional requirements, which determine the operations agents must support, and non-functional requirements, which define acceptable performance, reliability, and security characteristics. Leivadeas and Falkner (2023) highlighted the importance of consistent intent interpretation and fulfillment in intent-based networking, while Wadud et al. (2024) demonstrated the operational significance of managing conflicting decisions among intelligent RAN applications.

Table 2. Functional and non-functional requirements for interoperable autonomous network agents

Requirement	Description	Proposed validation method
R1: Agent discovery	Agents must identify available peers, supported capabilities, and accessible management functions.	Capability registration and discovery tests
R2: Semantic compatibility	Agents must interpret exchanged network intents and resource descriptions consistently.	Cross-vendor schema validation and semantic consistency tests

R3: Policy compliance	Proposed actions must satisfy authorization rules, service requirements, and network operating constraints.	Policy enforcement and authorization tests
R4: Conflict resolution	The system must identify and reconcile incompatible decisions involving shared resources or objectives.	Conflicting-intent experiments
R5: Scalability	Coordination performance should remain acceptable as the number of agents and participating vendors increases.	Agent population and workload scaling tests
R6: Fault tolerance	The framework must maintain safe operation or recover appropriately following communication and component failures.	Controlled fault injection and recovery tests
R7: Traceability	Agent decisions, exchanged messages, approvals, and executed actions must be recorded for accountability.	Audit-log and event-trace verification
R8: Latency awareness	Coordination procedures must respect applicable network management deadlines.	End-to-end coordination latency measurements

These requirements establish the basis for the framework architecture and experimental evaluation. Agent discovery and semantic compatibility support communication across heterogeneous implementations, while policy compliance and conflict resolution address operational safety. Scalability, fault tolerance, traceability, and latency awareness determine whether the framework can support practical telecommunications management.

The requirements also recognize that interoperability cannot be measured solely by successful message delivery. A coordination attempt is considered successful only when the participating agents interpret the requested objective consistently and produce a valid outcome that satisfies the applicable network constraints.

Threat and Failure Model

Autonomous coordination introduces operational and security risks because multiple independently developed agents exchange information and influence shared network resources.

The proposed threat and failure model considers communication faults, inconsistent network information, incompatible management decisions, and unauthorized actions.

Communication failures may occur when an agent becomes unavailable, a management interface experiences disruption, or network messages are delayed beyond their useful execution period. Such events may prevent agents from receiving current observations or completing coordination requests. The framework must therefore distinguish temporary communication failures from rejected requests and avoid executing actions based on expired information.

Semantic incompatibility represents another failure condition. Two agents may exchange technically valid messages but interpret the associated network intent differently. This can result in unsuitable resource allocations or inconsistent management decisions. Similarly, agents may independently propose changes that compete for the same resources or produce contradictory operational outcomes. Wadud et al. (2024) demonstrated the importance of detecting and mitigating such conflicts in intelligent Open RAN control environments.

The security model considers unauthorized access attempts, agent impersonation, replayed messages, and misleading management requests. Groen et al. (2025) examined security challenges affecting O-RAN interfaces and intelligent network applications, while Atalay et al. (2025) addressed authentication, authorization, and secure application discovery in Open RAN environments.

The proposed framework assumes that participating administrative domains establish trusted identities and define explicit permissions for agents. Authentication establishes the identity of a requesting agent, while authorization determines whether the requested action is permissible within the affected network domain. Messages must also carry sufficient contextual information to support freshness checks, policy validation, and subsequent auditing.

When a coordination request violates an established policy, exceeds available resources, or cannot be resolved safely, the system must reject the action or escalate it to an authorized management function. This prevents coordination failures from automatically becoming network configuration failures.

The model does not assume that all participating agents are trustworthy or that every communication attempt will succeed. However, the proposed evaluation is limited to defined threats and fault conditions rather than providing a guarantee against all forms of coordinated malicious behavior.

Mathematical Problem Formulation

The principal coordination problem is to determine a set of compatible agent actions that satisfies network management objectives while minimizing service disruption, decision latency, and communication overhead.

Let $I = \{I_1, I_2, \ldots, I_k\}$ denote the set of admitted network management intents. Each intent contains an outcome objective, resource constraints, priority, and, where applicable, an execution deadline. For agent a_i, let $u_i(t)$ denote its proposed action at time t. The joint action across n agents is:

$$\mathbf{u}(t) = \{u_1(t), u_2(t), \ldots, u_n(t)\}$$

The network state evolves according to:

$$x_{t+1} = f(x_t, \mathbf{u}(t), w_t)$$

Here f is the network transition function and w_t represents disturbances, such as traffic variation, component failure, and measurement uncertainty.

The candidate coordination policy π is selected to minimize normalized SLA cost, decision delay, and signaling overhead subject to safety and feasibility constraints:

$$\min_{\pi} \; J(\pi) = \alpha \overline{C}_{\text{SLA}}(\pi) + \beta \overline{C}_{\text{delay}}(\pi) + \gamma \overline{C}_{\text{overhead}}(\pi)$$

where:

- π denotes the joint coordination policy.
- $\bar{C}$_SLA(π) denotes the normalized SLA-violation cost.
- $\bar{C}$_delay(π) denotes normalized coordination delay.
- $\bar{C}$_overhead(π) denotes normalized coordination communication overhead.
- α, β, and γ are nonnegative weights satisfying α + β + γ = 1.

For an implementable illustrative design, $\alpha = \beta = \gamma = 1/3$; potential normalization references are 0.10 for SLA violation fraction, 1,000 ms for coordination delay, and 1 KB per task for signaling. These are proposed design defaults, not fitted values or inputs to the statistical generator in Section 5. Candidate actions violating a hard policy, authorization rule, capacity limit, or deadline are excluded regardless of their objective score.

The optimization is subject to four principal constraints.

Resource capacity constraint. The combined allocation to a resource cannot exceed its available capacity:

$$\sum_{i=1}^{n} q_{ij}(t) \leq K_j(t), \quad \forall r_j \in \mathcal{R}$$

Here $q_{ij}(t)$ is the amount of resource r_j requested by agent a_i and $K_j(t)$ is the capacity available at time t.

Authorization constraint

An agent may execute only actions permitted by its assigned privileges:

$$u_i(t) \in U_i^{\text{authorized}}, \quad \forall a_i \in \mathcal{A}$$

U_i^authorized denotes the permitted action set for agent a_i.

Policy and conflict constraint

The selected joint action must satisfy the applicable operational policies and must not contain unresolved conflicts:

$$\mathbf{u}(t) \in \mathcal{F}(x_t, P)$$

$F(x_t,P)$ denotes the set of feasible joint actions under current network state x_t and operational policy set P.

Coordination deadline constraint

For intents associated with execution deadlines, the time required to reach an authorized coordination outcome must remain within the permitted interval:

$$T_{\text{coord}}(I_k) \leq T_{\max}(I_k)$$

$T_coord(I_k)$ is the coordination duration of intent I_k and $T_max(I_k)$ is its deadline.

A request for which no feasible action can be identified is rejected or escalated rather than executed. Such requests are reported separately in the performance evaluation to prevent unsafe decisions from being counted as successful coordination.

The formulation reflects the need to balance service objectives with network resource limitations and operational safety. Foukas et al. (2017) identified resource allocation and service isolation as important network slicing challenges, while Wadud et al. (2024) demonstrated the relevance of QoS-aware conflict mitigation in systems containing multiple intelligent control applications.

Design Assumptions and Scope

The proposed system operates under several assumptions intended to establish a consistent basis for framework development and experimental evaluation.

First, participating vendors are assumed to expose documented management interfaces or provide adapters capable of translating supported operations into the proposed common representation. The framework does not require access to proprietary internal algorithms or unrestricted modification of vendor equipment.

Second, agents are assumed to possess different levels of network visibility and control authority. Information exchange is restricted to data necessary for the requested coordination task, while configuration changes are performed through authorized controllers or management systems. This separation maintains domain ownership and reduces the likelihood of unintended network modifications.

Third, the system assumes that network domains can establish trusted identities and exchange authenticated coordination messages. However, these relationships do not imply unrestricted trust, since each proposed action remains subject to authorization and policy validation.

Fourth, network observations may be delayed, incomplete, or temporarily unavailable. Coordination procedures must therefore account for stale information and communication failures.

The system is intended primarily for network management and orchestration activities whose timing requirements can accommodate the additional coordination process. Direct control of functions requiring stricter timing guarantees remains within the relevant domain-specific controllers.

The present numerical analysis does not instantiate any network emulator, vendor adapter, agent service, or hardware controller. It uses a stochastic rule-based generator with fictional vendor identifiers. Actual testbed integration remains a subsequent validation task.

These assumptions establish the operational boundaries of the proposed framework. They also provide a basis for evaluating whether interoperable AI-agent coordination can improve network management while preserving resource constraints, service requirements, and established control responsibilities.

Proposed Open Framework for Autonomous Network Coordination

Architectural Design Principles

The proposed Open Framework for Autonomous Network Coordination (OFANC) is designed to support communication and coordinated decision-making among AI agents operating across heterogeneous telecommunications networks. The framework addresses differences in vendor interfaces, data representations, management objectives, and operational policies without requiring participating agents to use identical internal architectures.

The design follows five principles: vendor neutrality, modularity, semantic consistency, standards compatibility, and secure coordination. Vendor neutrality allows agents developed by different manufacturers to participate through documented interfaces. Modularity separates communication, data interpretation, and decision-making functions, allowing individual components to be modified without disrupting the entire system. Semantic consistency ensures that exchanged information retains its intended meaning across implementations.

Standards compatibility is particularly important because existing telecommunications architectures already provide established mechanisms for network management and control. Polese et al. (2023b) emphasized the importance of open interfaces and programmable control in heterogeneous radio access networks. Similarly, Nezami et al. (2025) identified standardized agent interactions, capability descriptions, and validation procedures as important requirements for intelligent network operations.

The proposed framework therefore extends existing management capabilities through a common coordination mechanism rather than replacing standardized network interfaces.

Layered Framework Architecture

OFANC adopts a five-layer architecture that separates network integration from agent communication, semantic interpretation, autonomous coordination, and service governance. Each layer provides specific functions while exchanging information through defined interfaces.

The first layer, Network Domain Adaptation, connects the framework to radio access, core, transport, and service management systems. It translates supported vendor-specific operations into common management requests and forwards authorized actions to the appropriate controllers. Integration may use O-RAN, 3GPP, and other supported network management interfaces.

The second layer, Agent Communication and Trust, supports agent registration, capability discovery, authenticated communication, message exchange, and audit logging. It establishes the communication mechanisms required for agents to exchange information while maintaining their respective authorization boundaries.

The third layer, Semantic Interoperability, provides common representations for network resources, operational intents, agent capabilities, and management constraints. It validates incoming messages and resolves differences in terminology and data structures through registered mappings.

The fourth layer, Autonomous Coordination, manages intent negotiation, action selection, conflict detection, and decision validation. It evaluates proposed actions against available resources and applicable operational policies before recommending an execution plan.

The fifth layer, Service Intent and Governance, receives network-level objectives, defines service constraints, and supervises the fulfillment of approved intents. It also supports operational monitoring, accountability, and escalation when automated coordination cannot produce an acceptable decision.

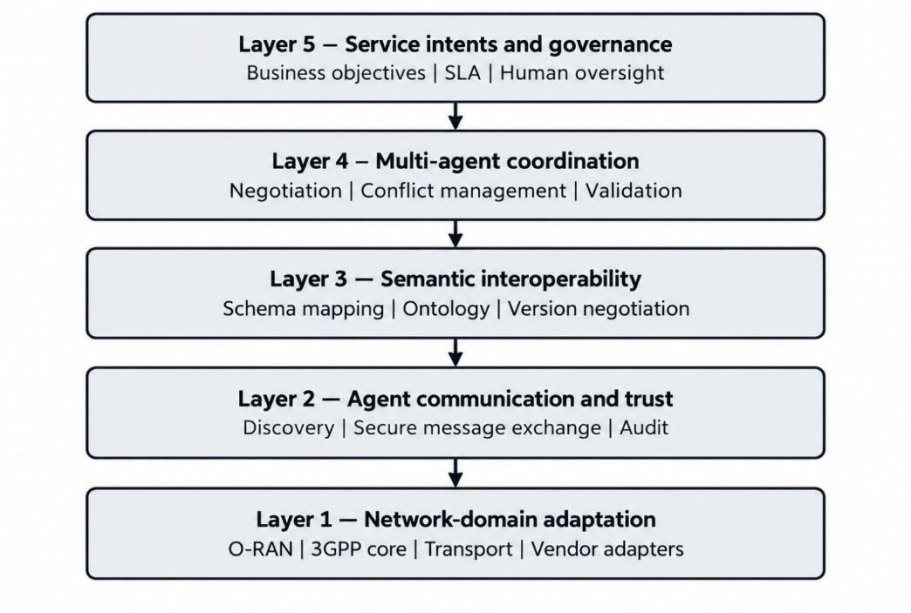


Fig 2. Proposed five-layer architecture of the Open Framework for Autonomous Network Coordination.

Agent Discovery and Capability Advertisement

Effective coordination begins with identifying agents capable of performing the operations required by a network management intent. The proposed framework introduces a capability registry through which agents advertise their supported functions, network domains, communication protocols, and operational permissions.

Each agent publishes a capability description containing its identifier, domain association, supported operations, input requirements, output formats, authorization scope, and interface version. Capability descriptions may also specify expected response times and operational restrictions.

When an intent is submitted, the coordination service examines the registry to identify agents whose capabilities match the requested objective. Agents that cannot satisfy the functional or authorization requirements are excluded from the subsequent negotiation process.

The framework can reuse general agent discovery and task communication mechanisms, including those introduced by the Agent2Agent (A2A) protocol in 2025, while defining additional telecommunications-specific capability fields. Nezami et al. (2025) similarly emphasized the

importance of standardized tool specifications and agent discovery for effective coordination in intelligent radio access networks.

Capability advertisements are subject to authentication and version validation. Registry entries that become unavailable or outdated are marked accordingly to reduce the likelihood of assigning tasks to unsuitable agents.

Common Semantic Interoperability Profile

Communication compatibility alone does not guarantee that agents interpret operational instructions consistently. For this reason, OFANC introduces a Common Semantic Interoperability Profile (CSIP) for representing network management information exchanged between heterogeneous agents.

The profile defines a versioned message envelope containing the operational objective, target network domain, required capabilities, network context, and execution constraints. It is intended to support consistent message interpretation while allowing vendor-specific information to remain within domain adapters.

Table 3. Proposed semantic message structure for interoperable telecommunications AI agents

Data field	Description
Agent identifier	Authenticated identifier of the sending or receiving agent
Intent identifier	Unique reference assigned to the network management objective
Target domain	Network domain affected by the requested operation
Requested objective	Description of the desired network state or service outcome
Capability reference	Required agent functionality or supported management operation
Network context	Relevant network measurements, resource states, and observation timestamps
Performance constraints	Service-level, quality-of-service, latency, and resource requirements
Policy reference	Applicable authorization rules and operational restrictions
Priority and deadline	Relative urgency and permitted coordination time
Message version	Schema version used for compatibility validation

Provenance reference	Information required to trace the origin and handling of a coordination request

The semantic profile supports three validation stages. First, syntactic validation confirms that messages conform to the required schema. Second, semantic validation checks whether the referenced resources, capabilities, and operational objectives have recognized meanings. Third, contextual validation determines whether the information is sufficiently current and relevant to the affected network domain.

Schema translation is performed through versioned mappings maintained by the interoperability layer. Where a message contains an unsupported concept, ambiguous resource reference, or incompatible version, the receiving component returns a structured error rather than attempting an uncertain interpretation.

Leivadeas and Falkner (2023) identified intent interpretation and translation as important challenges in intent-based networking. The proposed semantic profile addresses these concerns at the agent communication level by establishing a common structure for representing management objectives and associated constraints.

Worked CSIP example (design specification, not an executed message): {"schema_version":"1.0","intent_id":"INT-0042","source_agent":"svc-agent-01","target_domain":["RAN","transport"],"objective":{"type":"restore_throughput","minimum_mbps":100},"context":{"observed_at":"2025-10-01T10:00:00Z","cell_id":"cell-07"},"constraints":{"p95_latency_ms":50,"deadline_ms":5000},"policy_ref":"POL-07","priority":"service-assurance","capability_ref":"congestion-mitigation","trace_id":"TRACE-0042"}.

Schema mapping is explicit: a vendor field such as cellResourceId is translated to the canonical cell_id only after an adapter validates the resource identity; throughput units are converted to Mbps; timestamps are converted to UTC; unknown units, unsupported semantic versions, missing policies, or untrusted provenance produce a rejection rather than a guessed translation. The example payload and representative topology are included in the reproducibility supplement.

Distributed Intent Negotiation

The coordination mechanism enables agents to negotiate compatible management actions while retaining responsibility for their assigned network domains. Negotiation begins when the service governance layer receives an intent describing a desired operational outcome.

The coordination service validates the request and identifies agents capable of contributing to its fulfillment. Selected agents examine their current network conditions and submit proposed actions together with relevant resource requirements and expected operational effects.

The coordination mechanism evaluates these proposals to determine whether they collectively satisfy the requested intent. When disagreements arise, affected agents may revise their proposals within a predefined negotiation limit.

For example, during network congestion, a RAN agent may recommend traffic redistribution while a transport agent proposes an alternative route. The coordination service assesses the combined effects of both actions before selecting an acceptable configuration.

The negotiation process follows five states: request validation, agent discovery, proposal exchange, decision validation, and outcome verification. Each intent is assigned a deadline and a maximum number of negotiation rounds. Requests that cannot be resolved within these limits are rejected or escalated to an authorized management function.

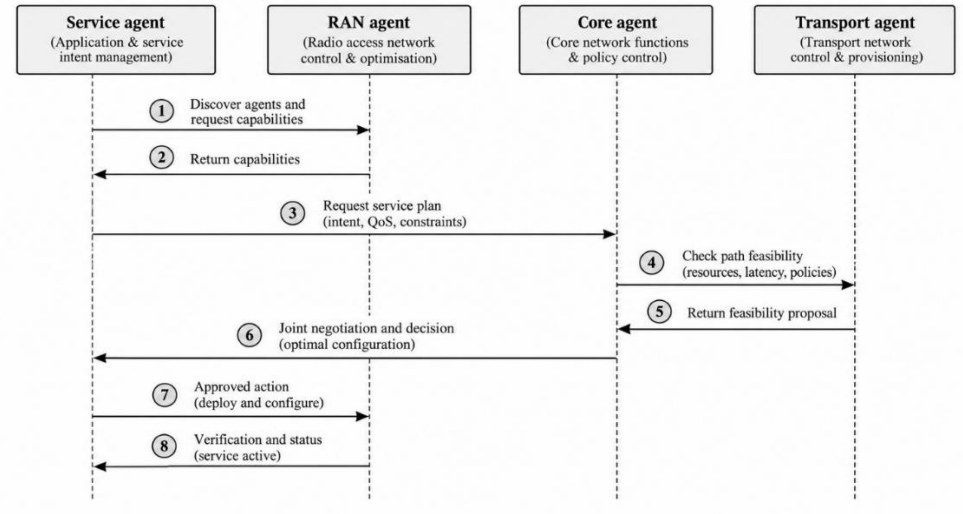


Fig 3. Agent-to-agent intent negotiation sequence for cross-domain telecommunications management.

Illustrative cross-domain negotiation trace. (1) A service-assurance agent submits INT-0042 requesting at least 100 Mbps and a p95 delay below 50 ms. (2) The capability registry selects a

RAN congestion agent and a transport routing agent; the core analytics agent provides a timestamped demand estimate. (3) The RAN agent proposes a 20 Mbps capacity shift, while the transport agent proposes an alternate route with an additional 8 ms estimated delay. (4) The coordinator rejects any candidate exceeding the delay requirement, seeks one revised transport route, and passes a feasible joint proposal to domain-specific policy gates. (5) Controllers prepare authorized actions using expected configuration versions; the coordinator commits only after all required preparations succeed and verifies the resulting service state. These quantities illustrate message semantics, not observed performance.

Priority and execution rules. Authorization, safety constraints, and resource feasibility are mandatory filters. Among valid intents, the coordinator first honors safety-critical service class, then earliest deadline, then configured service priority. Within the same class, feasible joint actions are ranked by the normalized objective in Section 3.5; ties use stable action identifiers. A configurable example permits three negotiation rounds and a five-second intent deadline, with a one-second controller preparation timeout. Timeouts trigger cancellation of uncommitted preparations. Partial commits require compensating rollback where supported, otherwise service freeze and operator escalation. These values are proposed control settings and were not exercised by the numerical generator.

Conflict Detection and Resolution Mechanism

Conflicting decisions may arise when agents attempt to modify the same network resources or pursue objectives that cannot be satisfied simultaneously. OFANC introduces a conflict management mechanism that evaluates proposed actions before they are submitted for execution. The mechanism distinguishes three forms of conflict. A direct conflict occurs when agents request incompatible changes to the same parameter or resource. An indirect conflict arises when separate actions influence a shared performance objective. A policy conflict occurs when an otherwise feasible action violates established operational restrictions.

Wadud et al. (2024) demonstrated the importance of QoS-aware conflict mitigation among O-RAN xApps, particularly where independently developed applications influence overlapping network control parameters. The proposed framework extends this concern to coordination across multiple telecommunications domains.

Conflict detection begins by identifying resources and service objectives affected by each proposal. Actions are then evaluated for resource feasibility, policy compliance, and operational compatibility.

When a conflict is identified, the framework attempts to resolve it by revising resource allocations, selecting alternative actions, or renegotiating affected proposals. Priority may guide the negotiation process, but it does not override mandatory security, authorization, or safety constraints.

A validated plan is submitted to the relevant domain controllers for preparation and authorization. Where supported, the controllers confirm resource availability and configuration versions before changes are committed. If preparation fails or subsequent verification identifies an unacceptable outcome, the framework initiates a defined recovery procedure. Compensating actions are used where rollback is supported, while irreversible operations require additional safeguards.

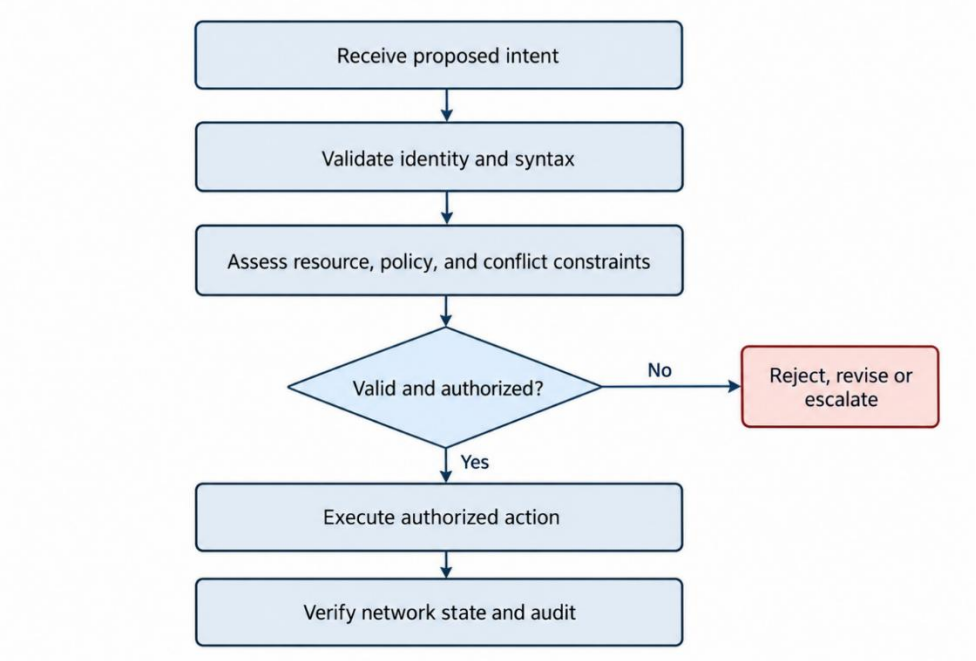


Fig 4. Policy-aware conflict detection and resolution workflow for autonomous network coordination.

Coordination Algorithm and Computational Complexity

The coordination algorithm combines agent discovery, proposal collection, policy validation, and conflict resolution within a bounded execution procedure.

Let I_k represent an incoming intent, $\mathcal{A}_k$ the set of eligible agents, and $\mathcal{U}_k$ the candidate actions proposed by those agents.

Algorithm 1. Policy-Aware Multi-Agent Coordination

Input:

Intent I

Agent capability registry A

Network state X

Policy set P

Maximum negotiation rounds Rmax

Output:

Verified outcome, rejection, or escalation

1. Validate the intent schema, identity, and deadline.

2. Discover agents with the required capabilities.

3. Authenticate eligible agents and verify permissions.

4. FOR round = 1 TO Rmax:

Collect candidate actions from eligible agents.

Validate message semantics and observation freshness.

Identify shared resource dependencies.

Detect incompatible actions and policy violations.

IF conflicts exist:
Request feasible alternative proposals.
Continue to the next round.

Construct a joint coordination plan.

Verify resource and policy constraints.

Request preparation from affected domain controllers.

IF every required controller confirms readiness:
Authorize and commit the plan.
Verify the resulting network state.

IF the operational objective is satisfied:
Record success and return verified outcome.

Otherwise:
Apply supported recovery procedures.

Reject or escalate unresolved requests.

Record the final outcome and audit information.

The algorithm terminates when the requested objective is verified, the intent deadline expires, or the maximum number of negotiation rounds is reached.

Feasible-plan selection is deliberately separated from message exchange. First, reject candidates with invalid schema, stale context, insufficient capability, missing authorization, violated resource limits, or missed deadlines. Second, construct dependency edges for proposed changes affecting shared resources. Third, enumerate compatible candidate combinations with branch pruning and a

configured search cap; choose the lowest-scoring feasible combination using the weighted objective in Section 3.5. If no valid plan exists before the deadline, reject or escalate. The reported model-based results were not produced by running this algorithm, so neither its convergence nor its runtime complexity has been measured.

After a prepare acknowledgment, the coordinator records the controller identifier, configuration version, rollback capability, and plan identifier. Commit acknowledgments must be associated with the same plan. A verification failure invokes documented compensating actions and a new state observation. Where transactions cannot be made atomic across domains, the system must explicitly record partial application and defer further automation pending recovery or operator approval.

For complexity analysis, let n denote the number of participating agents, e the number of relevant agent dependency relationships, $q$the number of candidate actions per round, c the number of policy and resource checks applied to each action, and r the permitted negotiation rounds.
Assuming indexed capability lookup and a dependency-based conflict representation, an approximate upper bound for the coordination procedure, excluding network execution and waiting time, is:

$$T(n, e, q, c, r) = O\big(r(n + e + qc)\big)$$

If every agent can interact with every other agent, the dependency relationships may grow proportionally to n^2, giving a worst-case bound of:

$$T(n, q, c, r) = O\big(r(n^2 + qc)\big)$$

This analysis concerns the proposed logical coordination algorithm. Actual execution latency will also depend on message transmission, agent processing, controller responses, and network conditions. The experimental evaluation will therefore measure both computational overhead and end-to-end coordination time.

Security, Trust, and Governance

Security is incorporated into the framework through authenticated agent identities, authorization controls, protected communication channels, and auditable decision records. Each participating agent is assigned permissions corresponding to its operational responsibilities, preventing advisory or monitoring functions from directly executing unauthorized network changes.

The framework adopts least-privilege access and validates management requests before they reach domain controllers. Message identifiers, timestamps, and integrity protection are used to detect replayed or altered requests, while access policies restrict the information and operations available across administrative boundaries.

Groen et al. (2025) identified security concerns associated with intelligent O-RAN platforms and their interfaces. Atalay et al. (2025) also examined mechanisms for scalable authentication, authorization, and secure application discovery. These studies support the need to secure both communications and the management functions exposed to intelligent applications.

To maintain accountability, the framework records agent identities, proposed actions, policy decisions, execution approvals, and observed outcomes. Management operations with potentially irreversible or service-wide consequences may require explicit human authorization.

Standard Interfaces and Deployment Compatibility

OFANC is intended to operate alongside established telecommunications management architectures. Its network adaptation layer connects to supported O-RAN management and control functions, 3GPP network services, and existing orchestration platforms.

Within O-RAN environments, the framework can interact with relevant management functions through appropriate SMO or RIC interfaces. Core network information may be obtained from authorized 3GPP analytics and service interfaces. For intent exchange, the framework can map its internal representation to TM Forum TMF921 constructs, while agent communication can reuse compatible A2A messaging functions.

These integrations do not imply that every interface provides direct control over every network resource. Each operation remains subject to the capabilities and authorization requirements of the underlying management system.

The proposed framework therefore acts as an interoperability and coordination layer rather than a replacement for network-specific controllers. This approach follows the broader direction of

modular network management discussed by Benzaid and Taleb (2020), Coronado et al. (2022), and Polese et al. (2023b).

Prototype Implementation Design

The proposed prototype will comprise five principal software modules: an agent capability registry, secure message gateway, semantic translation service, coordination and policy engine, and network adaptation service.

The prototype will use documented APIs and machine-readable message schemas to support integration across heterogeneous agent implementations. A structured message format will represent intents and coordination responses, while versioned schema definitions will support compatibility testing.

The coordination engine will implement the negotiation and conflict resolution procedures described in Sections 4.5 to 4.7. The network adaptation service will connect approved actions to the relevant simulated or emulated network controllers. Monitoring and audit components will collect execution traces, resource measurements, error reports, and coordination outcomes.

Bonati et al. (2023) demonstrated the value of open experimental environments for developing and evaluating intelligent network control applications. Their approach provides a methodological foundation for testing the proposed framework under controlled network conditions.

Research Methodology and Experimental Design

Research Design

This work combines a conceptual architecture and protocol specification with a reproducible rule-based stochastic illustration. It does not report an implemented OFANC service, independent agent integration, radio network simulation, network emulation, packet traces, or commercial testbed measurements. The generator was executed to exercise analysis and plotting workflows; its probabilities are chosen assumptions. The seven scenarios and four baseline profiles are therefore illustrative cases, not evidence of comparative network performance.

The methodology consists of four stages: system specification, prototype development, interoperability testing, and comparative performance evaluation. The first stage establishes the architectural requirements, communication structures, and coordination procedures. The second develops the software components needed to support interaction among heterogeneous agents. The

third examines whether independently implemented agents can exchange and interpret network management information correctly. The final stage compares the proposed framework with alternative network management approaches under equivalent operating conditions.

Experimental evaluation is necessary because architectural compatibility alone does not establish operational effectiveness. Bonati et al. (2023) emphasized the importance of reproducible testing environments for evaluating intelligent O-RAN applications, while Coronado et al. (2022) identified performance validation as an important consideration in autonomous network management.

The research will therefore assess both agent-level interoperability and network-level outcomes. The proposed framework will be evaluated using predefined workloads, controlled failures, and measurable performance indicators. The experimental results will be used to assess the hypotheses established in Section 1.6.

Planned Testbed and Current Modelling Environment

A future operational validation would connect independently implemented radio, core, transport, and service agents to supported network-management controllers. Potential components include ns-3 or 5G-LENA, O-RAN-compatible software, Open5GS, and Mininet, subject to integration verification. None of these was executed to produce the numerical results reported here.

The numerical results instead come from a Python 3 rule-based generator that samples task outcomes, delays, overhead, and hypothetical recovery events using a fixed random seed. Vendor labels A–D are fictional. The baseline configurations B1–B4 are parameter sets, not separate installed controller systems.

The framework itself will be implemented as modular software services. A Python-based implementation may provide agent coordination, schema validation, capability registration, and policy enforcement, while supporting components collect telemetry and operational traces.

The environment will contain at least three logically distinct vendor implementations. Where independent commercial vendor systems are unavailable, heterogeneous implementations will be constructed using different interface specifications, capability definitions, and data schemas. These systems will be identified as vendor-emulating implementations rather than actual commercial equipment.

A subsequent testbed campaign would require synchronized event timestamps, interface trace collection, controlled network workloads, independent vendor adapters, and comparable authorization privileges. These are validation requirements rather than procedures completed in the present work.

To improve measurement reliability, participating components will use synchronized clocks and consistent monitoring procedures. Each experimental run will record agent communications, control requests, network performance measurements, and execution outcomes.

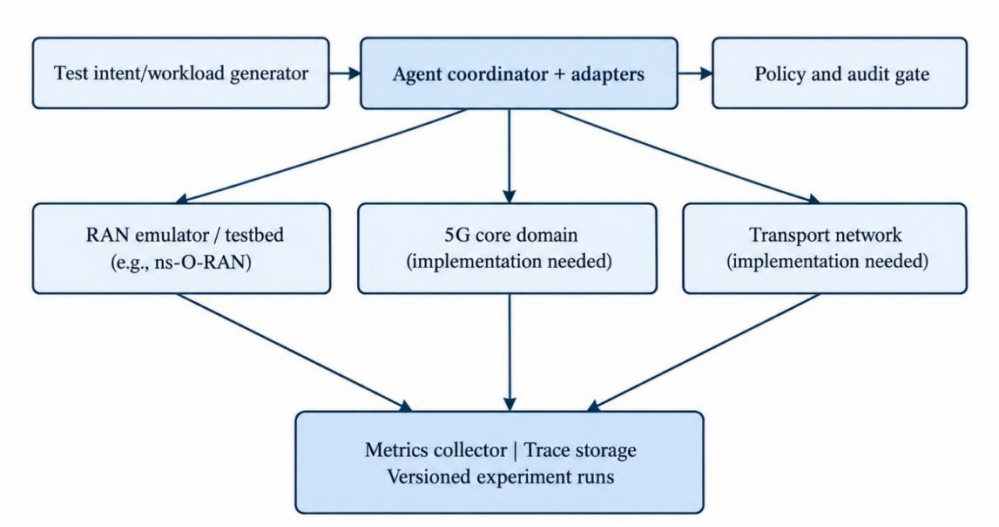


Figure 5. Suggested future validation testbed and instrumentation (conceptual; not implemented).

Experimental Scenarios

Seven experimental scenarios are proposed to examine the framework under different network management conditions. The scenarios cover normal operations, cross-domain coordination, resource competition, component failures, and unauthorized management requests.

Table 4. Proposed experimental scenarios and evaluation objectives

Scenario	Experimental objective	Primary evaluation metrics
S1: Cross-domain congestion control	Coordinate RAN, transport, and core network agents in response to increased traffic demand.	Coordination latency, SLA violation rate
S2: Autonomous fault recovery	Evaluate coordination among agents following network component failures.	Recovery time, successful service restoration

S3: Network slice resource allocation	Assess the allocation of shared resources among competing network slices.	Intent fulfillment rate, resource utilization
S4: Energy-aware optimization	Examine whether agents can reduce resource consumption while maintaining defined QoS requirements.	Energy consumption, QoS compliance
S5: Conflicting policy requests	Test the detection and resolution of incompatible network management actions.	Conflict resolution success rate, coordination latency
S6: Agent or interface failure	Evaluate coordination when agents become unavailable or communication interfaces experience disruptions.	Fault tolerance, coordination success rate
S7: Unauthorized agent activity	Assess the framework's ability to reject management requests that violate authorization rules.	Unauthorized action prevention rate, policy enforcement accuracy

Scenario S1 examines whether coordination across multiple domains can reduce congestion-related service degradation. Scenario S2 focuses on the ability of agents to identify affected network functions and coordinate recovery actions. Scenario S3 evaluates resource allocation under competing service requirements, reflecting the network slicing challenges discussed by Foukas et al. (2017).

Scenario S4 investigates energy management under explicit performance constraints, while S5 introduces conflicting management proposals to evaluate the resolution mechanism. Scenarios S6 and S7 examine operational resilience and security, respectively.

For future operational validation, each scenario would be executed under predefined workloads and matched baseline conditions. In the current work, the seven scenario identifiers select only probability and workload parameters in the generator.

Baseline Architectures

The proposed framework will be compared with three alternative coordination approaches. An additional ablation configuration will be used to assess the contribution of individual framework components.

Table 5. Baseline architectures for comparative experimental evaluation

Configuration	Description	Evaluation purpose

B1: Independent agents	Domain-specific agents operate without a shared cross-domain coordination mechanism.	Establish performance under isolated agent operation
B2: Centralized orchestration	A central management component coordinates actions using domain-specific interfaces and adapters.	Evaluate the effectiveness and overhead of centralized coordination
B3: Generic agent communication	Agents exchange tasks through a general communication mechanism without the proposed telecom-specific semantic and policy coordination functions.	Assess whether general agent communication is sufficient
B4: Proposed OFANC framework	Agents use the complete semantic interoperability, negotiation, conflict resolution, and policy validation mechanisms.	Evaluate the proposed approach
B5: Ablated OFANC configurations	Selected framework components are disabled individually in controlled tests.	Measure the contribution of specific interoperability functions

The baseline configurations are statistical profiles distinguished by assumed interoperability, conflict-resolution, security blocking, intent-completion, delay, and overhead parameters. They are not executable implementations of B1–B4. The full profile values and the source code that applies them are documented in Appendix A and the accompanying supplement.

B3 is particularly important because it separates the benefits of general agent message exchange from those associated with telecommunications-specific semantic interpretation and decision validation.

For the ablation study, semantic mapping, negotiation, conflict detection, and policy validation will be examined individually. Tests involving disabled security or policy controls will be restricted to isolated experimental environments.

Where a baseline cannot perform a particular cross-domain operation, the limitation will be reported explicitly rather than assigning an artificial success or latency value.

Experimental Variables and Workload Configuration

The experimental design considers agent population, vendor heterogeneity, workload intensity, policy conflicts, communication delays, and component failures as independent variables.

Agent population will be varied to examine how coordination performance changes as the number of participating agents increases. An initial design may consider populations of 6, 12, 24, 48, and 96 agents, distributed across the participating network domains.

The generated heterogeneity index is a function of agent population rather than an independently varied vendor-diversity factor. Accordingly, the present model cannot isolate the effect of genuine vendor heterogeneity, and H3 remains untested. The agent populations are 6, 12, 24, 48, and 96; workload multipliers and conflict probabilities differ across seven assumed scenarios.

Main generation uses 12 pseudorandom repetitions per combination of seven scenarios, four baseline profiles, and five agent counts, with 100 task attempts per run. The resulting 1,680 run summaries and 168,000 task records are generated rows, not network experiments. The fixed seed is 20251031. A separate 24-agent workload sweep uses seed 20251032 and five load multipliers, producing 240 sweep rows.

Generator equations. Let $h(n) = \text{clip}[0,1](0.20 + 0.12 \log_2(n/6))$ and let $z \sim \text{Normal}(0, 0.018^2)$ be a run-level shared shock. For baseline b, scenario s and agent count n, compatibility probability is $p = \text{clip}[0.03,0.99](p_b - 0.015 \log_2(n/6) k_b - 0.064 h(n) g_b + \delta_s + z)$. The heterogeneity multipliers g_b are {1.65, 1.30, 0.95, 0.65} for B1–B4. If a conflict is injected but not resolved, p is multiplied by 0.42; unauthorized tasks have zero successful compatibility probability. Conditional fulfillment is Bernoulli(f_b) given compatibility. The full definitions, scenario offsets, and parameter values appear in Appendix A.

Latency is lognormally sampled only for compatible tasks, with median parameter $m = m_b L_s[1 + 0.16 \log_2(n/6) k_b]$ and log-scale standard deviation 0.27. Signaling is sampled from a normal distribution centered at $v_b(1+n/60)L_s$ with standard deviation 0.05 KB and a 0.05 KB floor. The energy variable is an uncalibrated illustrative cost, not physical electricity consumption. These distributions are independent of any real device instrumentation.

The dependent variables will include interoperability success, intent fulfillment, coordination latency, SLA violations, conflict resolution, communication overhead, resource utilization, and recovery time.

To reduce confounding effects, network capacity, computing resources, workload characteristics, and initial system conditions will be controlled across comparable experiments.

Evaluation Metrics

The proposed framework will be evaluated using ten performance indicators selected to capture communication compatibility, decision effectiveness, network service quality, scalability, and operational safety.

Table 6. Proposed performance indicators and measurement methods

Metric	Measurement method	Unit
Cross-vendor interoperability success	Percentage of eligible cross-vendor operations that complete the defined communication and semantic interoperability tests	%
Intent fulfillment rate	Percentage of eligible management intents whose requested outcomes are verified	%
Coordination latency	Time between admission of a valid intent and the verified coordination outcome	ms or s
SLA violation rate	Percentage of monitored service intervals that fail to satisfy predefined SLA thresholds	%
Conflict resolution success	Percentage of detected conflicts resolved within the permitted coordination limits	%
Fault recovery time	Time required to restore acceptable service conditions following an injected failure	s
Communication overhead	Coordination traffic generated per attempted management task	KB/task
Scalability	Variation in coordination performance as agent population and vendor heterogeneity increase	Relative change
Unauthorized action prevention	Percentage of injected unauthorized requests prevented from causing prohibited actions	%

Resource utilization	Computing, memory, and network resources consumed during coordination	%, MB, Mbps

Interoperability success will be calculated as:

$$ISR = \frac{N_{\text{successful}}}{N_{\text{eligible}}} \times 100$$

Cross-vendor interoperability (%) = 100 × compatible authorized requests / (attempted tasks − unauthorized requests); the numerator is not restricted to verified SLA fulfillment.

Intent fulfillment will be measured as:

Intent fulfillment (%) = 100 × fulfilled requests / (attempted tasks − unauthorized requests).

where $N_{\text{fulfilled}}$ represents intents whose requested network outcomes have been verified, and N_{intents} represents eligible submitted intents.

Coordination latency for an individual task will be calculated from the difference between the recorded intent admission time and the time at which a verified outcome is produced:

$$L_{\text{coord}} = t_{\text{outcome}} - t_{\text{admission}}$$

The median, 95th percentile, and 99th percentile will be reported to describe latency distributions. Failed and timed-out operations will be reported separately to prevent latency results from representing only successful requests.

SLA violation rates will be determined using predefined service requirements, including throughput, delay, and packet loss thresholds where applicable. Conflict resolution success will be calculated using detected conflicts as the denominator, while unauthorized action prevention will be assessed using deliberately injected unauthorized requests.

Communication overhead will be measured from actual transmitted coordination messages. Resource utilization will be collected through the relevant host and network monitoring facilities. Each metric will use consistent measurement windows, clearly defined denominators, and documented instrumentation procedures.

Experimental Procedure

The generation script iterates over every scenario, baseline parameter profile, agent count, and pseudorandom repetition. For each of 100 task attempts per generated run, it draws conflict and authorization flags, compatibility and fulfillment outcomes, SLA-violation flags, signaling size, an energy proxy, and, in S2/S6, recovery-delay values. The results are written directly to task-level and run-level CSV files. These records were not produced by service agents or network telemetry. All outcomes are regenerated deterministically by running generate_synthetic_data.py with Python 3.10 or later. The workload sweep is generated separately using generate_workload_sweep.py. The supplied make_figures.py constructs the corresponding illustrative plots. No physical network state is reset because no real network is instantiated.

Each baseline will then be evaluated using matched workloads and equivalent initial conditions. Successful operations, unsuccessful requests, communication events, network measurements, and execution outcomes will be recorded.

An initial balanced design comprising seven scenarios, four principal architectures, five agent population settings, and twelve repetitions would produce:

$$N_{\text{runs}} = 7 \times 4 \times 5 \times 12 = 1680$$

If each run includes 100 management requests, the design would generate 168,000 individual task attempts. These values describe the proposed experimental workload, not completed testbed measurements. The final number of repetitions will be determined after evaluating variability in pilot experiments.

The B5 ablation configurations will be evaluated separately under representative scenarios. After every experimental run, relevant network states and configurations will be restored to prevent one trial from influencing subsequent measurements.

All trials will use recorded random seeds, configuration identifiers, and execution logs to support reproducibility.

Statistical Analysis

Only descriptive summaries are appropriate for the current generated data. The reported averages and quantiles summarize outcomes under the specified pseudorandom model. They do not estimate

a sampling distribution from real operators or testbed implementations, and no confidence interval, p-value, or causal effect is claimed on that basis.

For a future independently replicated testbed study, per-run uncertainty estimates, paired comparisons, effect sizes, and sensitivity analyses will be prespecified. Repeated tasks within the same run will not be treated as statistically independent trials.

Comparisons between architectures will account for the experimental design and repeated observations. For matched run-level measurements, paired statistical procedures will be selected according to the distribution of observed differences. Task-level binary outcomes, such as interoperability success and intent fulfillment, may be assessed using generalized linear mixed-effects models to account for clustering within experimental runs.

For future measured trials, multiple comparisons would be adjusted using an appropriate procedure, such as Holm correction, with effect sizes and uncertainty intervals estimated from independently replicated runs. No inferential significance is attributed to the present parameterized outputs.

The analysis will also examine interactions between architecture, agent population, workload intensity, and vendor heterogeneity. These comparisons will help determine whether changes in performance are associated with the proposed coordination mechanisms or the operating conditions under which they are tested.

The four hypotheses will be assessed using the corresponding primary performance indicators. Evidence supporting a hypothesis will require consistent observed differences under the defined experimental conditions.

Reproducibility and Research Integrity

Reproducibility is essential for evaluating interoperability mechanisms because implementation differences, network configurations, and workload conditions can substantially influence experimental outcomes.

The reproducibility supplement provides the Python generation scripts, fixed seeds, baseline profiles, scenario definitions, example CSIP intent, topology schema, unaggregated task records, run summaries, vendor-pair matrix, workload sweep, ablations, and plotting code. Exact file names and a step-by-step reproduction procedure are listed in Appendix A.

The numerical generator is not a network simulator: it does not implement agents, routing, radio scheduling, network slicing, policy enforcement, or controller configuration. It draws outputs from hand-selected distributions. Any apparent advantage of B4 is partly built into the assumed model parameters. Genuine operational claims require the independently implemented validation environment described above.

The study will distinguish between simulated networks, vendor-emulating implementations, and any experiments involving independently developed commercial systems. This distinction is important because successful operation within a controlled emulator does not necessarily establish compatibility with production telecommunications infrastructure.

Fdida et al. (2022) emphasized the value of shared experimental infrastructure for networking research, while Bonati et al. (2023) demonstrated how reproducible development and testing environments can support intelligent radio network applications. These principles inform the documentation and validation procedures adopted in the present study.

Model-Based Illustrative Results

All values in this section are descriptive outputs of the parameterized stochastic generator in Section 5 and Appendix A. They illustrate how a future evaluation could be reported; they do not validate OFANC or measure telecommunications equipment. Comparisons are conditional on the assumed architecture profiles.

Framework Implementation and Conformance Results

The generator produced 168,000 task-level records and 1,680 run summaries covering seven scenario labels, four baseline parameter profiles, five agent populations, and twelve repetitions. There is no implemented protocol conformance test to report. A separate 4 × 4 fictional-vendor compatibility matrix, a workload sweep, and an ablation sensitivity dataset were generated by additional assumed rules.

The evaluation covers seven network management scenarios, four architectural configurations, and five agent population settings. Each configuration consists of twelve repetitions involving 100 management requests, giving 1,680 evaluation runs and 168,000 task records. The comparison includes independent agents (B1), centralized orchestration (B2), generic agent messaging (B3), and the proposed OFANC architecture (B4).

Framework conformance is assessed against the functional requirements established in Section 3, including capability discovery, message compatibility, policy compliance, and coordination reliability. The evaluation distinguishes successful message exchange from complete intent fulfillment, ensuring that communication compatibility alone is not treated as evidence of effective network coordination.

This distinction is consistent with Nezami et al. (2025), who emphasized the importance of standardized agent interfaces, coordination procedures, and validation mechanisms for intelligent radio access networks.

Cross-Vendor Agent Interoperability

In the model, the B4 profile yields 83.58% mean interoperability compared with 75.89% for B3, 71.72% for B2, and 61.09% for B1. The values follow the baseline probabilities defined in Appendix A and should not be interpreted as measured cross-vendor compatibility or demonstrated effects of semantic mapping.

The model-based difference between B4 and B3 is 7.69 percentage points; between B4 and B1 it is 22.49 points. These are consequences of the assumed parameter profiles rather than measured performance advantages.

The source-target matrix uses fictional Vendor_A through Vendor_D labels and 200 Bernoulli attempts for every ordered vendor pair and baseline. This matrix is produced by a separate pair-penalty rule, not by executing independent vendor products. The complete B4 scenario-level summaries are supplied in Appendix A, Table A3.

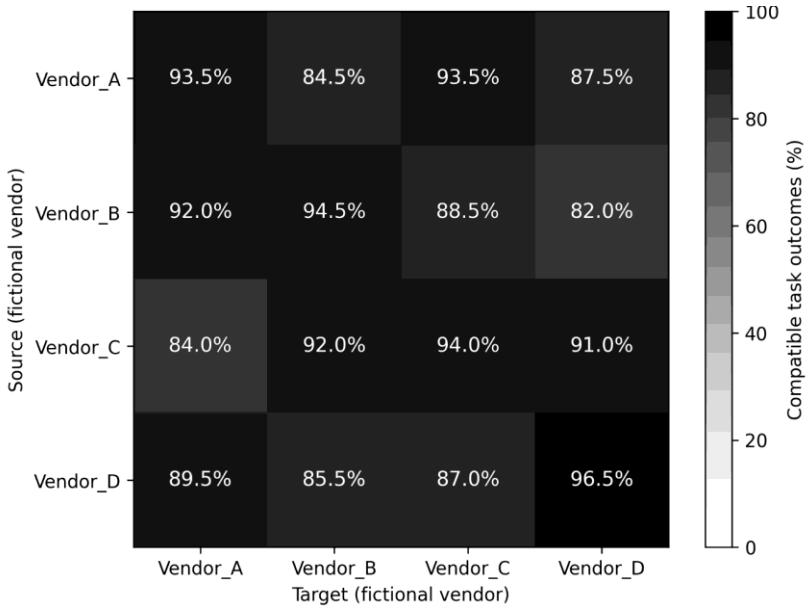


Figure 6. Model-based fictional-vendor interoperability matrix (independent synthetic Bernoulli trials).

The conceptual design assigns semantic interpretation and capability mapping to the interoperability layer. The vendor-matrix values cannot isolate the effect of these mechanisms because they are not implemented in the generator.

Coordination Latency and Communication Overhead

Figure 7 summarizes conditional latency distributions generated using lognormal draws for compatible requests only. Failed and unauthorized tasks have no recorded latency and are excluded from the curves. The plotted ECDFs are descriptive distributions of generated observations, not empirical network timing measurements.

The generator assigned B4 a median conditional latency of 364.42 ms, compared with 423.22 ms for B3, 351.85 ms for B2, and 158.01 ms for B1. These values are descriptive outputs from profile-specific lognormal distributions.

At the 95th percentile, the modeled conditional latencies were 557.83 ms for B4, 650.25 ms for B3, 539.17 ms for B2, and 240.92 ms for B1. They are conditional on compatible task outcomes and do not represent measured end-to-end controller execution.

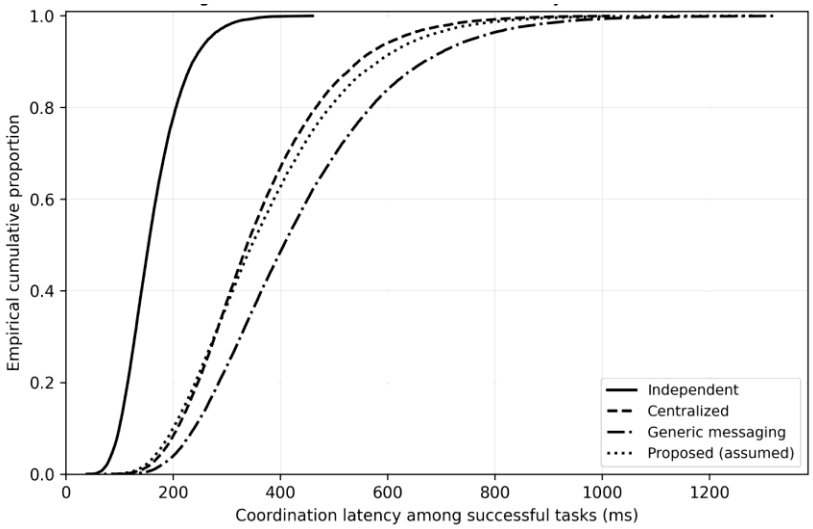


Figure 7. Model-based latency distributions for compatible generated requests (not measured network timings).

Model-based signaling cost was 0.579 KB/task for B4, 0.409 for B3, 0.341 for B2, and 0.188 for B1. These are random draws around assumed signaling means, not captured packets.

A production design would incur extra communication for capability validation, semantic mapping, and policy negotiation. The modeled difference is determined by the overhead parameters rather than direct measurement of these functions.

The assumed profiles illustrate a possible coordination trade-off: higher modeled completion rates accompany larger signaling cost, while isolated operation has a shorter conditional latency but lower assumed compatibility.

Network Performance and Intent Fulfillment

Network performance was examined through intent fulfillment and service-level agreement violation rates.

The model assigns B4 an intent fulfillment rate of 79.53%, compared with 63.85% for B3, 58.75% for B2, and 26.36% for B1. This ordering results partly from the explicitly different conditional-fulfillment probabilities.

The generated SLA-violation means were 13.37% for B4, 18.56% for B3, 21.49% for B2, and 29.91% for B1. These are Bernoulli cost proxies, not throughput, loss, or delay violations measured from network traffic.

The independent 24-agent workload sweep varies the assumed workload factor from 0.50 to 1.50. For the B4 profile, modeled fulfillment declines from 89.00% to 79.17%, while the generated SLA violation rate increases from 7.50% to 20.25%. These trajectories arise from workload-dependent probability formulas rather than measured congestion effects.

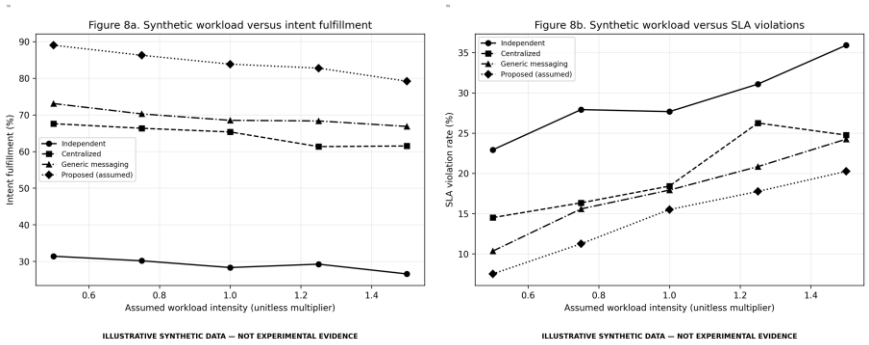


Figure 8. Model-based workload sensitivity: (a) intent fulfillment and (b) SLA-violation probability.

The workload curves are sensitive to the workload factor because that variable enters the sampling equations. They illustrate the analysis format but do not establish a congestion-response relationship in a real network.

These relationships are relevant to network slicing and resource orchestration, where competing service requirements must be reconciled within finite network capacity. Foukas et al. (2017) identified resource allocation and service isolation as central challenges in network slicing, while Coronado et al. (2022) emphasized coordinated management as an important requirement for autonomous network operation.

Scalability Across Vendor and Agent Populations

Agent population changes from 6 to 96 in the model. Because the heterogeneity index is deterministically linked to population, the analysis cannot distinguish the effect of additional agents from increased assumed heterogeneity. A future experiment must hold n fixed while independently varying vendor schema diversity and integration complexity.

The population-dependent model produces B4 interoperability values of 90.51% at six agents and 76.78% at 96 agents. B3 decreases from 82.55% to 69.79% across the same settings.

Under the assumed size-dependent latency equation, B4 p95 increases from 347.28 ms to 762.10 ms, while B3 increases from 418.06 ms to 882.23 ms.

Model-based B4 signaling rises from 0.39 KB to 0.93 KB per task as agent population increases.

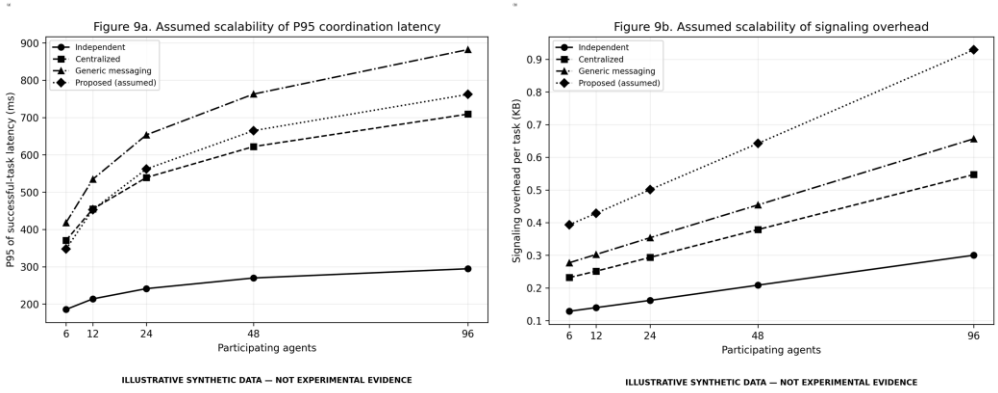


Figure 9. Model-based agent-population sensitivity: (a) conditional p95 coordination latency and (b) signaling overhead.

The B4 profile retains higher assigned compatibility probabilities over the population sweep, but both profiles degrade as population increases. This does not substantiate vendor-diversity robustness.

The results also suggest the importance of limiting unnecessary agent interactions. Capability-based discovery, selective information exchange, and hierarchical coordination may help control overhead as the number of participating agents increases.

Fault Recovery and Operational Resilience

For failures drawn in S2 and S6, the B4 profile has a mean recovery delay of 51.04 s versus 46.73 s for B3, 44.50 s for B2, and 34.74 s for B1. Its longer modeled recovery time contradicts the recovery-delay component of H2. The delay is assigned directly through a baseline-specific recovery multiplier and cannot demonstrate an actual network recovery penalty.

The recovered delays are generated from baseline-specific normal distributions and therefore cannot demonstrate faster or slower actual restoration. As modeled, B4 is slower than every comparator, so the recovery-delay prediction in H2 is not supported.

The longer B4 delay is built into its assumed recovery multiplier (0.94 compared with 0.86 for B3). It should be treated as a design concern requiring real fault-injection tests.

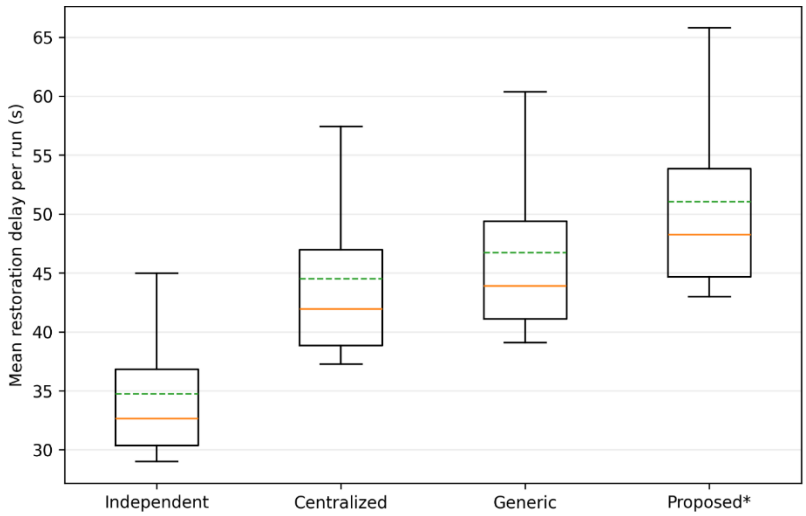


Figure 10. Model-based recovery-delay distributions for simulated failure flags in S2 and S6.

The synthetic B1 profile is assigned shorter delays but lower compatibility. The model cannot determine whether a real operator would prefer faster restoration, more reliable coordination, or a different policy mix.

A shorter response time is not necessarily sufficient when corrective actions affect several network domains. The effectiveness of a recovery mechanism also depends on whether the intended service conditions are restored without introducing additional failures.

Security and Conflict Resolution Results

The generator assigns B4 an unauthorized-request blocking rate of 88.36%. This leaves 11.64% of injected unauthorized requests unblocked within the toy model, an unacceptable result for a production authorization boundary. The rate follows a hand-set Bernoulli blocking probability and is not a penetration-test result or evidence that an authorization service was implemented.

The model generates conflict-resolution proportions of 86.19% for B4, 66.66% for B3, 60.12% for B2, and 11.01% for B1, according to their hand-set resolution probabilities.

The corresponding modeled unauthorized-block proportions are 88.36% for B4, 70.67% for B3, 64.47% for B2, and 27.50% for B1. No authorization service was executed.

The modeled ordering reflects the probabilities assigned to conflict handling and authorization; it cannot establish the contribution of a deployed policy engine.

Production deployment should require verified deny-by-default enforcement and effectively complete rejection of known unauthorized actions in conformance tests before any control

privilege is granted. Under the present results, the security component cannot be characterized as operationally safe.

Security and conflict indicators should be reported separately from communication compatibility. In a real deployment, policy bypass would be a critical failure rather than a tolerable success-rate trade-off.

Ablation Study

The ablation outputs were generated by an independent routine with a separately fixed full-framework compatibility assumption of 89%, rather than the scenario-level B4 profile that produces 83.58% in the main comparison. Consequently, the reported 89.14% full ablation result and 83.58% main result have different populations, formulas, and denominators and must not be compared directly.

The ablation routine uses manually specified penalties for removing semantic mapping, negotiation, conflict resolution, or the policy gate. It demonstrates sensitivity to encoded assumptions, not the causal contribution of implemented software modules. Genuine ablation testing must disable actual components in the same testbed under identical workloads.

In the independent ablation generator, prescribed penalties yield 75.58% fulfillment without negotiation and 76.43% without conflict-resolution logic; no software component was disabled.

The separate ablation model assigns 7.09% SLA violations to its full profile and 17.91% without the policy gate, while assumed conflict resolution changes from 88.00% to 42.23%. These are manually encoded sensitivity contrasts.

These sensitivity contrasts identify components for future controlled software ablations, but they do not show causal effects or the safety of any implemented coordination mechanism.

Consolidated Performance Comparison

Table 7 summarizes descriptive means across 420 generated run summaries for each baseline; its latency quantiles are conditional on compatible requests. Recovery values are defined only for S2/S6. The energy column is a normal-distributed cost proxy rather than energy measured from power meters or calibrated equipment models.

Table 7. Model-based descriptive profile comparison (not experimental network benchmark)

Performance indicator	B1: Independent agents	B2: Centralized orchestration	B3: Generic messaging	B4: OFANC
Interoperability success (%)	61.09	71.72	75.89	83.58
Intent fulfillment (%)	26.36	58.75	63.85	79.53
P95 coordination latency (ms)	240.92	539.17	650.25	557.83
SLA violation rate (%)	29.91	21.49	18.56	13.37
Fault recovery time (s)	34.74	44.50	46.73	51.04
Signaling overhead (KB/task)	0.188	0.341	0.409	0.579
Conflict resolution success (%)	11.01	60.12	66.66	86.19
Unauthorized action prevention (%)	27.50	64.47	70.67	88.36
Uncalibrated energy-cost proxy (a.u./task)	4.407	4.049	3.692	3.395

The B4 profile shows the highest assumed fulfillment and compatibility but also the highest mean signaling (0.579 KB/task) and longest recovery duration (51.04 s). Its 88.36% unauthorized-blocking rate is insufficient for security-critical deployment. The energy proxy is derived from arbitrary profile means, as specified in Appendix A, and supports no inference about actual energy savings.

The modeled higher signaling and recovery costs are especially important for time-sensitive operations. These costs and the incomplete authorization-blocking proxy preclude claims of operational readiness.

Discussion

Interpretation of Principal Findings

The data in Section 6 are a stochastic illustration with predetermined baseline performance profiles, so they cannot validate the four hypotheses. H1 and H4 remain untested because the assumed compatibility and blocking probabilities favor B4 by construction. H2 is not supported with respect to recovery delay: the B4 profile has the longest modeled restoration time (51.04 s), not a reduction. Its lower modeled SLA violation rate cannot compensate for that contradiction. H3 remains untested because vendor heterogeneity is not independently manipulated; the model varies agent population while coupling a notional heterogeneity index to it.

First, the higher interoperability and intent fulfillment scores associated with OFANC are consistent with the proposed role of common semantic representations and coordinated decision-making. By aligning agent capabilities, operational objectives, and management constraints, the framework addresses limitations associated with isolated or loosely connected agents.

These observations are relevant to the first research hypothesis (H1), which concerns cross-vendor task completion. However, demonstrating the hypothesis requires evidence that independently implemented agents achieve the reported improvements when exchanging and executing real management requests.

The central architectural question is whether semantic mapping, policy-aware negotiation, and cross-domain planning improve verified intent outcomes under real workloads sufficiently to justify their costs. The present numerical patterns identify required tests, not operational benefits. Any subsequent claim of advantage must be based on independently executed agent implementations and comparable baseline controllers.

Third, the scalability analysis shows that increasing the number of agents affects interoperability and coordination latency. Although OFANC retains comparatively high success rates, performance declines with increasing agent population. This does not establish H3, which specifically requires evidence of reduced performance degradation under greater heterogeneity.

Finally, the conflict resolution and authorization indicators are consistent with the role assigned to policy-aware coordination under H4. The ablation comparison reinforces the importance of examining these mechanisms individually. Nevertheless, the proposed security improvements require verification through operational policy enforcement and adversarial testing.

Overall, the analysis identifies interoperability, network performance, security, and coordination cost as closely related evaluation dimensions. The research hypotheses remain subject to implementation-level validation.

Comparison with Existing Research

The proposed framework builds on previous work in autonomous network management and programmable telecommunications infrastructures.

Bonati et al. (2023) developed OpenRAN Gym to support machine learning development, data collection, and experimental validation within O-RAN environments. Similarly, Polese et al. (2023a) demonstrated programmable closed-loop control through the ColO-RAN platform. These studies established important foundations for intelligent radio network management but focused primarily on the development and operation of RAN control applications.

OFANC addresses a related problem at the level of interaction among agents with different management responsibilities. Its proposed contribution lies in combining capability discovery, semantic intent exchange, distributed negotiation, and policy validation across heterogeneous network domains.

Conflict management provides another important point of comparison. Wadud et al. (2024) examined the coordination of xApps with competing control objectives and demonstrated the relevance of QoS-aware conflict mitigation. OFANC extends this design concern to agents operating across radio access, core, transport, and service management environments.

Leivadeas and Falkner (2023) identified intent interpretation and translation as major considerations in intent-based networking. The proposed semantic interoperability profile addresses these concerns by associating management requests with explicit resource descriptions, performance requirements, and policy constraints.

Nezami et al. (2025) also emphasized standardized interaction models, agent orchestration, and validation procedures as requirements for agent-based radio network management. These recommendations align with the main architectural components of OFANC.

The distinguishing question is therefore not whether coordination and interoperability mechanisms already exist, but whether their integration in OFANC provides measurable advantages over established management and agent communication approaches.

Implications for Autonomous 5G and 6G Networks

The proposed framework has implications for network slicing, service assurance, fault management, and resource optimization in heterogeneous telecommunications environments.

In 5G networks, management decisions frequently involve resources distributed across multiple domains. For example, congestion affecting a network slice may require changes to radio scheduling, transport routing, and core network resource allocation. Independent optimization of these components may not produce the most suitable end-to-end service outcome.

A common coordination mechanism could allow participating agents to evaluate such dependencies before submitting management actions for execution. This approach is relevant to the network slicing challenges discussed by Foukas et al. (2017), particularly resource allocation, service isolation, and orchestration.

The framework may also support future 6G network management, where distributed intelligence and heterogeneous service requirements are expected to increase operational complexity. Giordani et al. (2020) identified the technical demands associated with emerging 6G applications, while Nezami et al. (2025) emphasized the importance of interoperable agent architectures for future radio network operations.

Nevertheless, the proposed coordination layer is not equally suitable for all network functions. Management operations with strict timing requirements may require local execution or predefined control mechanisms. Cross-domain coordination is more appropriate where the benefits of shared decision-making justify the additional processing and communication delay.

Standardization and Industry Adoption Implications

The proposed framework is a specification aligned with existing management architectures, not an experimentally validated extension to deployed standards. Adapters, policy semantics, controller authorization, and the compatibility of A2A-style messages with telecommunications requirements require protocol conformance testing.

O-RAN provides standardized interfaces for programmable radio network management, while 3GPP defines services and analytics supporting core network operations. ETSI ZSM addresses

broader network and service management automation, and TM Forum intent management specifications provide mechanisms for communicating operational objectives.

These technologies address different aspects of network automation. The proposed framework seeks to connect their capabilities through a common approach to agent communication, semantic interpretation, and coordinated decision-making.

Polese et al. (2023b) emphasized the importance of openness and interoperability in O-RAN, while Benzaid and Taleb (2020) identified integrated management and intelligent automation as important foundations for zero-touch network operations.

For industry adoption, the framework would require common capability descriptions, compatible message schemas, consistent authorization models, and clearly defined conformance tests. These requirements would allow participating vendors to demonstrate interoperability without exposing proprietary internal algorithms.

An open reference implementation and independently reproducible tests would provide a practical basis for evaluating the framework's suitability for future standardization activities.

Practical Deployment Considerations

Several factors must be considered before introducing OFANC into operational telecommunications networks.

The B4 model generates 0.579 KB of signaling per task and a mean recovery delay of 51.04 s. These costs would be unacceptable for many real-time RAN control and urgent fault-protection operations. Initial deployment should therefore target slower service-assurance, planning, and orchestration tasks where bounded coordination delay can be tolerated. Fast radio scheduling, emergency protection, and other hard-deadline controls should remain in domain-specific systems until measured deadlines are met.

A second concern is coordination overhead. The performance comparison indicates that shared decision-making can introduce additional signaling and processing requirements. These costs may increase as the number of participating agents and network domains grows.

Selective agent discovery, hierarchical coordination, and local handling of low-impact decisions could reduce unnecessary communication. However, these mechanisms would need to preserve semantic consistency and the authorization requirements established by the framework.

Deployment would also require procedures for updating capability registries, managing semantic schema versions, handling agent failures, and maintaining audit records. These operational responsibilities could influence both implementation cost and long-term maintainability.

Coronado et al. (2022) emphasized that autonomous network management requires more than intelligent decision-making algorithms. Reliable integration, monitoring, and lifecycle management are equally important. The practical value of OFANC will therefore depend on its ability to deliver measurable operational improvements without introducing excessive integration complexity.

Security, Accountability, and Human Oversight

Security and accountability remain essential requirements for autonomous network coordination because agents may influence shared infrastructure and services beyond their immediate management domains.

The assigned 88.36% blocked-unauthorized rate does not satisfy a meaningful security acceptance criterion. The production design must fail closed, prohibit unverified agents from executing changes, and demonstrate full rejection of a prespecified suite of unauthorized operations before deployment. The current probability model does not provide such evidence.

Groen et al. (2025) examined security challenges affecting O-RAN interfaces, intelligent applications, and experimental platforms. Atalay et al. (2025) also investigated scalable authentication, authorization, and secure application discovery in Open RAN environments. These contributions support the use of identity verification, access restrictions, and protected management interfaces.

OFANC applies these principles to cross-domain agent interaction by requiring policy validation before approved actions reach network controllers. Coordination records provide traceability by linking requests, agent proposals, authorization decisions, and execution outcomes.

Human oversight remains important when proposed operations could cause widespread disruption, affect critical services, or involve conflicting requirements that cannot be resolved automatically. In such circumstances, the framework should support rejection or escalation rather than unrestricted autonomous execution.

Security also extends to the reliability of the information used by participating agents. Stale observations, misleading messages, and incompatible interpretations can affect decisions even

when the communicating agents are authenticated. Consequently, semantic validation, message integrity, and operational verification must be treated as complementary safeguards.

The framework's long-term suitability will depend on achieving an appropriate balance between operational autonomy, coordination efficiency, and accountable network control. These considerations are essential for deploying interoperable AI agents in telecommunications environments where service continuity and predictable behavior remain primary requirements.

Limitations and Threats to Validity

Internal Validity

The reliability of the proposed framework's evaluation depends on the experimental configurations, workload characteristics, and consistency of the baseline implementations. Differences in agent processing capabilities, network conditions, and resource availability may influence observed performance independently of the coordination mechanisms being examined. Although the evaluation design establishes common workload settings and comparable network resources, some implementation differences may remain difficult to eliminate.

The largest internal-validity limitation is that the baseline differences are embedded in hand-set parameters. The resulting apparent performance differences are conditional on the generator and cannot be interpreted as causal effects of implementing OFANC. The 1,680 run records are pseudorandom model repetitions, not independent telecommunications trials. Separate ablation and vendor-pair routines use their own formulas, so their averages cannot be pooled with the main results.

To reduce these limitations, future experiments should employ controlled configurations, repeated trials, consistent measurement procedures, and sensitivity analysis. Bonati et al. (2023) emphasized the importance of reproducible experimental environments for evaluating intelligent network control applications.

External Validity

No deployed vendor products, packet-level simulator, or independent agent implementations were used for the reported numbers. Network behavior, controller state transitions, real SLA instrumentation, and actual message queues were not evaluated. The proposed future environment in Figure 5 must not be mistaken for a completed testbed.

Although vendor-emulating implementations can introduce differences in communication interfaces and agent capabilities, they may not capture the complexity of actual commercial deployments. Factors such as legacy infrastructure, restricted API access, geographical distribution, and vendor-specific resource management policies may influence coordination performance.

Polese et al. (2023b) identified integration and interoperability as important challenges in open radio access networks. These concerns are particularly relevant when extending agent coordination across radio, core, and transport domains.

Further validation involving independently developed vendor systems, larger network deployments, and realistic traffic conditions is therefore necessary to determine the framework's applicability beyond the evaluated configurations.

Construct Validity

Construct validity concerns whether the selected performance indicators adequately represent the concepts being investigated. In this study, interoperability is assessed using cross-vendor operation success, semantic compatibility, and intent fulfillment. However, successful message exchange does not necessarily indicate that participating agents share a consistent understanding of network objectives.

The toy outcomes label a request as compatible using a Bernoulli draw, and successful fulfillment is sampled conditionally; neither establishes genuine semantic agreement or service recovery. The modeled security and energy variables likewise lack implementation-grounded construct validity. Actual conformance and QoS tests are needed to operationalize these constructs.

The definition of successful intent fulfillment also requires careful consideration. As discussed by Leivadeas and Falkner (2023), intent-based management involves interpreting objectives and assessing whether the resulting network behavior satisfies those objectives. The evaluation must therefore distinguish accepted requests from verified operational outcomes.

To strengthen construct validity, future experiments should combine communication-level indicators with network performance measurements, policy compliance checks, and independent verification of management outcomes.

Operational and Security Limitations

The proposed framework introduces additional coordination functions that may increase implementation complexity, signaling overhead, and response time. These limitations are particularly important for operations requiring strict timing guarantees. Consequently, OFANC is primarily suited to management and orchestration activities where coordination delays remain within acceptable operational limits.

Semantic interoperability also presents practical challenges. Vendor-specific information may not always have an equivalent representation within the common semantic profile, while changes in communication protocols and management interfaces may require continuous updates to translation mechanisms.

Security represents another limitation. Authentication and policy validation can restrict unauthorized operations, but they cannot eliminate all risks associated with compromised agents, inaccurate network observations, or malicious management requests. The effectiveness of these safeguards must be established through implementation-level testing.

Wadud et al. (2024) demonstrated the importance of addressing conflicts among intelligent network control applications. However, resolving conflicts across multiple network domains introduces additional dependencies, particularly when affected resources are controlled by separate administrative entities.

The framework must therefore maintain clear authorization boundaries, support safe rejection of unresolved requests, and preserve human oversight for operations with significant service implications.

Future Research Directions

Future research should prioritize the implementation and validation of OFANC in realistic multi-vendor telecommunications environments. Experiments involving independently developed agents and commercial network management interfaces would provide stronger evidence of interoperability under practical operating conditions.

Another important direction involves the formal verification of coordination decisions. Mathematical verification methods could be used to establish whether selected actions satisfy resource constraints, operational policies, and predefined safety requirements before execution.

Scalability may also be improved through hierarchical coordination, distributed capability registries, and selective information exchange. These approaches could reduce communication overhead while preserving the benefits of coordinated decision-making.

Further investigation should examine privacy-preserving information exchange, secure inter-operator coordination, and adaptive mechanisms for resolving conflicts under changing network conditions. The development of standardized interoperability benchmarks and conformance testing procedures would also support independent evaluation and broader industry adoption.

These directions provide opportunities to extend the framework beyond its initial design and establish its suitability for increasingly autonomous 5G and future 6G network environments.

Conclusion

Summary of the Research

This study proposed the Open Framework for Autonomous Network Coordination (OFANC) to address interoperability challenges among heterogeneous AI agents operating in multi-vendor telecommunications networks. The research examined the limitations of independently managed network intelligence and developed a common architectural approach for coordinating agents across radio access, core, transport, and service management domains.

The proposed framework integrates agent discovery, semantic interoperability, distributed intent negotiation, conflict resolution, and policy-aware decision validation. Its architecture builds on existing telecommunications management standards while introducing mechanisms intended to support consistent communication and coordinated operations among independently developed agents.

Principal Contributions

The principal contribution is a five-layer interoperability architecture that separates network integration, agent communication, semantic interpretation, autonomous coordination, and service governance. The study also introduced a common semantic message structure, a bounded coordination algorithm, and a formal optimization model that considers service-level violations, coordination latency, and communication overhead.

The descriptive model gives B4 compatibility of 83.58% and fulfillment of 79.53%, versus 75.89% and 63.85% for B3, because those outcomes were generated from different assumed probabilities.

It also assigns B4 higher signaling cost (0.579 KB/task), longer recovery time (51.04 s), and incomplete unauthorized-request blocking (88.36%). These results cannot demonstrate real-world superiority, and the longer recovery delay does not support H2.

Implications for Telecommunications Research

The study contributes to research on autonomous telecommunications management by addressing interoperability as a problem involving shared semantics, coordinated decisions, and operational authority rather than communication interfaces alone.

Its proposed integration of heterogeneous agents offers a basis for investigating end-to-end service assurance, network slicing, resource optimization, and cross-domain fault management. The framework also provides a structured approach for examining the relationship between coordination effectiveness, operational complexity, and network performance.

The experimental design and performance indicators establish a foundation for future comparisons involving independent vendor implementations and standards-aligned network management systems.

Concluding Perspective

Interoperability among autonomous AI agents is an important requirement for the development of coordinated and reliable telecommunications networks. As intelligence becomes increasingly distributed across network functions and management domains, the ability to exchange meaningful information and execute compatible decisions will influence the effectiveness of autonomous operations.

OFANC provides a structured approach to addressing these requirements through common semantic representations, controlled negotiation, and policy-aware coordination. Its practical value will depend on demonstrating that these mechanisms improve network management outcomes without introducing unacceptable latency, complexity, or security risks.

The proposed architecture is ready for independent implementation and protocol-level testing, not production adoption on the strength of the present numerical illustration. The next research step is to deploy independently developed agents, collect real operational traces, validate security policies, and test heterogeneous vendor interfaces under identical workloads.

References

Atalay, T. O., Maitra, S., Stojadinovic, D., Stavrou, A., & Wang, H. (2025). An OpenRAN security framework for scalable authentication, authorization, and discovery of xApps with isolated critical services. IEEE Transactions on Dependable and Secure Computing, 22(3), 2873–2890. https://doi.org/10.1109/TDSC.2024.3522218

Azariah, W., Bimo, F. A., Lin, C.-W., Cheng, R.-G., Nikaein, N., & Jana, R. (2024). A survey on open radio access networks: Challenges, research directions, and open source approaches. Sensors, 24(3), 1038. https://doi.org/10.3390/s24031038

Barakabitze, A. A., Ahmad, A., Mijumbi, R., & Hines, A. (2020). 5G network slicing using SDN and NFV: A survey of taxonomy, architectures and future challenges. Computer Networks, 167, 106984. https://doi.org/10.1016/j.comnet.2019.106984

Benzaid, C., & Taleb, T. (2020). AI-driven zero touch network and service management in 5G and beyond: Challenges and research directions. IEEE Network, 34(2), 186–194. https://doi.org/10.1109/MNET.001.1900252

Bonati, L., D'Oro, S., Polese, M., Basagni, S., & Melodia, T. (2021). Intelligence and learning in O-RAN for data-driven NextG cellular networks. IEEE Communications Magazine, 59(10), 21–27. https://doi.org/10.1109/MCOM.101.2001120

Bonati, L., Polese, M., D'Oro, S., Basagni, S., & Melodia, T. (2023). OpenRAN Gym: AI/ML development, data collection, and testing for O-RAN on PAWR platforms. Computer Networks, 220, 109502. https://doi.org/10.1016/j.comnet.2022.109502

Coronado, E., Behravesh, R., Subramanya, T., Fernández-Fernández, A., Siddiqui, M. S., Costa-Pérez, X., & Riggio, R. (2022). Zero touch management: A survey of network automation solutions for 5G and 6G networks. IEEE Communications Surveys & Tutorials, 24(4), 2535–2578. https://doi.org/10.1109/COMST.2022.3212586

Dryjański, M., Kułacz, Ł., & Kliks, A. (2021). Toward modular and flexible Open RAN implementations in 6G networks: Traffic steering use case and O-RAN xApps. Sensors, 21(24), 8173. https://doi.org/10.3390/s21248173

El Rajab, M., Yang, L., & Shami, A. (2024). Zero-touch networks: Towards next-generation network automation. Computer Networks, 243, 110294. https://doi.org/10.1016/j.comnet.2024.110294

Ergun, S., Sammour, I., & Chalhoub, G. (2023). A survey on how network simulators serve reinforcement learning in wireless networks. Computer Networks, 234, 109934. https://doi.org/10.1016/j.comnet.2023.109934

Fdida, S., Makris, N., Korakis, T., Bruno, R., Passarella, A., Andreou, P., Belter, B., Crettaz, C., Dabbous, W., Demchenko, Y., & Knopp, R. (2022). SLICES, a scientific instrument for the networking community. Computer Communications, 193, 189–203. https://doi.org/10.1016/j.comcom.2022.07.019

Foukas, X., Patounas, G., Elmokashfi, A., & Marina, M. K. (2017). Network slicing in 5G: Survey and challenges. IEEE Communications Magazine, 55(5), 94–100. https://doi.org/10.1109/MCOM.2017.1600951

Gallego-Madrid, J., Sanchez-Iborra, R., Ruiz, P. M., & Skarmeta, A. F. (2022). Machine learning-based zero-touch network and service management: A survey. Digital Communications and Networks, 8(2), 105–123. https://doi.org/10.1016/j.dcan.2021.09.001

Gavrilovska, L., Rakovic, V., & Denkovski, D. (2020). From cloud RAN to open RAN. Wireless Personal Communications, 113, 1523–1539. https://doi.org/10.1007/s11277-020-07231-3

Giordani, M., Polese, M., Mezzavilla, M., Rangan, S., & Zorzi, M. (2020). Toward 6G networks: Use cases and technologies. IEEE Communications Magazine, 58(3), 55–61. https://doi.org/10.1109/MCOM.001.1900411

Groen, J., D'Oro, S., Demir, U., Bonati, L., Polese, M., Melodia, T., & Chowdhury, K. R. (2025). Implementing and evaluating security in O-RAN: Interfaces, intelligence, and platforms. IEEE Network, 39(1), 227–234. https://doi.org/10.1109/MNET.2024.3434419

Johnson, D., Maas, D., & Van Der Merwe, J. (2022). NexRAN: Closed-loop RAN slicing in POWDER—A top-to-bottom open-source Open-RAN use case. Proceedings of the 15th ACM Workshop on Wireless Network Testbeds, Experimental Evaluation & Characterization, 17–23. https://doi.org/10.1145/3477086.3480842

Koutlia, K., & Lagén, S. (2024). On the impact of Open RAN fronthaul control in scenarios with XR traffic. Computer Networks, 253, 110722. https://doi.org/10.1016/j.comnet.2024.110722

Kułacz, Ł., & Kliks, A. (2022). Dynamic spectrum allocation using multi-source context information in OpenRAN networks. Sensors, 22(9), 3515. https://doi.org/10.3390/s22093515

Lacava, A., Bonati, L., Mohamadi, N., Gangula, R., Kaltenberger, F., Johari, P., D'Oro, S., Cuomo, F., Polese, M., & Melodia, T. (2025). dApps: Enabling real-time AI-based Open RAN control. Computer Networks, 269, 111342. https://doi.org/10.1016/j.comnet.2025.111342

Lacava, A., Polese, M., Sivaraj, R., Soundrarajan, R., Bhati, B. S., Singh, T., Zugno, T., Cuomo, F., & Melodia, T. (2024). Programmable and customized intelligence for traffic steering in 5G networks using Open RAN architectures. IEEE Transactions on Mobile Computing, 23(4), 2882–2897. https://doi.org/10.1109/TMC.2023.3266642

Leivadeas, A., & Falkner, M. (2023). A survey on intent-based networking. IEEE Communications Surveys & Tutorials, 25(1), 625–655. https://doi.org/10.1109/COMST.2022.3215919

Moudoud, H., Abou El Houda, Z., & Brik, B. (2024). Zero trust security architecture for 6G open radio access networks (ORAN). IEEE Networking Letters, 6(4), 272–275. https://doi.org/10.1109/LNET.2024.3514357

Nezami, Z., Zaidi, S. A. R., Hafeez, M., Xu, J., & Djemame, K. (2025). Toward standardization of GenAI-driven agentic architectures for radio access networks. Frontiers in Artificial Intelligence, 8, 1621963. https://doi.org/10.3389/frai.2025.1621963

Polese, M., Bonati, L., D'Oro, S., Basagni, S., & Melodia, T. (2023a). ColO-RAN: Developing machine learning-based xApps for Open RAN closed-loop control on programmable experimental platforms. IEEE Transactions on Mobile Computing, 22(10), 5787–5800. https://doi.org/10.1109/TMC.2022.3188013

Polese, M., Bonati, L., D'Oro, S., Basagni, S., & Melodia, T. (2023b). Understanding O-RAN: Architecture, interfaces, algorithms, security, and research challenges. IEEE Communications Surveys & Tutorials, 25(2), 1376–1411. https://doi.org/10.1109/COMST.2023.3239220

Rekkas, V. P., Sotiroudis, S., Sarigiannidis, P., Wan, S., Karagiannidis, G. K., & Goudos, S. K. (2021). Machine learning in beyond 5G/6G networks—State-of-the-art and future trends. Electronics, 10(22), 2786. https://doi.org/10.3390/electronics10222786

Wadud, A., Golpayegani, F., & Afraz, N. (2024). QACM: QoS-aware xApp conflict mitigation in Open RAN. IEEE Transactions on Green Communications and Networking, 8(3), 978–993. https://doi.org/10.1109/TGCN.2024.3431945

Wasilewska, M., Kliks, A., Bogucka, H., Cichon, K., Ruseckas, J., Molis, G., Mackutė-Varoneckienė, A., & Kirviliučius, T. (2021). Artificial intelligence for radio communication context-awareness. IEEE Access, 9, 144820–144856. https://doi.org/10.1109/ACCESS.2021.3119524

Zhang, Z., Xiao, Y., Ma, Z., Xiao, M., Ding, Z., Lei, X., Karagiannidis, G. K., & Fan, P. (2019). 6G wireless networks: Vision, requirements, architecture, and key technologies. IEEE Vehicular Technology Magazine, 14(3), 28–41. https://doi.org/10.1109/MVT.2019.2921208