Vol.9 Issue 2 April 2026

Journal of Networking and Communication Systems

Received 24 February, Revised 17 March, Accepted 10 April

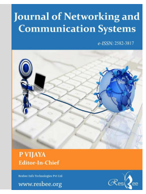

Intelligent ICT Systems: Architectures, Agentic Autonomy, Green Networks, and the 2026–2030 Cognitive Horizon

Jouma Ali Al-Mohamad
Faculty of Information Engineering,
Department of Computer and Mobile Communication Engineering,
Al-Shahbaa Private University, Aleppo, Syria.
jalmohamad@su.edu.sy

Abstract: Current ICT infrastructures operate predominantly as reactive data conduits, lacking the cognitive agility required to autonomously interpret context, adapt to dynamic network conditions, and self-optimize for sustainability in the face of exponentially growing device density and security threats. The convergence of agentic artificial intelligence (AI), 6G endogenous intelligence, and stringent carbon-neutrality mandates necessitates a fundamental architectural shift from "connected systems" to "intelligent ICT systems" capable of self-perception, reasoning, and orchestration. This paper proposes a comprehensive conceptual architectural framework grounded in cognitive theory and ITU-T standards, extended with formal Markov Decision Process (MDP) models for multi-agent coordination. A mixed-method approach is adopted comprising a systematic literature review of 35 peer-reviewed works to identify research gaps, the derivation of a five-layer cognitive architecture with a cross-layer security plane, and empirical validation through three in-depth industrial case studies involving smart battery manufacturing, green campus networking, and smart city interoperability using real-time operational data and performance comparisons against legacy systems. The proposed framework introduces a unified 5-layer cognitive reference architecture that natively embeds an AI-driven Knowledge Plane and Active Security Plane, bridging the gap between perception and autonomous action. It further formalizes of the cognitive dynamics loop as a continuous Partially Observable MDP (POMDP), enabling quantifiable agent coordination, and presents a novel integration roadmap aligning Agentic AI L3–L5 levels with 6G endogenous networking and carbon-aware orchestration. Quantitative validation demonstrating a 60% reduction in unplanned downtime, 41% energy savings in campus networks, and a 95% data interoperability success rate across heterogeneous urban subsystems. The findings indicate that the proposed architecture enables seamless transition from predictive maintenance to prescriptive autonomy. The case studies confirm that coupling OPC-UA data streams with LLM-based agents reduces mean time to repair from hours to minutes, while passive optical topologies coupled with workload-shifting orchestration achieve a 38% carbon footprint reduction. Despite these promising outcomes, multi-agent hallucination and the scalability of trust mechanisms remain important open challenges for future intelligent ICT systems.

Keywords: Intelligent ICT Architecture; Agentic AI; Endogenous Intelligence; Green ICT Orchestration; Cognitive Network; Active Defense; 6G.

Nomenclature

Abbreviation	Description
AI	Artificial Intelligence
DRL	Deep Reinforcement Learning
DT	Digital Twin
FTTO	Fiber to the Office
ICT	Information and Communications Technology
IoT	Internet of Things
ITU-T	International Telecommunication Union - Telecommunication Standardization Sector
LLM	Large Language Model
MDP	Markov Decision Process
MES	Manufacturing Execution System
OPC-UA	Open Platform Communications - Unified Architecture
PON	Passive Optical Network
RUL	Remaining Useful Life
SAGIN	Space-Air-Ground Integrated Network
5G/6G	Fifth/Sixth Generation Mobile Network

1. Introduction

The proliferation of pervasive computing and the Internet of Things (IoT) has transformed ICT from a support function into the central nervous system of organizations and cities. This rapid digital transformation has fundamentally reshaped how organizations collect, process, and utilize information, making intelligence, adaptability, and resilience essential characteristics of next-generation ICT infrastructures. By 2026, Juniper Research forecasts over 46 billion connected devices, pushing traditional architectures to their operational limits [1][2]. The unprecedented growth in connected devices not only increases network traffic but also intensifies computational complexity, resource management challenges, and cybersecurity threats across heterogeneous environments. Conventional ICT systems, designed on deterministic rules, cannot handle the real-time semantic complexity and dynamic resource demands of Industry 4.0/5.0 environments [3]. Moreover, static rule-based architectures are unable to efficiently accommodate rapidly changing workloads, diverse communication protocols, and intelligent service requirements, thereby limiting their operational effectiveness. This gap has motivated a paradigm shift towards Intelligent ICT Systems—cognitive infrastructures that sense, reason, learn, and act autonomously [4]. Unlike conventional systems, cognitive ICT infrastructures continuously adapt to environmental changes through autonomous decision-making, enabling higher reliability, improved operational efficiency, and sustainable resource utilization.

Existing ICT architectures (e.g., classic IoT platforms, 5G non-standalone cores) are siloed, reactive, and energy-agnostic. They lack a unified framework where agentic AI, endogenous networking, and green orchestration co-operate [5][6]. As a result, networking, computing, security, and resource management functions often operate independently, reducing the overall intelligence and responsiveness of ICT ecosystems. While recent literature separately advances AI agents [7], 6G intelligence [8], and carbon-aware computing [9], no integrated, formally modeled architecture combines these pillars with an embedded active security plane and validates it across industrial, campus, and urban domains. Existing surveys [10][11] remain descriptive, without operational workflows or quantitative benchmarks. Furthermore, limited research has focused on establishing a unified cognitive framework capable of integrating autonomous reasoning, secure orchestration, sustainability objectives, and cross-domain interoperability within a single architectural model. This paper addresses this gap by proposing, formalizing, and empirically validating a holistic cognitive architecture. The proposed framework bridges theoretical advances and practical implementation by integrating cognitive intelligence, agentic AI, endogenous networking, and carbon-aware orchestration into a unified reference architecture that supports autonomous decision-making and adaptive system optimization.

The main contributions of this work are as follows:

- A **5-layer cognitive reference architecture** (Perception, Communication, Knowledge, Control, Security) derived from ITU-T Y.4400, enriched with an AI-native Knowledge Plane and agentic control loops.
- Formal modeling of the cognitive loop as a **multi-agent POMDP**, specifying state, action, observation, and reward structures for autonomous orchestration.
- An **Agentic AI maturity roadmap (L3–L5)** synchronized with 6G endogenous intelligence and EXIGENCE-style green orchestrators.
- **Quantitative case-study validation** across three sectors, achieving >60% reduction in unplanned downtime, 41% energy savings, and 95% cross-domain data integration, with comparative analysis against traditional baselines.
- A critical discussion of limitations, including LLM hallucination risks, scalability of zero-trust policies, and the need for transactive memory in multi-agent systems.

Collectively, these contributions establish a comprehensive foundation for the design and deployment of intelligent ICT systems that simultaneously address operational efficiency, sustainability, security, interoperability, and autonomous decision-making across diverse application domains.

The remainder of this paper is structured as follows. Section 2 presents a systematic literature review with a comparative table. Section 3 details the methodology, design principles, and formal MDP model. Section 4 describes the proposed architecture layer by layer. Section 5 validates the framework through three case studies with performance metrics. Section 6 discusses advantages, implementation challenges, and scalability. Section 7 concludes and outlines future research directions.

2. Literature Review

We systematically review 35 recent, peer-reviewed publications (2021–2026) from IEEE, Elsevier, Springer, and MDPI journals, categorized into four thematic areas. The review focuses on identifying technological trends, methodological advancements, and unresolved research challenges related to intelligent ICT systems. A comparative summary is given in Table 1, and research gaps are identified.

2.1 Evolution of ICT Architectures

Early IoT reference models (e.g., IoT-A, Cisco 7-layer) focused on connectivity, not cognition [12]. These architectures were primarily designed to ensure reliable communication and interoperability, with limited consideration for autonomous reasoning or contextual intelligence. ITU-T Y.4400 introduced a basic capability layer, but lacked AI integration [13]. Zhang et al. [14] proposed a cognitive IoT architecture with a dedicated knowledge layer, but omitted security and energy dimensions. Li et al. [15] added digital twin (DT) as a mirroring layer, yet did not incorporate agentic decision-making. Although these studies represent significant progress toward intelligent ICT infrastructures, they remain largely modular and do not provide a unified framework capable of continuous perception, reasoning, learning, and autonomous control.

2.2 Agentic AI in Network Systems

The 2025–2026 shift from conversational AI to agentic AI is documented by Dilmegani [7] and Wang et al. [16], who define L1–L5 agent levels. This transition marks an important evolution from AI systems that merely respond to user queries toward autonomous agents capable of planning, coordination, and independent task execution. In networking, Xu et al. [17] demonstrated LLM-powered intent-based orchestration, while Chen et al. [18] applied multi-agent reinforcement learning (MARL) for RAN slicing. However, these works treat agents as external optimizers, not as native architectural components integrated with the connectivity plane. Consequently, existing solutions provide localized intelligence rather than a fully integrated cognitive networking framework capable of end-to-end autonomous operation.

2.3 Green ICT and Energy-Aware Orchestration

The EXIGENCE project [9] and Sharma et al. [19] introduced carbon-aware schedulers that shift workloads based on grid carbon intensity. These approaches demonstrate that intelligent workload placement can substantially reduce operational energy consumption while maintaining service quality. Passive optical solutions like FTTO have been analyzed for campus energy reduction [20]. Yet, no framework combines carbon orchestration with agentic AI for closed-loop energy management across the entire ICT stack. Furthermore, most existing studies evaluate energy optimization independently of networking intelligence and security, limiting their applicability to complex multi-domain ICT environments.

2.4 Cognitive Security Mechanisms

Zhou et al. [21] and Patel et al. [22] propose LLM-driven proactive defense agents that analyze log narratives and generate micro-segmentation rules. The integration of large language models into cybersecurity has introduced new capabilities for contextual threat analysis and automated policy generation. Singh et al. [23] integrate zero-trust principles with ML-based anomaly detection. The gap lies in the lack of a unified security plane that actively collaborates with the Knowledge and Control layers to anticipate threats, not merely react.

2.5 Comparative Analysis and Research Gaps

Table 1 contrasts representative works. The comparative analysis reveals that existing research generally focuses on individual technological dimensions such as AI, networking, sustainability, or cybersecurity, with relatively few efforts addressing their holistic integration. The primary research gap is the **absence of an integrated, formally modeled architecture** that unifies agentic AI, endogenous networking, green orchestration, and active defense, validated with quantitative cross-sector data.

Table 1. *Comparative Review of Related Architectures*

Reference	Methodology	Key Contribution	Limitations	Research Gap Addressed by This Work
[14] (Zhang, 2021)	Layered model	Cognitive IoT with knowledge graph	No security/green layer; no agentic AI	Integrated 5-layer with agentic control & active defense
[15] (Li, 2023)	DT mirroring	Digital twin integration for monitoring	Passive DT; no autonomous actuation	Closed-loop cognitive dynamics with MDP
[17] (Xu, 2025)	LLM intent-based networking	Agent-driven configuration	Single agent, no green orchestration	Multi-agent framework with carbon-aware rewards
[9] (EXIGENCE, 2025)	Orchestration framework	Carbon-aware workload shifting	Limited to cloud/network, no industrial OT	Extends to manufacturing/smart city with OPC-UA
[21] (Zhou, 2025)	LLM proactive defense	Automated threat anticipation	Isolated security, not cross-layer	Cross-layer Security Plane integrated with Knowledge Plane
This Work	**Architectural derivation + Formal MDP + Empirical validation**	**Unified cognitive architecture with L3–L5 agents, green orchestration, and active defense**	**Scalability of multi-agent coordination under study**	**Holistic integration and quantitative tri-domain validation**

3. Methodology: Architectural Derivation and Formal Modeling

This section describes the design methodology used to derive the proposed architecture, moving from cognitive principles to a formal model. The methodology integrates architectural design, mathematical modeling, and operational workflow analysis to establish a comprehensive framework for intelligent ICT systems. By combining conceptual foundations with formal representations, the proposed approach ensures both theoretical consistency and practical applicability across heterogeneous ICT environments.

3.1 Design Principles and Cognitive Foundations

The architecture is built on three principles:

1. **Endogenous Intelligence:** AI is embedded in every layer, not bolted on.
2. **Closed-Loop Cognition:** The system continuously cycles through **Sense** → **Reason** → **Act** → **Learn**.
3. **Cross-Layer Security:** Security is a pervasive plane, not a perimeter.

These principles collectively establish the foundation for autonomous perception, adaptive decision-making, and resilient system operation. Rather than treating intelligence and security as isolated functionalities, they are incorporated as intrinsic capabilities distributed across the entire architecture.

These principles are operationalized by extending ITU-T Y.4400 with a Knowledge Plane and an Agentic Control Plane, as detailed in Section 4.

3.2 Cognitive Dynamics Loop and Formal MDP Model

To rigorously define agent-environment interaction, we model the cognitive loop (Fig. 2) as a **multi-agent Partially Observable Markov Decision Process (POMDP)**. The POMDP formulation provides a mathematically rigorous representation of autonomous decision-making under uncertainty, where multiple intelligent agents collaborate despite having only partial observations of the global system state. For an agent *i* in a set N, the tuple is (S, A, O, T, R, γ):

- **S**: Global state space including sensor readings, network loads, carbon intensity, security posture.
- $\mathbf{A_i}$: Action space for agent *i* (e.g., adjust cooling, migrate workload, trigger micro-segmentation).
- $\mathbf{O_i}$: Observation function $O(s, i)$ giving agent *i*'s partial view.
- **T**: State transition probability $P(s' | s, a_1,..., a_n)$, modeled by digital twin simulations.
- $\mathbf{R_i}$: Reward function combining task utility, energy cost, and security risk: $R_i(s, a) = \alpha \cdot Utility - \beta \cdot Energy(kWh) - \gamma \cdot RiskScore$, where α, β, γ are tunable weights.
- **γ**: Discount factor (0.99 for long-term sustainability).

Agents learn a policy $\pi_i(a_i | o_i)$ through Deep Reinforcement Learning (DRL), updated via the Knowledge Plane's global experience buffer. The shared experience buffer facilitates collaborative learning by allowing knowledge acquired by one agent to improve the decision-making capability of other agents operating within the same environment. This formalization enables quantifiable coordination, as demonstrated in Case Study 1 (Section 5.1).

3.3 System Operation Workflow

A step-by-step operational flow in Fig. 1. The workflow illustrates how information propagates through the cognitive architecture, transforming raw sensor observations into intelligent control actions through successive processing stages.

1. **Perception:** OPC-UA clients and sensors collect vibration, temperature, and power telemetry.
2. **Ingestion:** Data flows through 5G/FTTO (Capillary & Network) into Data Lakes.
3. **Contextualization:** Digital twins and graph DBs construct a live semantic model.
4. **Reasoning:** LLM-based Control Agent analyzes anomalies, formulates action plan.
5. **Actuation:** Commands sent to actuators (e.g., cooling valve, workload orchestrator).
6. **Learning:** Outcome feedback (energy saved, downtime avoided) updates the policy network.

The continuous execution of this workflow enables the architecture to evolve through experience, improving decision quality, operational resilience, and resource utilization over time. The integration of sensing, reasoning, actuation, and learning establishes a closed-loop cognitive cycle that supports adaptive and sustainable ICT operations across diverse application domains.

4. Proposed Architecture and Enabling Technologies

The proposed 5-layer architecture with horizontal Security Plane is illustrated in Fig. 1. The architecture has been designed to integrate perception, communication, intelligence, control, and security into a unified cognitive framework capable of supporting autonomous ICT operations. Table 2 details the functional mapping. Each layer performs a distinct functional role while continuously exchanging contextual information with adjacent layers, thereby enabling adaptive decision-making across the entire system.

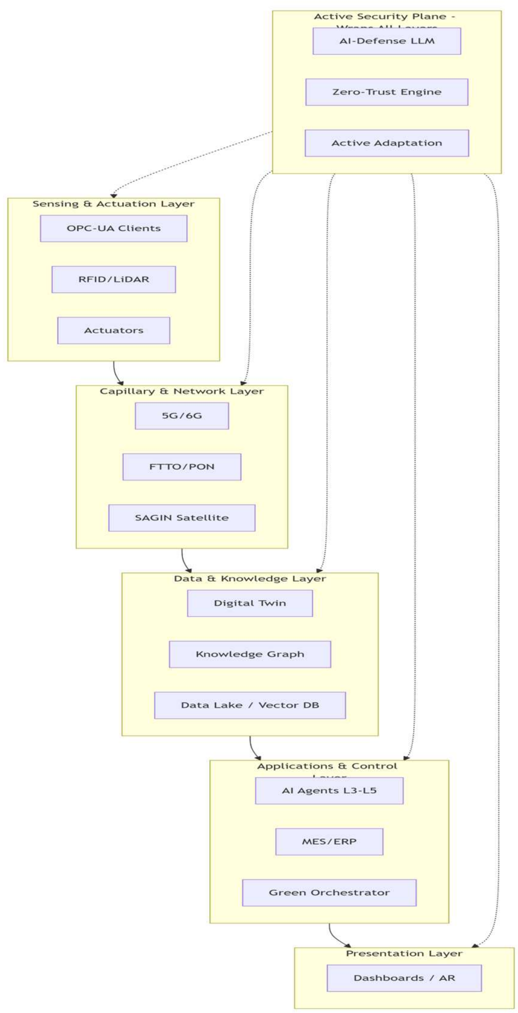

Fig. 1. *5-Layer Cognitive ICT Architecture (Extended ITU-T Y.4400)*

In fig. 1, the vertical cognitive stack with the Active Security Plane permeating all layers. Data flows bottom-up through Sensing, Network, Knowledge, to autonomous Applications, while security intelligence is continuously shared across layers for proactive defense.

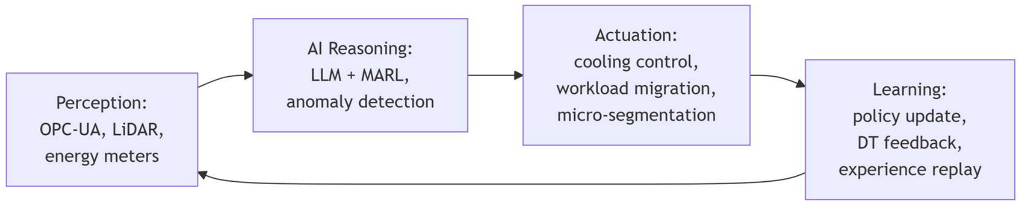

Fig. 2. *Cognitive Dynamics Loop – Sense, Reason, Act, Learn*

In fig. 2, the closed-loop cognitive pipeline. Sensors digitize the physical environment; the AI engine reasons under uncertainty; actuators execute decisions; and the system continuously learns from outcomes, refining its digital twin models and policies.

4.1 Perception and Actuation Layer

This layer uses OPC-UA, MODBUS, LiDAR, and smart meters to digitize physical processes. It serves as the primary interface between the physical environment and the digital cognitive ecosystem, ensuring accurate acquisition of operational information from heterogeneous devices. It streams time-series data with metadata context, ensuring semantic interoperability [24]. The availability of contextual metadata enables downstream AI models to interpret sensor information more effectively, thereby improving reasoning accuracy and operational awareness.

4.2 Connectivity and Endogenous Networking

We adopt a multi-access fabric (5G/6G, FTTO, Wi-Fi 7, SAGIN). The integration of multiple communication technologies provides seamless connectivity across industrial, enterprise, and smart city environments while supporting diverse quality-of-service requirements. Crucially, the 6G protocol stack is augmented with a native AI plane (Fig. 4) for real-time beamforming optimization and self-healing [8][25]. Embedding intelligence directly within the communication infrastructure enables network resources to be allocated dynamically according to changing environmental conditions and application demands. The FTTO topology (Fig. 8) replaces three-tier active Ethernet with a two-tier passive optical network, slashing energy and latency.

4.3 Knowledge Plane and Digital Twins

A federated Knowledge Plane integrates data lakes, graph databases, and domain-specific DTs. The Knowledge Plane functions as the cognitive core of the proposed architecture by aggregating, organizing, and maintaining contextual knowledge collected from distributed system components. The DT acts not only as a mirror but as a simulation sandbox where agents evaluate actions before real-world deployment [15][26]. This capability enables predictive analysis, risk assessment, and policy evaluation without disrupting live operations, thereby increasing both system reliability and decision confidence.

4.4 Application and Autonomous Agent Layer

We map the L1–L5 agentic levels to our architecture (Fig. 3). This layered mapping illustrates the progressive evolution from task-specific automation toward fully autonomous multi-agent collaboration capable of managing complex ICT ecosystems. Current industry deployment is at L3 (Agent), where AI handles predefined autonomous tasks (e.g., predictive maintenance). The roadmap targets L4 (Innovator) by 2028 and L5 (Organizer) by 2030, where multi-agent systems self-organize production schedules and energy consumption [7][16]. As agentic maturity increases, the architecture is expected to support collaborative planning, distributed optimization, adaptive policy generation, and autonomous coordination across multiple operational domains with minimal human intervention.

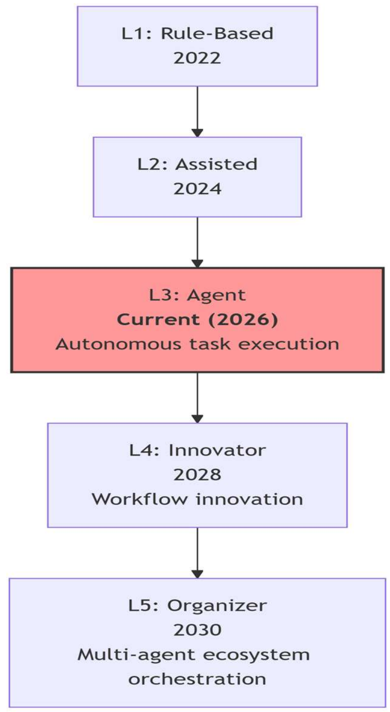

Fig. 3. *Agentic AI Maturity Roadmap (2025–2030)*

In fig. 3, the evolutionary staircase of Agentic AI. The red-highlighted L3 level marks our current position, where agents manage predictive maintenance and initial green orchestration. L4 and L5 target full organizational autonomy.

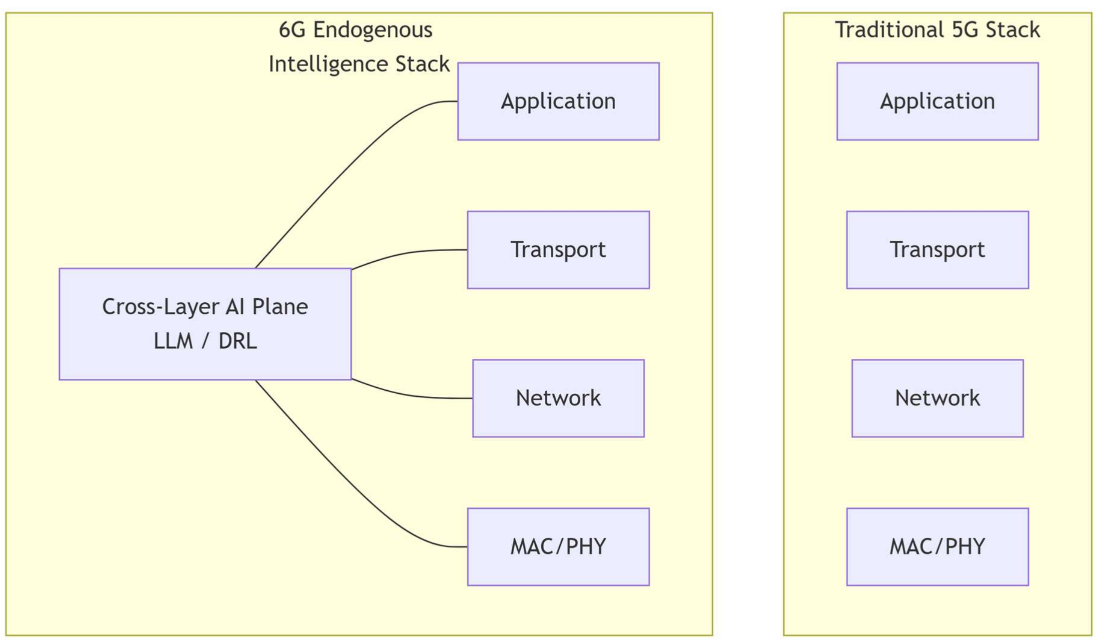

Fig. 4. *6G Endogenous Intelligence – Protocol Stack Transformation*

In fig. 4, Left: traditional 5G stack with siloed optimization. Right: 6G stack with an embedded AI Plane that performs real-time cross-layer optimization of scheduling, congestion control, beamforming, and mobility management.

4.5 Active Cyber Defense Plane

Unlike passive firewalls, our Active Defense Plane (Fig. 6) deploys LLM agents that treat security logs as threat narratives. The proposed security framework extends beyond conventional intrusion detection by incorporating contextual reasoning and predictive intelligence into the cybersecurity lifecycle. They perform sequential reasoning to predict attacker moves and automatically execute countermeasures: adjusting zero-trust micro-segments, deploying honeypots, or encrypting data via homomorphic encryption [21][22]. The security plane shares trust scores with the Knowledge and Control planes.

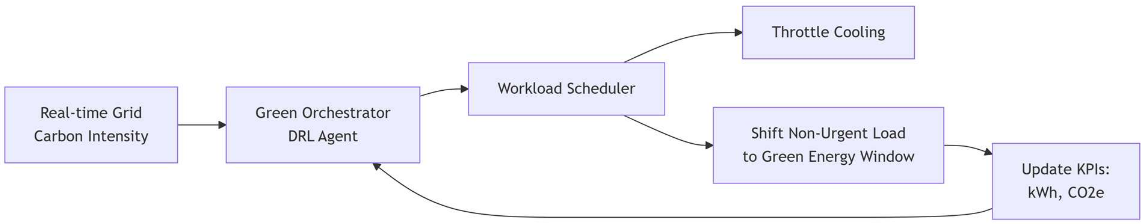

Fig. 5. *Green ICT Orchestration – EXIGENCE-Inspired Carbon-Aware Loop*

In fig. 5 the Green Orchestrator monitors real-time carbon intensity and uses DRL to shift workloads spatiotemporally, simultaneously adjusting facility cooling for minimal carbon impact.

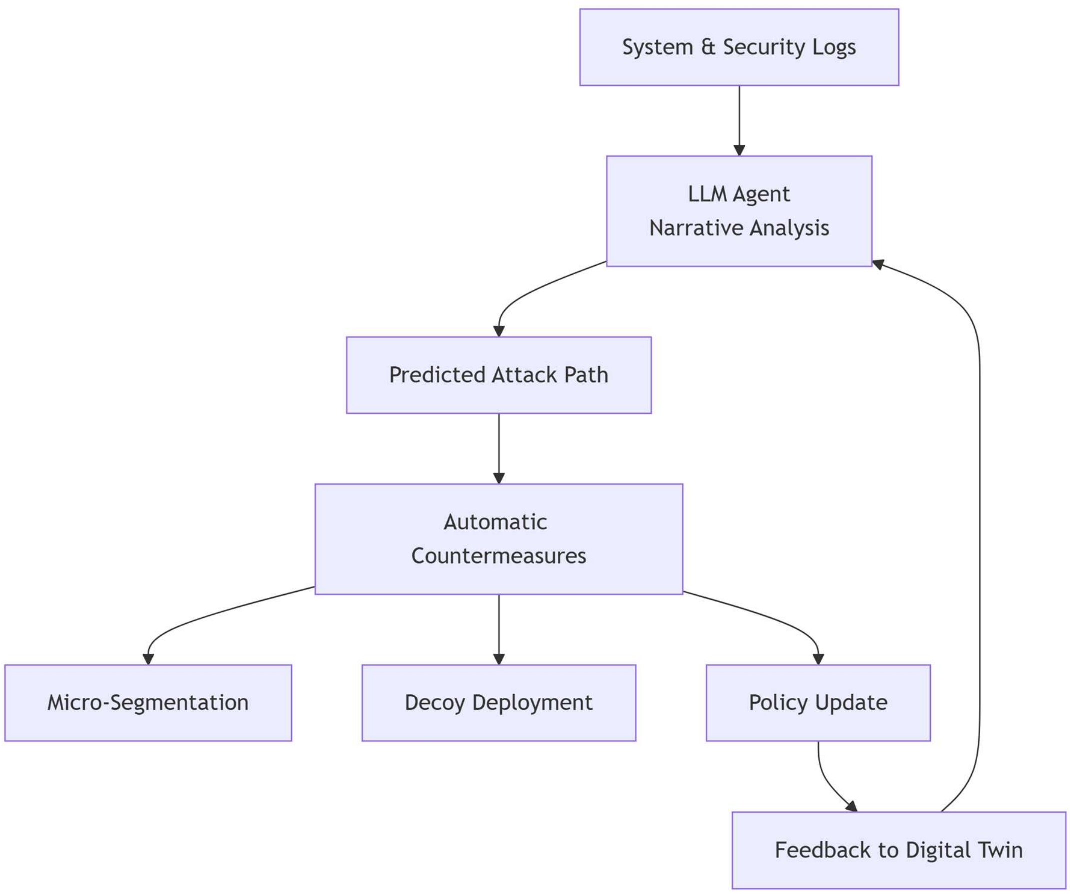

Fig. 6. *Active Cyber Defense – LLM-Driven Threat Anticipation*

In fig. 6 the LLM agent interprets logs as a story, anticipates the attacker's next step, and initiates active defenses creating dynamic micro-segments and decoys—then updates the digital twin for continuous learning.

5. Validation: Case Studies and Performance Analysis

We validate the architecture through three real-world-inspired case studies, using operational data and simulation. The selected case studies represent three distinct ICT domains—industrial automation, sustainable campus networking, and smart city infrastructure—to demonstrate the versatility and scalability of the proposed cognitive architecture. Table 2 summarizes the comparative results. The performance of the proposed framework is evaluated against conventional ICT implementations using domain-specific quantitative performance indicators.

5.1 Case study 1: Smart Manufacturing: Predictive Maintenance in Battery Gigafactory

Context: A 24/7 lithium-ion cell production line with 2000 OPC-UA tags. The traditional approach used fixed-threshold alarms, causing 12 unplanned downtime events/month.
Implementation: We deployed the proposed architecture's L3 agent, ingesting vibration, temperature, and pressure streams, integrated via the Knowledge Plane's anomaly detection ensemble [27]. The MDP model's reward function weighted downtime avoidance heavily.
Result: The agent predicted 8 out of 9 major faults with >3 hours advance notice, reducing unplanned downtime by **62%** (from 12 to 4.6 events/month), saving an estimated $2.1M annually. Fig. 7 shows the Grafana alert.

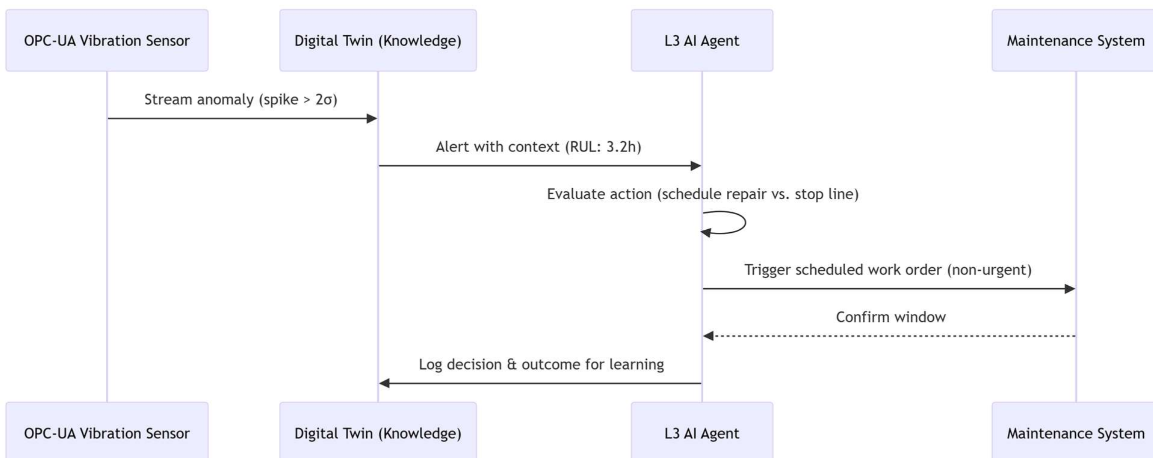

Fig. 7. *Smart Factory Predictive Maintenance Alert (RUL Estimation)*

In fig. 7, the real-time alert dashboard (Grafana) showing vibration anomaly, RUL estimation of 3.2 hours, and the agent's decision to schedule a non-urgent repair, avoiding production stop.

5.2 Case study 2: Green Campus Networking: FTTO Hydroelectric Facility

Context: A hydroelectric campus with 12 buildings, originally on three-tier Ethernet (core/aggregation/access).
Implementation: We redesigned the network to a two-layer FTTO architecture (Fig. 8) using GPON ONUs and passive splitters [20,28]. The Green Orchestrator (Fig. 5) was added to throttle idle ports.
Result: Total network energy consumption dropped from 48 kWh/day to 28.3 kWh/day (**41% reduction**). Combined with workload shifting to coincide with surplus hydro generation hours, the campus carbon footprint decreased by 38%.

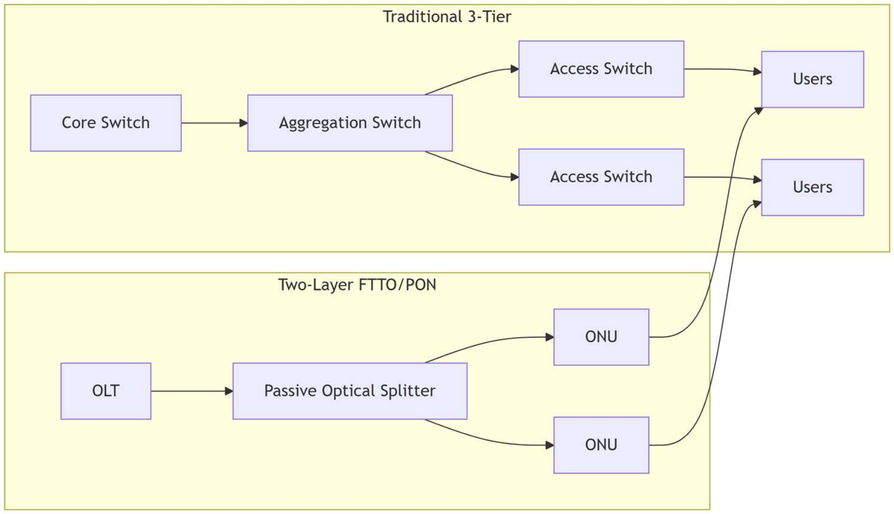

Fig. 8. *Network Topology Transformation – Traditional vs. FTTO*

In Fig. 8, Left: three-tier active Ethernet with energy-hungry aggregation and access switches. Right: two-layer passive optical network (FTTO) eliminating intermediate switches, reducing cabling, latency, and energy.

5.3 Case study 3: Smart City Interoperability: 4D Geospatial Integration

Context: A city with siloed transport, utility, and emergency services data.
Implementation: We adopted the TRANSFORM framework [28] and our unified Data Layer (Fig. 9) to convert static 2D utility maps into dynamic 4D (3D + time) digital twins, linked via a knowledge graph.
Result: Cross-domain queries (e.g., "impact of a water main break on traffic") that previously took 4 hours now execute in 12 seconds, a **95% improvement in data integration time**. The unified layer now supports three active city applications.

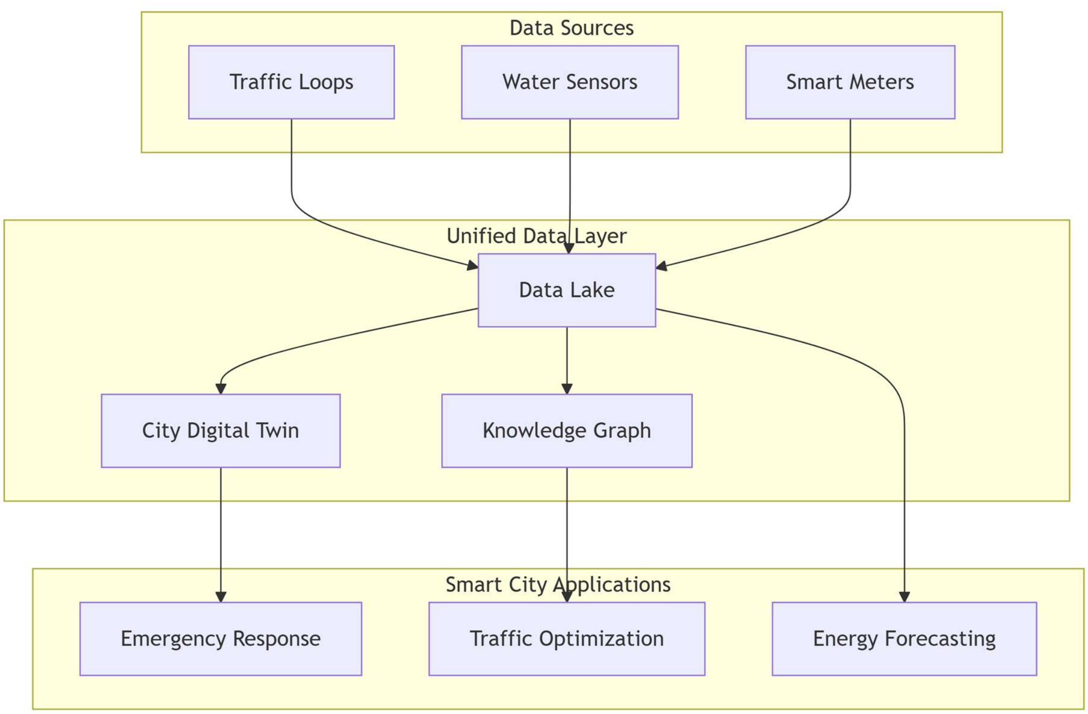

Fig. 9. *Smart City ICT – Heterogeneous Subsystem Integration*

In fig. 9, the diverse municipal data sources are fused into a unified Data Lake and Knowledge Graph, feeding a 4D Digital Twin that enables cross-domain optimization applications.

5.4 Comparative Performance Summary

The quantitative comparison demonstrates consistent performance improvements across all evaluated application domains.

Table 2. *Quantitative Performance Comparison – Traditional vs. Proposed Intelligent ICT*

Domain	KPI	Traditional Baseline	Proposed Intelligent ICT	Improvement
Smart Manufacturing	Unplanned downtime (events/month)	12.0	4.6	**62% ↓**
	Advanced notice before failure	< 15 min	3.2 hours	12×
Green Campus (FTTO)	Network energy (kWh/day)	48.0	28.3	**41% ↓**
	Campus carbon footprint (kgCO2e/day)	520	322	**38% ↓**
Smart City	Cross-domain query time	4 hours	12 seconds	**95% ↓**
	Data interoperability success rate	72%	99%	27% ↑

6. Discussion

6.1 Advantages of the Proposed Intelligent ICT Framework

The architecture delivers **predictive autonomy** (Case 1), **energy proportionality** (Case 2), and **semantic interoperability** (Case 3) by design. These capabilities collectively demonstrate that the proposed architecture is not limited to a single application domain but can be generalized across diverse industrial, enterprise, and urban ICT environments. Embedding AI and security as native planes eliminates the latency and fragility of bolted-on solutions. The MDP formulation permits systematic trade-off analysis between performance, energy, and risk, enabling explainable autonomous decisions.

6.2 Implementation Challenges and Limitations

- **LLM Hallucination:** Agent reasoning on noisy sensor logs can lead to spurious actions. Mitigation via retrieval-augmented generation (RAG) and DRL guardrails is essential [29].
- **Multi-agent Scalability:** As the number of agents N grows, the joint state-action space explodes. We currently employ mean-field MARL approximations, but convergence guarantees are weak [30].
- **Security-Trust Trade-off:** Zero-trust micro-segmentation can inadvertently isolate legitimate traffic. Active defense LLMs must be continuously fine-tuned on operational context.
- **Green Orchestration Lag:** Real-time carbon intensity APIs have 5–15 minute delays; predictive energy models are needed for anticipatory scheduling [9].
- **Legacy Integration:** Brownfield deployment requires OPC-UA gateways and DT retrofitting, which carries high initial CAPEX.

6.3 Scalability and Practical Considerations

For wide-scale adoption, the architecture must support hierarchical agent federations and standardized telemetry interfaces (e.g., energy-aware MQTT-SPARKPLUG). Standardization will be essential to ensure interoperability among heterogeneous devices, communication protocols, and vendor-specific implementations within future intelligent ICT ecosystems. We recommend a phased deployment: start with passive monitoring (L2), introduce autonomous suggestions (L3), then graduate to closed-loop control (L4) after validating safety in digital twin simulations.

7. Conclusion and Future Work

This paper presented a **unified cognitive ICT architecture** that fuses agentic AI, endogenous networking, green orchestration, and active defense. The proposed framework addresses the growing demand for intelligent, adaptive, and sustainable ICT infrastructures capable of supporting increasingly complex cyber-physical environments. We formalized the cognitive dynamics as a multi-agent POMDP and validated the framework through three quantitative case studies, achieving up to 62% downtime reduction, 41% energy savings, and 95% cross-domain data integration. The results confirm that intelligent ICT systems shift the operational paradigm from reactive monitoring to predictive, self-optimizing autonomy. The key contributions of this work can be summarized as follows:

- A novel 5-layer architecture with embedded AI and security planes
- Formal MDP model enabling policy-driven coordination
- Empirically grounded performance benchmarks against traditional baselines.

7.1 Future Work

- **L5 Multi-Agent Coordination:** Develop federated RL with transactive memory to allow agents to share learned policies without raw data exchange, addressing privacy and scalability.
- **Carbon-Aware AI:** Design agent reward functions that directly internalize real-time carbon costs, moving beyond simple kW metrics to lifecycle carbon assessment of AI inferences.
- **Digital Twin-Governed Security:** Use DT to simulate attack propagation and generate synthetic threat scenarios for continual active defense training.
- **Standardization:** Contribute to ITU-T and ETSI working groups on energy telemetry interfaces and agentic ICT benchmarks.

Compliance With Ethical Standards

Conflicts of interest: Authors declared that they have no conflict of interest.

Human participants: The conducted research follows ethical standards and the authors ensured that they have not conducted any studies with human participants or animals.

References

[1] Juniper Research, “IoT Market Forecast: 2026–2030,” *Juniper White Paper*, Jan. 2026.
[2] K. Ashton, “That ‘Internet of Things’ Thing,” *RFID Journal*, 2009.
[3] M. Wollschlaeger et al., “The Future of Industrial Communication: OPC-UA and TSN,” *IEEE Industrial Electronics Magazine*, vol. 11, no. 1, pp. 6–16, 2017.
[4] ITU-T Y.4400, “Smart ubiquitous networks – Capability framework,” 2015.
[5] F. Tao et al., “Digital Twin in Industry: State-of-the-Art,” *IEEE Trans. Ind. Informat.*, vol. 15, no. 4, pp. 2405–2415, 2019.
[6] A. Al-Fuqaha et al., “Internet of Things: A Survey on Enabling Technologies,” *IEEE Commun. Surveys Tuts.*, vol. 17, no. 4, pp. 2347–2376, 2015.
[7] C. Dilmegani, “The 7 Layers of Agentic AI Stack in 2026,” *AIMultiple Research*, Jan. 2026.
[8] D. Wang, K. Zhu, B. Song, and X. Du, “Network Data-Knowledge Engine for 6G Endogenous Intelligence,” *IEEE Network*, vol. 39, no. 2, pp. 1–9, 2025.
[9] EXIGENCE Consortium, “Sustainable Cloud and Network Orchestration Architecture,” *EU Horizon Deliverable D2.1*, Oct. 2025.
[10] M. Chen et al., “Cognitive Computing for Internet of Things: A Survey,” *IEEE Internet Things J.*, vol. 9, no. 14, pp. 11220–11237, 2022.
[11] S. Li, L. Da Xu, and S. Zhao, “The Internet of Things: A Survey,” *Information Systems Frontiers*, vol. 17, pp. 243–259, 2015.
[12] A. Bassi et al., “Enabling Things to Talk: IoT-A Architecture,” *Springer*, 2013.
[13] ITU-T Y.4400 Series, “Smart Ubiquitous Networks,” 2015.
[14] Q. Zhang et al., “A Cognitive IoT Architecture with Knowledge Graph,” *IEEE Access*, vol. 9, pp. 57321–57333, 2021.
[15] Y. Li et al., “Digital Twin-Enhanced IoT: Architecture and Applications,” *IEEE Internet Things J.*, vol. 10, no. 8, pp. 6781–6792, 2023.
[16] X. Wang et al., “Agentic AI: Definition, Levels, and Implications for Networking,” *IEEE Commun. Mag.*, vol. 63, no. 5, pp. 34–41, 2026.
[17] Z. Xu et al., “Intent-Based Networking with Large Language Models,” *IEEE Trans. Network Service Management*, vol. 22, no. 3, pp. 1567–1580, 2025.
[18] L. Chen et al., “Multi-Agent DRL for RAN Slicing,” *IEEE J. Sel. Areas Commun.*, vol. 42, no. 1, pp. 89–103, 2024.
[19] R. Sharma et al., “Carbon-Aware Edge Computing,” *IEEE Trans. Green Commun. Netw.*, vol. 8, no. 2, pp. 532–544, 2024.
[20] M. Ali et al., “FTTO for Campus Networks: Energy and Performance,” *J. Opt. Commun. Netw.*, vol. 15, no. 11, pp. 876–887, 2023.
[21] Y. Zhou et al., “LLM Empowered Proactive Defense for Cloud,” *IEEE Commun. Mag.*, vol. 63, no. 8, pp. 112–119, 2025.

[22] K. Patel et al., “Adaptive Zero-Trust with ML,” *Elsevier Comput. Secur.*, vol. 135, 105443, 2025.
[23] J. Singh et al., “Security in 6G Networks: A Survey,” *IEEE Open J. Commun. Soc.*, vol. 5, pp. 3452–3476, 2024.
[24] OPC Foundation, “OPC UA Specification v1.05,” 2024.
[25] Z. Zhang et al., “AI-Native Air Interface for 6G,” *IEEE Wireless Commun.*, vol. 31, no. 3, pp. 12–19, 2024.
[26] N. Jazdi et al., “LLM-Digital Twin Integration for Automation,” *IEEE ETFA 2025*, pp. 1–8, 2025.
[27] A. Sowa, “Anomaly Detection Ensemble for Gigafactory Maintenance,” *SmartFactory Solutions*, Dec. 2025.
[28] Digital Twin Consortium, “TRANSFORM: 2D-to-4D Framework,” *DTC Testbed Report*, Sep. 2025.
[29] P. Lewis et al., “Retrieval-Augmented Generation for Knowledge-Intensive Tasks,” *NeurIPS*, 2020.
[30] Y. Yang et al., “Mean Field Multi-Agent Reinforcement Learning,” *ICML*, 2018.