Telco AI

Autonomous Networks: Transforming Telecomms for the Agentic AI Era

TelcoFutures.net

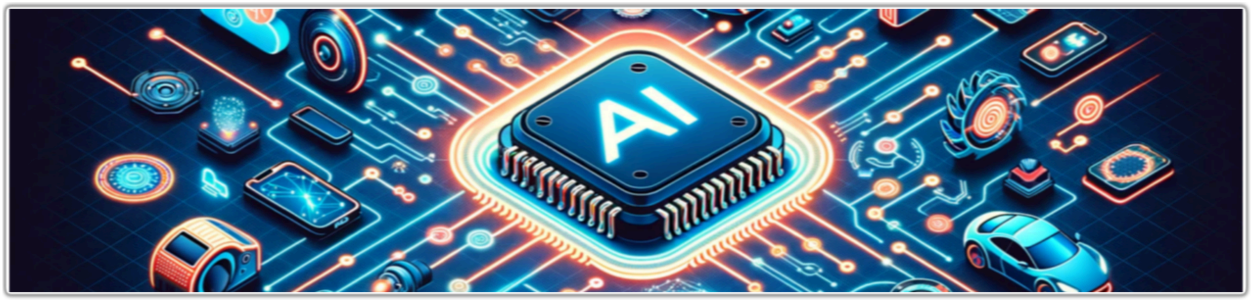

Autonomous Networks:

Transforming Telecomms for the Agentic AI Era

Executive Summary

This report articulates a comprehensive technology strategy and roadmap for the transition to the **Agentic AI Era**. Unlike the preceding wave of Generative AI, which focused on content creation and summarization, Agentic AI introduces autonomous software entities capable of perception, reasoning, planning, and action.

These agents do not merely follow pre-defined scripts; they pursue high-level business intents within a specialized environment, dynamically navigating constraints to achieve outcomes such as zero-touch network slicing, self-healing assurance, and hyper-personalized customer experiences.

1. The Strategic Inflection Point: From Automation to Autonomy

The global telecommunications industry is currently navigating a period of unprecedented architectural entropy. The simultaneous maturation of 5G Standalone (SA) networks, the disaggregation of the Radio Access Network (Open RAN), and the cloudification of core network functions has created an operational environment of staggering complexity.

Traditional operational models, predicated on human-centric management and static, rule-based automation, are rapidly approaching their theoretical limits. As Communication Service Providers (CSPs) face the dual pressures of commoditized connectivity revenues and rising operational expenditures (OpEx), the strategic imperative has shifted from mere automation—doing the same things faster—to *autonomation*, or true autonomous operation.

This report articulates a comprehensive technology strategy and roadmap for the transition to the **Agentic AI Era**. Unlike the preceding wave of Generative AI, which focused on content creation and summarization, Agentic AI introduces autonomous software entities capable of perception, reasoning, planning, and action.

These agents do not merely follow pre-defined scripts; they pursue high-level business intents within a specialized environment, dynamically navigating constraints to achieve outcomes such as zero-touch network slicing, self-healing assurance, and hyper-personalized customer experiences.

1.1 The Complexity Crisis and the Limits of Scripted Automation

For the past decade, the industry's approach to complexity has been the implementation of "runbook automation" or Robotic Process Automation (RPA). These systems function deterministically: *if* event A occurs, *then* execute script B. While effective for repetitive, stable tasks, this paradigm is brittle in the face of the dynamic topology of modern cloud-native networks. A minor change in a vendor's API or a deviation in network latency can cause rigid automation scripts to fail, requiring manual intervention and negating the efficiency gains.

The deployment of 5G and the forthcoming research into 6G introduces requirements that exceed human cognitive reaction times. Use cases such as ultra-reliable low-latency communications (URLLC) for autonomous vehicles or industrial robotics demand network reconfiguration in milliseconds—speeds attainable only by autonomous systems residing at the edge. Furthermore, the sheer volume of managed entities—from massive MIMO arrays to containerized microservices—has exploded, creating a management surface area that cannot be covered by human operators, regardless of team size.

The industry response, led by bodies such as the TM Forum and ETSI, is the vision of the **Autonomous Network (AN)**. This vision targets a progression from Level 2 (Partial Automation) and Level 3 (Conditional Automation), where humans are still heavily involved, to Level 4 (High Autonomy) and Level 5 (Full Autonomy), where the network is self-governing, driven by high-level intent rather than low-level commands.

1.2 The Rise of the Agentic Enterprise

The catalyst for bridging the gap between current automation and future autonomy is Agentic AI. While Generative AI (GenAI) demonstrated the power of Large Language Models (LLMs) to reason over unstructured data, it remained a passive tool—an oracle waiting for a prompt. **Agentic AI** transforms the LLM from a passive reasoning engine into the "brain" of an active agent.

An AI Agent is defined by four core capabilities that distinguish it from a chatbot or a script:

1. **Perception:** The ability to ingest real-time telemetry, logs, and state data from the network environment.
2. **Memory:** The retention of context over time, allowing the agent to learn from past interactions and maintain state across multi-step workflows.
3. **Reasoning & Planning:** The capacity to decompose a high-level goal (e.g., "Optimize coverage in Sector 4") into a sequence of technical steps, adapting the plan if obstacles arise.
4. **Action:** The ability to manipulate the environment through tool use—executing API calls, running diagnostics, or modifying configurations.

By 2030, the leading CSPs will have evolved into "Agentic Enterprises," where the workforce shifts from manual configuration to the orchestration of Multi-Agent Systems

(MAS). In this model, agents are not merely tools but "digital workers" that collaborate, negotiate, and execute complex operations.

Feature	Traditional Automation	Generative AI (2023-2024)	Agentic AI (2025+)
Trigger Mechanism	Explicit command or static rule match.	Prompt-based user query.	Autonomous goal/intent perception.
Workflow Structure	Linear, pre-defined, rigid sequence.	Content generation or summarization.	Dynamic planning, loops, and self-correction.
Adaptability	Low; breaks on unexpected errors.	Moderate; capable of regenerating text.	High; seeks alternative paths to goal.
Primary Outcome	Execution of a specific script.	Information/Content/Code.	Action and Outcome Realization.
Human Role	Executor/Monitor.	Reviewer/Prompter.	Architect/Governor.

Table 1.1: The Evolution from Automation to Agentic AI.

2. The Unified Architectural Framework: ODA and ZSM Convergence

To deploy Agentic AI safely and effectively, CSPs must adopt a robust architectural framework. Placing autonomous agents into a legacy, monolithic OSS/BSS environment is a recipe for chaos. The strategy necessitates the harmonization of two primary industry standards: **TM Forum's Open Digital Architecture (ODA)** and **ETSI's Zero-touch Network and Service Management (ZSM)**. Together, they form the "body" and "nervous system" of the autonomous network.

2.1 ODA: The Environment for Componentized Intelligence

TM Forum's ODA is the structural foundation. It mandates the decoupling of monolithic systems into modular, cloud-native components that expose standardized Open APIs. For the Agentic Era, ODA serves a dual purpose: it provides the *environment* in which agents exist and the *tools* they manipulate.

The ODA Canvas as Agent World

The **ODA Canvas** is the execution environment, typically built on Kubernetes, that hosts ODA components. In an agentic strategy, the Canvas becomes the "world" for the agents. It abstracts the underlying infrastructure complexity, providing agents with standardized methods for discovery, observability, and security.

- **Interoperability:** The Canvas ensures that an "Inventory Agent" can communicate with a "Billing Agent" because both inhabit the same standardized runtime and speak the same API language.
- **Observability:** The Canvas automatically generates the telemetry that "Observer Agents" rely on to perceive the state of the system.

Components as Agent Wrappers

In the traditional view, an ODA Component (e.g., a Product Catalog) is a passive software block. In the Agentic view, we envision **Agentic ODA Components**. Here, the functional block is wrapped or paired with a specialized "Micro-Agent."

- **Example:** A "Service Assurance Component" is no longer just a dashboard. It contains an embedded agent that proactively monitors for SLA breaches and can negotiate with other components to resolve them. This decentralization of intelligence prevents the bottleneck of a single central AI trying to manage the entire telco stack.

2.2 ETSI ZSM: The Governance of Closed Loops

While ODA defines the blocks, ETSI ZSM defines the *flows* and *governance*. ZSM is explicitly designed for closed-loop automation, distinguishing between the specific **Management Domains** (e.g., RAN, Core, Transport) and the overarching **End-to-End (E2E) Service Management**.

Hierarchical Agentic Loops

ZSM's architecture aligns perfectly with a hierarchical Multi-Agent System (MAS).

1. **Domain-Specific Agents:** Residing within the ZSM Management Domains (e.g., a RAN Domain Controller), these agents are highly specialized. They understand the physics of RF (Radio Frequency), beamforming, and local interference. Their autonomy is bounded; they optimize local parameters (Self-Optimization).
2. **E2E Orchestration Agents:** Residing in the ZSM E2E Service Management domain, these agents manage cross-domain services. For example, ensuring a network slice maintains low latency across the RAN, Transport, and Core. They do not micromanage the RAN; instead, they send *Intents* (e.g., "Guarantee 5ms latency for Slice X") to the Domain Agents.

Cross-Domain Coordination

A critical risk in autonomous systems is **loop conflict**. An "Energy Saving Agent" might power down a cell to save electricity (OpEx), while a "Customer Experience Agent" tries to boost power to improve throughput for a VIP user. ZSM specifies **Cross-Domain Coordination** services to mediate these conflicts.

- **Strategy:** Implement a "Governance Agent" or "Arbiter" within the ZSM Integration Fabric. This agent evaluates conflicting actions against high-level business policies (e.g., "During peak hours, Performance > Energy Saving") and grants permission to the winning loop.

2.3 The Unified Reference Architecture

The target architecture for the 2025-2030 horizon fuses these elements into a cohesive stack:

1. **Infrastructure Layer:** Physical and virtual resources (CNFs/VNFs) managed by ZSM Domain Controllers.
2. **Component Layer (ODA):** Modular OSS/BSS functions hosted on the ODA Canvas. Each component is "Agent-Ready," exposing Model Context Protocol (MCP) interfaces or ODA Open APIs.
3. **Agent Fabric:** The communication and discovery layer (A2A Protocol) where agents find each other and negotiate tasks.
4. **Intelligence Layer:** The Multi-Agent System (MAS) comprising Planning Agents, Execution Agents, and Observer Agents. This layer utilizes the **Active Ontology** (SID) for semantic reasoning.
5. **Governance Layer:** The "Trust Sandwich," implementing ZSM governance policies and deterministic guardrails to ensure AI safety.

3. The Cognitive Stack: Technical Deep Dive

Building the Agentic Enterprise requires a specific technical stack that extends beyond standard IT requirements. This section details the protocols, data structures, and system designs required for robust agentic operations.

3.1 Multi-Agent Systems (MAS) Design Patterns

Telecom networks are too large and latency-sensitive for a single "God Agent" (Centralized AI) to manage. The strategy demands a **Multi-Agent System (MAS)** approach, leveraging swarm intelligence and specialization.

Decomposition of Agent Roles

To manage complexity, agent responsibilities must be strictly scoped:

- **Planning Agents (The "Brain"):** Reside in the orchestration layer. They receive high-level business intents (e.g., "Launch V2X service in Munich"). They decompose this intent into a graph of sub-tasks but do not execute them directly.

- **Execution Agents (The "Hands"):** Reside closer to the resources. They receive specific sub-tasks (e.g., "Instantiate UPF in Edge Node 4") and interact with the ODA component APIs to execute them.
- **Observer Agents (The "Eyes"):** Continuously scan network telemetry and logs. They do not act; they only filter noise and alert Planning Agents to anomalies or state changes (e.g., "Link utilization > 80%").
- **Broker Agents (The "Negotiators"):** Manage the economic or resource allocation negotiations between domains (e.g., negotiating the "price" of bandwidth between the Core and Transport domains).

3.2 The Communication Fabric: A Hybrid Protocol Stack

For agents to collaborate, they must share a common language. Traditional REST APIs are insufficient because they support only *commands* ("Do this"), not *negotiation* ("Can you do this? If not, what can you offer?"). The roadmap recommends a hybrid protocol stack.

1. Discovery and Delegation: Agent-to-Agent (A2A) Protocol

The **A2A Protocol** is the modern, cloud-native standard for agent interoperability. It operates over HTTP/JSON and focuses on the lifecycle of tasks.

- **Mechanism:** Agents publish "Agent Cards" (Manifests) describing their capabilities (e.g., "I can configure Cisco routers"). Other agents search the A2A registry to find partners.
- **Usage:** Used for the *logistics* of collaboration—finding an agent, authenticating, and assigning a task.

2. Negotiation Semantics: FIPA ACL

While A2A handles the transport, the *content* of the conversation requires richer semantics. **FIPA ACL (Foundation for Intelligent Physical Agents - Agent Communication Language)** remains the gold standard for defining *communicative acts*.

- **Performatives:** Instead of just POST/GET, agents use performatives like `PROPOSE`, `REFUSE`, `ACCEPT-PROPOSAL`, `CFP` (Call for Proposal).
- **Example:** A Service Agent sends a `CFP` for bandwidth. A Resource Agent replies with a `PROPOSE` (offering 100Mbps at Cost X). The Service Agent sends

a REJECT-PROPOSAL (too expensive), triggering a new round. This negotiation logic is encapsulated within the JSON payload of an A2A message.

3. Tool Binding: Model Context Protocol (MCP)

Agents need to interact with non-agentic systems (databases, legacy OSS). **MCP** provides a standardized way for LLM-based agents to "read" the schema of a tool and understand how to use it without custom code integration.

Protocol	Layer	Role in Strategy
A2A (Agent-to-Agent)	Transport & Discovery	Enables agents to find each other on the ODA Canvas and manage task lifecycles (HTTP/JSON/SSE).
FIPA ACL	Semantic/Logic	Defines the *intent* of the message (Negotiation, Proposal, Refusal) enabling complex reasoning.
MCP (Model Context)	Tooling/Data	Connects agents to static ODA components, databases, and legacy systems.

Table 3.1: The Hybrid Agent Communication Stack.

3.3 The Semantic Core: Active Ontology and SID

Agents cannot reason effectively if the data they consume is ambiguous. If one agent refers to a "Circuit" and another to a "Bearer," collaboration fails.

From Passive SID to Active Ontology

The TM Forum **Shared Information/Data Model (SID)** provides the necessary vocabulary. However, it typically exists as a static UML diagram. The strategy requires implementing SID as an **Active Ontology** or Knowledge Graph.

- **Machine-Readability:** The ontology must be machine-readable (e.g., RDF/OWL) and hosted in a Graph Database.
- **Reasoning:** Agents query this graph to understand relationships. For example, "If I reboot *Device A*, the ontology shows it hosts *Function B*, which supports *Service C*. Therefore, rebooting A impacts Service C." Without this ontological grounding, agents are prone to hallucination.

Ontology Mediation

In multi-vendor environments, disparate ontologies are inevitable. The architecture must include **Ontology Mediation Services**. These services use AI to align concepts between domains (e.g., mapping a Nokia "Cell" to an Ericsson "Sector") in real-time, ensuring semantic interoperability across the ZSM fabric.

4. High-Value Use Cases and Agentic Workflows

The adoption of Agentic AI must be driven by use cases that deliver measurable ROI, specifically where traditional automation fails due to complexity or latency.

4.1 Dynamic Network Slicing Negotiation (The "Contract Net" Pattern)

Static network slicing is inefficient; it locks up resources regardless of usage. Agentic AI enables **Dynamic Micro-Slicing**, where resources are negotiated in real-time.

Scenario: A stadium requires a temporary, high-bandwidth slice for a 30-minute VR event.

Workflow:

1. **Request:** The "Stadium Enterprise Agent" broadcasts a **Call for Proposal (CFP)** via the A2A protocol: `CFP: {Bandwidth: 10Gbps, Latency: <5ms, Duration: 30min}`.
2. **Feasibility Check:** "RAN Agents" and "Core Agents" receive the CFP. They query their internal state and Digital Twins to predict if they can support the load

without violating existing SLAs.

3. **Bidding:**
   - *RAN Agent A* replies: REFUSE (Congested).
   - *RAN Agent B* replies: PROPOSE {Cost: $500, QoS: Gold}.
4. **Selection:** The Enterprise Agent evaluates the proposals using a utility function (Price vs. Performance) and sends an ACCEPT-PROPOSAL to Agent B.
5. **Instantiation:** The accepted RAN Agent triggers the ZSM provisioning workflow to spin up the slice resources instantly.
6. **Termination:** Upon contract expiry, the agents auto-negotiate the teardown and billing settlement.

Impact: Monetization of transient opportunities and maximal resource utilization efficiency.

4.2 Autonomous Service Assurance and Self-Healing

Moving beyond "Root Cause Analysis" (RCA) to "Root Cause Resolution" (RCR).

Scenario: A degradation in video streaming quality is detected in a metro area.

Workflow:

1. **Perception:** "Observer Agents" detect a spike in re-buffering events. They correlate this with a "Link Down" alarm in the Transport domain.
2. **Reasoning:** A "Healer Agent" analyzes the topology graph (Active Ontology). It identifies that the primary fiber route is cut.
3. **Planning:** The agent generates a plan: "Reroute traffic via Microwave Backup."
4. **Simulation:** Before acting, the agent simulates the reroute in the **Network Digital Twin**. It discovers that the microwave link will be congested.
5. **Refinement:** The agent refines the plan: "Reroute *Gold* traffic to Microwave; drop *Best Effort* traffic." It simulates again—Success.
6. **Action:** The agent executes the configuration change via ODA APIs and simultaneously opens a ticket for the physical fiber repair.

Impact: Prevention of major outages and reduction of MTTR from hours to seconds.

4.3 Proactive Customer Concierge

Transforming customer support from a reactive cost center to a proactive value driver.

Scenario: A known outage affects a residential area.

Workflow:

1. **Trigger:** The "Network Operations Agent" confirms a sector outage and estimates a 4-hour repair time. It publishes this event to the A2A fabric.
2. **Proactive Care:** "Customer Concierge Agents" (assigned to affected subscribers) consume this event. They immediately check the customers' profiles.
3. **Action:** The Concierge Agents proactively message the customers via WhatsApp/SMS: *"We've detected a network issue in your area. Our engineers are fixing it (ETA: 4 hours). We have automatically credited 10GB of mobile data to your account so you can tether your devices in the meantime."*
4. **Resolution:** Once the Network Agent confirms the fix, the Concierge Agents notify the customers and ask for feedback.

Impact: massive reduction in call center volume (OpEx) and improved Net Promoter Score (NPS).

5. Strategic Roadmap 2025-2030

The transformation to an Agentic Enterprise is a multi-year journey. This roadmap defines three horizons to manage risk and maturity.

Horizon 1: Augmented Intelligence & Co-Pilots (2025)

- **Focus: "Human-in-the-Loop."** Agents act as assistants to human operators, reducing cognitive load but having no autonomous write access to the network.
- **Key Initiatives:**
  - **NOC Co-Pilots:** Deploying GenAI agents in the Network Operations Center to summarize alarm floods, query knowledge bases (RAG), and draft RCA reports for human review.
  - **ODA Canvas Pilot:** Establishing the ODA Canvas for non-critical OSS applications to test the environment.
  - **Data Foundation:** Building the "Data Nerve Center"—unifying telemetry and inventory data into a format accessible by agents (MCP).

    - **Ontology Setup:** Initial implementation of the Telco Ontology to map key domains (RAN, Core).
- **Target Autonomy:** TM Forum Level 2/3 (Conditional Automation).

Horizon 2: Domain Autonomy & Policy-as-Code (2026-2028)

- **Focus: "Human-on-the-Loop."** Agents execute tasks autonomously within strictly bounded domains. Humans focus on setting policies and handling exceptions.
- **Key Initiatives:**
    - **Single-Domain Agents:** Deploying autonomous agents for specific tasks like "RAN Energy Saving" or "Transport Route Optimization." These agents can sense and act (closed-loop) but only within their silo.
    - **Agentic Workflows:** Implementing the A2A protocol to allow limited interaction between agents (e.g., Customer Support Agent querying Network Status Agent).
    - **Policy Governance:** Implementing the ZSM Governance framework. Agents must pass "deterministic guardrails" (code-based safety checks) before execution.
    - **Digital Twin Integration:** Mandating that all autonomous changes be simulated in a Digital Twin before application.
- **Target Autonomy:** TM Forum Level 4 (High Autonomy in specific domains).

Horizon 3: Cross-Domain Ecosystems (2029-2030)

- **Focus: "Human-out-of-the-Loop" (Routine).** Swarm intelligence manages end-to-end services. Humans act as strategic architects.
- **Key Initiatives:**
    - **Cross-Domain Negotiation:** Full implementation of FIPA ACL negotiation. Agents dynamically trade resources across RAN, Core, and Cloud domains to optimize E2E SLAs.
    - **Self-Evolving Networks:** Agents collaborate to design new network topologies and service products. The network "re-programs" itself based on long-term demand trends.
    - **Agent Marketplace:** Opening the A2A fabric to third parties. An "Autonomous Vehicle Fleet Agent" from a logistics company can negotiate

connectivity directly with the Telco's agents.

- **Target Autonomy:** TM Forum Level 5 (Full Autonomy).

6. Governance, Trust, and Security

The introduction of non-deterministic AI agents into deterministic network infrastructure introduces significant risk. A hallucinating agent cannot be permitted to misconfigure a core router.

6.1 The "Trust Sandwich" Architecture

To mitigate risk, we propose a governance architecture known as the **Trust Sandwich**:

1. **Top Layer (Intent):** The human operator defines the high-level Intent and Policy constraints (e.g., "Optimize for Power, but never drop below 98% coverage").
2. **Middle Layer (Agent):** The Agent perceives, reasons, and proposes a plan of action.
3. **Bottom Layer (Deterministic Guardrails):** A non-AI, code-based Policy Engine (e.g., OPA - Open Policy Agent) validates the proposed action against strict safety rules.
    - *Rule:* "Configuration changes must not affect Emergency Services interfaces."
    - *Rule:* "Power transmission cannot exceed X dBm."
    - Only if the action passes these deterministic checks is it executed.

6.2 Explainability and Auditability

Every action taken by an agent must be auditable.

- **Chain of Thought Logging:** Agents must be architected to log their internal reasoning steps (e.g., "I detected congestion at Node A. I considered Reroute and Throttle. I chose Reroute because Policy B prefers availability over speed").
- **Trust Index:** Agents should operate on a "earned autonomy" model. A new agent starts in "Recommendation Mode" (human approves all actions). As its **Trust Index** (success rate) improves, it graduates to "Supervised Autonomy" and finally "Full Autonomy".

7. Workforce and Organizational Transformation

The shift to the Agentic Enterprise is as much cultural as it is technological. The skills required to manage a fleet of autonomous agents are fundamentally different from those required to configure routers via Command Line Interface (CLI).

7.1 The Skills Gap and Reskilling

Current network engineers are experts in vendor-specific syntax and protocols (BGP, OSPF). Future engineers must be experts in **AI Orchestration** and **Data Modeling**.

- **Impact:** Research suggests that while AI will augment productivity, it will displace roles focused on repetitive configuration. However, it creates a demand for high-level problem solvers who can manage the AI.
- **Reskilling Strategy:** CSPs must invest in "AI Literacy" programs. Engineers must learn **Prompt Engineering**, **Ontology Design**, and **Logic Governance**.

7.2 Emerging Roles

New job families will emerge to support the Agentic stack:

- **Agent Architect:** Responsible for designing the "persona," goals, and tool-access permissions of an agent. They define *what* the agent should try to achieve.
- **Ontology Engineer:** Responsible for maintaining the Active Ontology (SID). They ensure that the data model accurately reflects the physical network, preventing agent confusion.
- **AI Governance Officer:** A compliance role responsible for auditing agent behaviors, ensuring they adhere to ethical guidelines and safety regulations.

7.3 Cultural Shift: From CLI to Intent

The organization must move from a culture of "How" (scripting the exact steps) to a culture of "What" (defining the desired outcome). This requires a shift in mindset where operators trust the system to find the optimal path, focusing their efforts on defining the boundaries and goals rather than the execution mechanics.

Conclusion

The convergence of ODA, ZSM, and Agentic AI represents the definitive architecture for the future of telecommunications. This is not merely an incremental upgrade but a transition to a "Cognitive Network"—an organism that perceives, reasons, and adapts.

By 2030, the telecom operator will no longer be a provider of static pipes but an orchestrator of intelligent, adaptive resources. The "Agentic Enterprise" will be characterized by fluid, self-optimizing infrastructure, hyper-personalized customer experiences, and a workforce elevated to the role of strategic architects. The roadmap presented here—moving from Co-pilots to Domain Autonomy to Cross-Domain Ecosystems—provides the structured path to realize this vision. The risk of inaction is obsolescence; the reward for execution is the unlocking of the true value of the programmable network.

Detailed Strategic Analysis & Reference Data

Appendix A: Protocol Comparison for Agentic Networks

Feature	REST/OpenAPI (Legacy)	A2A Protocol (Modern)	FIPA ACL (Semantic)	Strategic Recommendation
Primary Function	Command execution ("Do X").	Discovery & Lifecycle.	Negotiation ("Discuss X").	**Hybrid Stack**: Use A2A for discovery/transport, carrying FIPA payloads for logic.

Interaction Style	Request/Response (Synchronous).	Async, Event-driven (SSE).	Asynchronous, Multi-turn.	Shift to Async to support long-running agent reasoning tasks.
Flexibility	Low; requires rigid schema.	High; adaptable manifests.	High; content-agnostic.	Essential for multi-vendor environments where schemas differ.
Key Artifact	Swagger/OAS file.	Agent Card/Manifest.	Performative (Propose/Refuse).	Develop standardized "Agent Cards" for all ODA components.
Status	Ubiquitous.	Emerging Standard.	Mature Standard.	Pilot A2A in 2025; adopt FIPA semantics for Slicing in 2026.

Table A.1: Strategic comparison of communication protocols.

Appendix B: Value Realization Projections

Metric	2024 Status	2026 Target (Horizon 2)	2030 Target (Horizon 3)	Driver

	(Baseline)			
Agentic AI Adoption	< 1% of apps	15% of apps	> 80% of apps	Maturity of ODA Canvas & A2A.
Decision Autonomy	Human-driven	15% Autonomous	40-50% Autonomous	Trust in ZSM Governance Framework.
Network Maturity	TM Forum L2/3	TM Forum L3/4	TM Forum L5	Deployment of MAS in RAN/Core.
MTTR (Repair)	Hours (Manual)	Minutes (Co-pilot)	Seconds (Self-healing)	Automated RCA & Closed-loop action.
OpEx Reduction	-	10-15%	> 40%	Removal of manual configuration tasks.
Key Technology	GenAI (Chatbots)	Single-Domain Agents	Swarm Intelligence	Evolution of AI reasoning models.

Table B.1: Projected milestones and value KPIs for the Agentic Network transformation.