White Paper June 2025

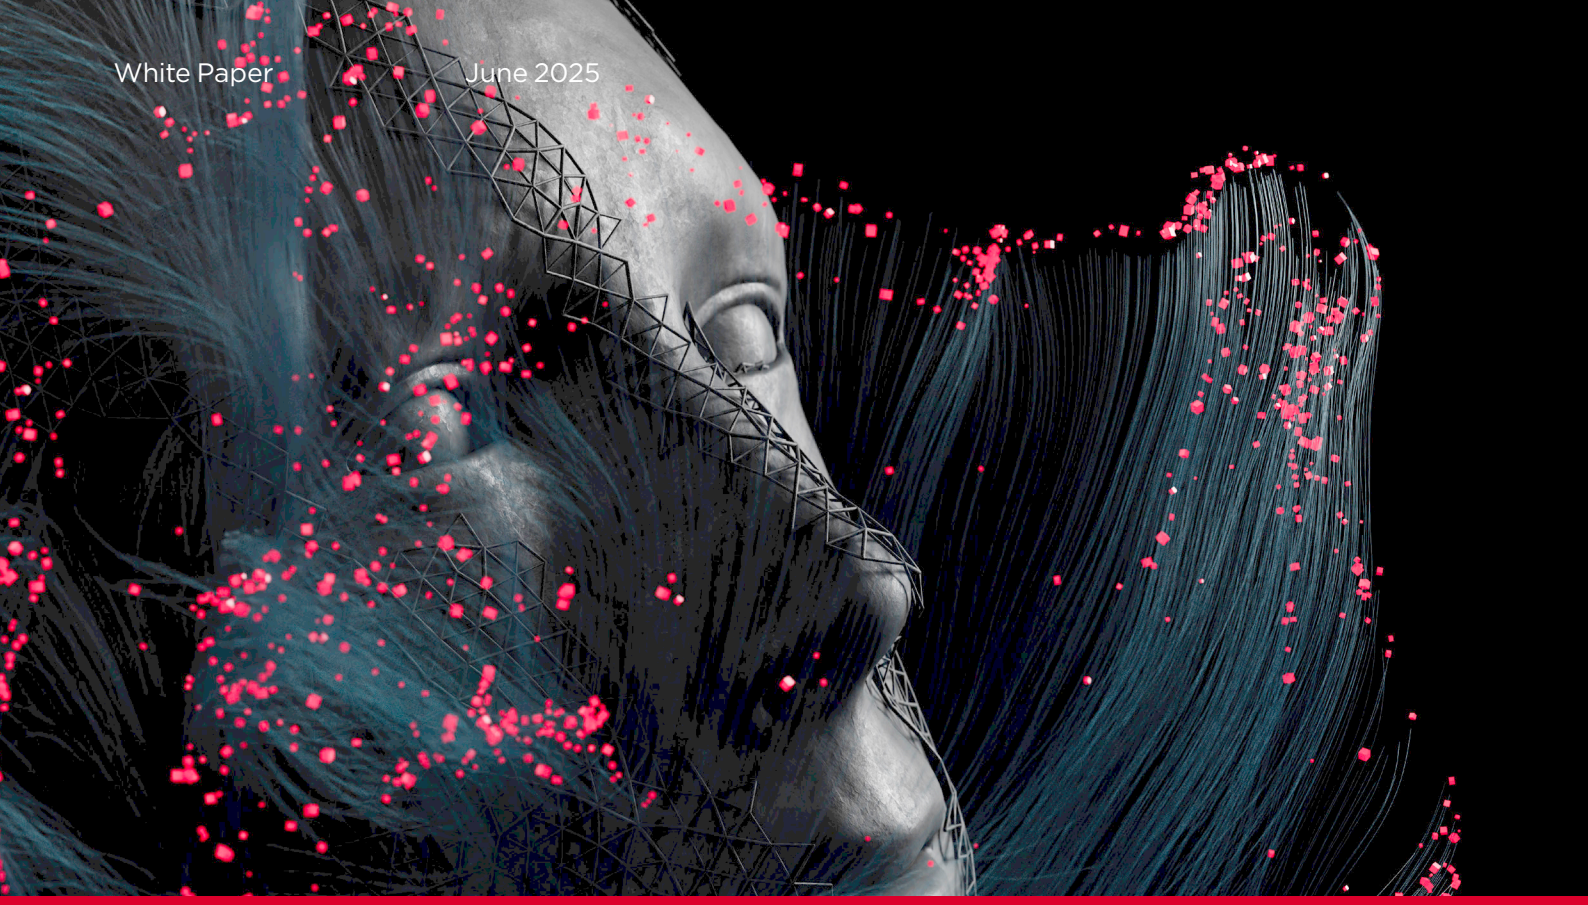

Agentic AI for Telecom: Charting the Course for an Intelligent Future

GSMA™

About the GSMA

The GSMA is a global organisation unifying the mobile ecosystem to unlock the full power of connectivity so that people, industry and society thrive.

Led by our members, we represent the interests of over 1,100 operators and businesses in the broader ecosystem. The GSMA also unities the industry at world-leading events, such as MWC (in Barcelona, Kigali, Las Vegas and Shanghai) and the M360 Series.

Unlock the benefits of GSMA membership

As a member of the GSMA, you join a vibrant community of industry leaders and visionaries – helping to shape the future of mobile technology and its transformative impact on societies worldwide.

Our unique position at the heart of the mobile industry means you get exclusive access to our technical experts, data and analysis – as well as unrivalled opportunities for networking, innovation support and skills acceleration.

For more information, please visit:
http://www.gsma.com/membership/

Security Classification: Non-confidential

Access to and distribution of this document is restricted to the persons permitted by the security classification. This document is confidential to the Association and is subject to copyright protection. This document is to be used only for the purposes for which it has been supplied and information contained in it must not be disclosed or in any other way made available, in whole or in part, to persons other than those permitted under the security classification without the prior written approval of the Association.

Antitrust Notice

The information contained herein is in full compliance with the GSM Association's antitrust compliance policy.

Copyright Notice

Disclaimer

The GSM Association ("Association") makes no representation, warranty or undertaking (express or implied) with respect to and does not accept any responsibility for, and hereby disclaims liability for the accuracy or completeness or timeliness of the information contained in this document. The information contained in this document may be subject to change without prior notice.

Contents

Foreword

The telecommunications industry stands at the cusp of another profound transformation, driven by the rapid advancements in Agentic Artificial Intelligence. This is not merely an incremental change; it represents a paradigm shift towards networks and services that are inherently autonomous, adaptive, and deeply integrated into the fabric of our digital lives. For Communication Service Providers, Agentic AI unlocks unprecedented opportunities to create new value, enhance customer experiences, and drive operational excellence.

However, seizing these opportunities requires a clear vision, strategic foresight, and, most critically, industry-wide collaboration. The path to a thriving Agentic AI ecosystem is one we must forge together. This white paper serves as a vital contribution to that journey, outlining the strategic landscape, key technological enablers, and the pressing challenges we face. It also issues a crucial call for alignment on standardised frameworks, essential for fostering innovation and ensuring a cohesive, interoperable future.

GSMA is committed to facilitating this dialogue and spearheading collaborative initiatives that will empower our members to lead in the Agentic AI era. I encourage you to engage with the insights and proposals within these pages as we collectively chart the course for an intelligent, connected world.

Louis Powell
Director of AI Initiatives, GSMA

1. Agentic AI for Telecom: Charting the Course for an Intelligent Future

Executive Framework

Strategic Positioning

Agentic Artificial Intelligence (Agentic AI) is emerging as a pivotal direction in global technological evolution and market development. The core strength of Agentic AI lies in its autonomous decision-making capabilities and automated execution features that transcend the creative functions of generative AI (Gen AI), as it directly integrates into enterprise workflows to deliver quantifiable business value.

Communication Service Providers (CSPs) that secure early access to infrastructure and industry scenario gateways are poised to become key players in the rapidly expanding Agentic AI market, projected by some analysts to reach trillion-dollar valuations. This potential naturally leads to a critical question for the industry: What concrete steps must CSPs take to achieve this pivotal role?

With the rapid evolution of generative AI, artificial intelligence is now transitioning into the era of Agentic AI, characterised by autonomy and intent-driven decision-making. Unlike traditional AI limited to predefined tasks, Agentic AI leverages workflows that decompose complex goals, iteratively optimise actions, and actively adapt to dynamic environments, positioning itself as the cornerstone of next-generation digital infrastructure. This technological leap addresses efficiency bottlenecks in complex scenarios and fosters cross-industry applications through open ecosystems, shifting AI from a 'tool' to a 'collaborative partner'. Agentic AI will deeply integrate into digital infrastructure, driving revenue productivity, resource efficiency, and high-level autonomy.

CSPs are poised to play a pivotal role in AI-native networks by embedding Agentic AI throughout their infrastructure. By leveraging end-to-end control and real-time data, they enable autonomous decision-making, predictive fault management, and contextual services. This deep integration unlocks adaptive, resilient, and user-centric communication systems, driving innovation and growth in telecommunications.

Given this transformative potential, how can CSPs strategically position themselves to invest in and derive maximum value from the Agentic AI revolution?

Realising this potential and effectively navigating the Agentic AI era cannot be a solitary endeavour for any single CSP. The development of a robust, interoperable, and trusted Agentic AI ecosystem necessitates a unified approach. Crucially, cross-market collaboration

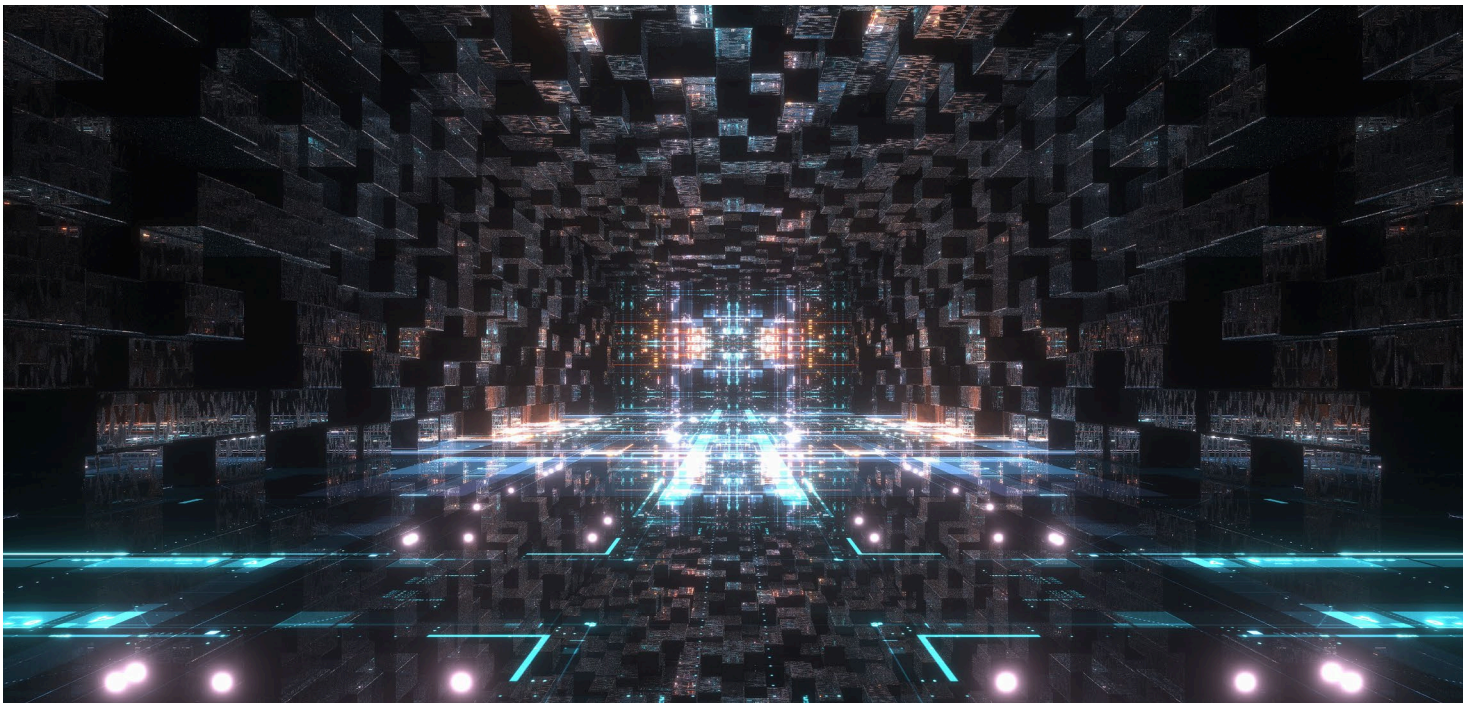

between CSPs is essential to develop a standardised framework that can support this new ecosystem. Such a framework, built on co-operation, will foster innovation, ensure seamless service integration, and provide the foundation for the sector to thrive, to the long-term benefit of businesses, consumers and society more broadly.

This white paper explores the vision and challenges for CSPs to re-position themselves in the Agentic AI era, outlining three strategic tiers for consideration:

1. **Tool-oriented (Conservative):** Expose traditional deterministic services via Application Programming Interfaces (APIs).
2. **Service-oriented (Proactive):** Transform core services with Large Language Models (LLMs) into proactive agents (e.g., AI customer service**).**
3. **Ecosystem-oriented (Radical):** Build customer assistants as primary gateways, integrating internal/external services (e.g., cross-platform AI agents)

Acknowledgements

Leads

— Lingli Deng, China Mobile
— Soonmin Bae, KT
— Jerry Jae Yeol Kim, LG U+
— Gabriele Elia, TIM

Contributors:

— Gyuseob Lee, KT
— Carlo Licciardi, TIM
— Tangqing Liu, [Affiliation]
— Takeshi Kato, NTT DOCOMO
— Issei Nakamura, NTT DOCOMO
— Idil Cilbir, Turkcell
— Chi Ren, China Unicom

Contributing Organisations:

— China Mobile
— China Unicom
— KT
— LG U+
— NTT DOCOMO
— TIM
— Turkcell

2. Industry Transformation Context

2.1. Paradigm Shift Analysis

Evolution from reactive AI to agentic systems

Traditional Reactive AI, constrained by its passive response to predefined instructions, is being surpassed by Agentic AI, which achieves proactive decision-making and complex goal fulfillment through intent-driven task decomposition and iterative optimisation. This shift propels enterprise IT from a 'tool-assisted' paradigm toward 'autonomous collaboration'

Looking ahead, CSPs aiming to become effective agentic service providers must increasingly build modular technical stacks and federated governance models to balance innovation with risk mitigation. Simultaneously, cultivating cross-disciplinary talent will be critical to navigating the organisational transformations driven by agentic systems.

Security by design in agentic systems

Autonomous agent collaboration demands enhanced security. However, traditional industry best practices don't always produce the best results in this new landscape. Anyone adopting agentic systems into any workflow must adopt the concept of the threat actor as internal to the workflow, rather than as an external party interacting with it.

For this purpose, we will introduce 3 main definitions:

1. Boundary Collapse, where an agentic model's output is executed by tooling that treats the model as authoritative;
2. Supply Chain Substitution, tampered dependencies, images, or model weights; and
3. Privilege Misconfiguration, overly permissive service accounts that let an agent or user facing tool modify the model scheduler itself, or any sensitive data crossing a boundary from deterministic permissions to probabilistic.

These elements will be expanded in section 6.2

Research Note (GSMA & 3GPP on Security):

- **GSMA:** GSMA has undertaken extensive work on IoT security, network security (NESAS), and security for 5G. Many principles like zero-trust, security by design, and multi-layered defence are echoed in GSMA guidelines (e.g., FS.31 IoT Security Guidelines, FS.07 GSMA 5G Cybersecurity Knowledge Base). The concept of "increasing attack costs" is a well-understood security principle. This paragraph aligns well with established GSMA security stances.
- **3GPP:** 3GPP security specifications (e.g., TS 33.xxx series) for 5G and beyond heavily emphasise security by design, authentication, authorisation, and integrity

5G-Advanced/6G network readiness for AI agent ecosystems

The 3rd Generation Partnership Project (3GPP) standards integrate AI agents into next-generation networks. 3GPP's Service Requirements working group (SA1) defines them as 'autonomous entities that interact with environments, reason, and execute tasks (e.g., dynamic resource allocation)'. 5G-Advanced (Release 18) enables agent operations via AI automation, enhanced slicing, and edge computing, while 6G natively embeds intelligence. This evolution transitions Agentic AI from external tools to intrinsic network components.

3. Telco-Centric AI Agent Opportunities

The emergence of the Agentic AI Era presents significant opportunities for CSPs. To capitalise on these, CSPs can act with agility and adopt one or a combination of the following three strategic positioning tiers within the Agentic AI industrial ecosystem:

Tool-Oriented Approach (Conservative Strategy)

Maintain existing service frameworks while exposing capabilities via standardised interfaces (APIs/Software Development Kits (SDKs)), packaging existing Network as a Service (NaaS) APIs as Agent-invokable tools to serve as foundational modules in the Agent ecosystem.

Agent-Oriented Approach (Proactive Strategy)

Revolutionise core services and network infrastructures through deep LLM integration, creating proactive cognitive Agentic systems for service delivery and network operations. Examples fall into two main categories, detailed further in Section 3 (Telco-Specific Implementation Scenarios):

- For Agentic Service Evolution:
  - **Omnipresent Life Assistant:** For individual consumers, leveraging cloud phones and home compute hosts as key enablers to deliver seamless AI service functionalities from wide-area mobile scenarios to home environments.
  - **All-Capable Smart Home Hub:** For household subscribers, unifying smart home devices, connectivity, and content services into intent-driven experiences.
- For Agentic Network Evolution:
  - **Proactive Customer Care:** Which analyses network Quality of Service (QoS), device status, and user behaviour to pre-emptively resolve service issues before customer complaints arise.
  - **Dynamic SLA Management:** Where multi-agent systems autonomously adjust SLAs based on real-time conditions.
  - **Personalised Connectivity:** Where a customised user experience is provided via a generated network.

Ecosystem-Oriented Approach (Radical Strategy

This approach involves providing general-purpose customer agents that build cognitive moats through continuous customer habit learning. These agents also serve as the primary entrance to other agents and tools, dynamically integrating proprietary offerings (e.g., connectivity/content ecosystems) and third-party capabilities (e.g., delivery/transportation). Such a strategy dictates the construction and maintenance of a Telco-centred multi-agent collaboration ecosystem.

4. Telco-Specific Implementation Scenarios

4.1 Multimodal Communication Suite

4.1.1 Context-aware IP Multimedia Subsystem (IMS) assistant with multi-modal interaction.

Combined with Gen AI, spatial computing, and split rendering capabilities, the enhanced IMS network can provide a multi-dimensional immersive experience for users. For example, during a call, the enhanced IMS network can help convert 2D to 3D and display content through Augmented Reality (AR) glasses. With these capabilities, the enhanced IMS network can also provide immersive shared spaces, virtual-real navigation, and collaborative design capabilities.

To handle context-rich multi-modal interactions, IMS must go beyond traditional session control to understand user intent, emotional state, device conditions, and environmental context, and intelligently route sessions to the appropriate AI agents through AI-based orchestration. This enables the following capabilities:

- Simultaneously manage a massive number of sessions driven by the proliferation of AI-powered smart devices.
- Synchronise interactions across modalities, including voice, text, and video.
- Dynamically allocate AI agents (e.g., translation, summarisation, emotion analysis) based on contextual information.

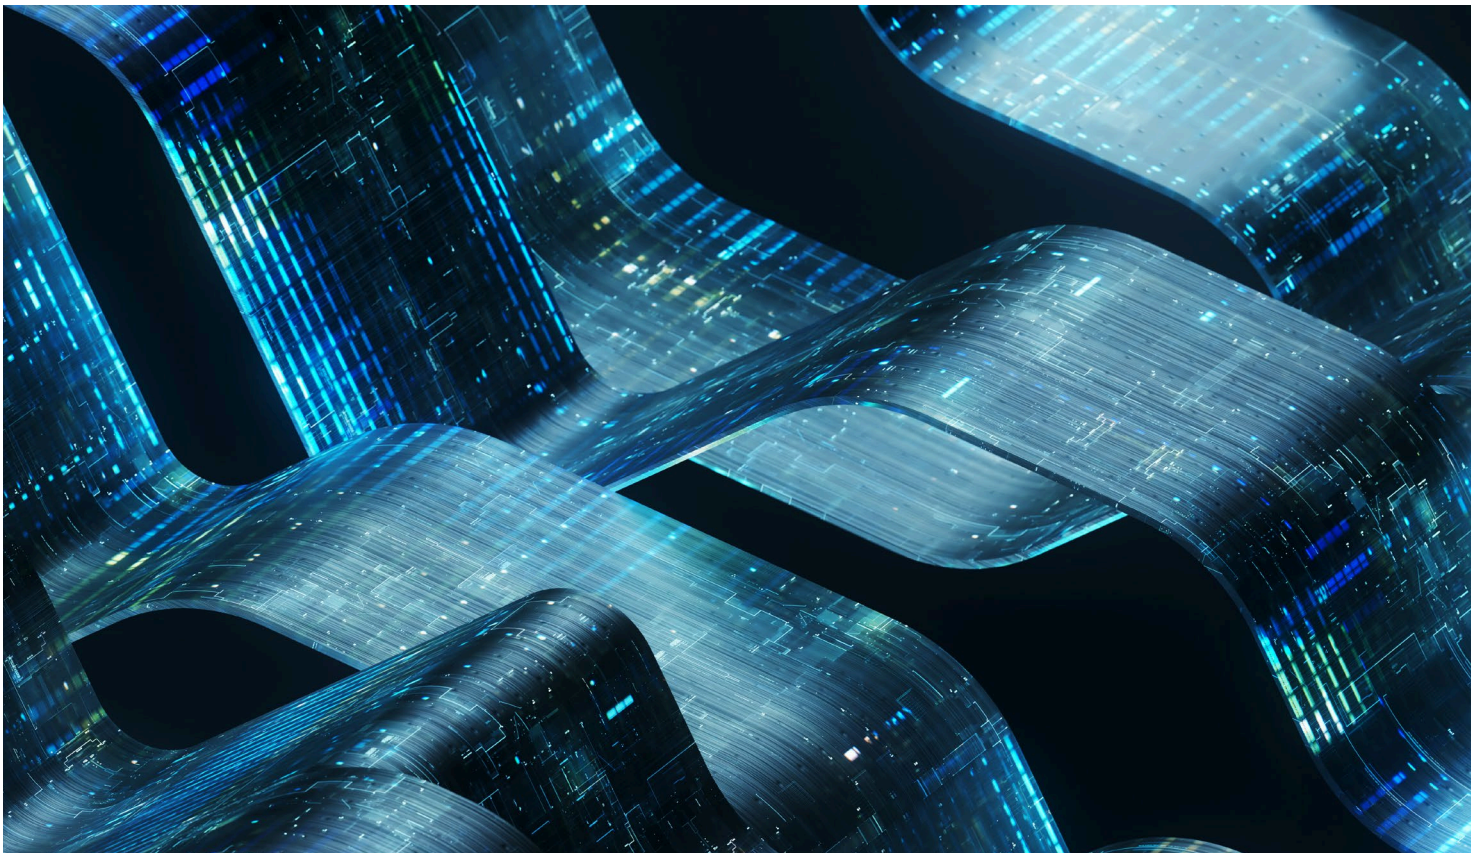

— Enhance security through AI-based threat detection, behaviour-based authentication, and advanced session encryption.

With these capabilities, IMS evolves from a traditional communication infrastructure into an intelligent service hub that supports personalised and adaptive multi-modal interaction experiences.

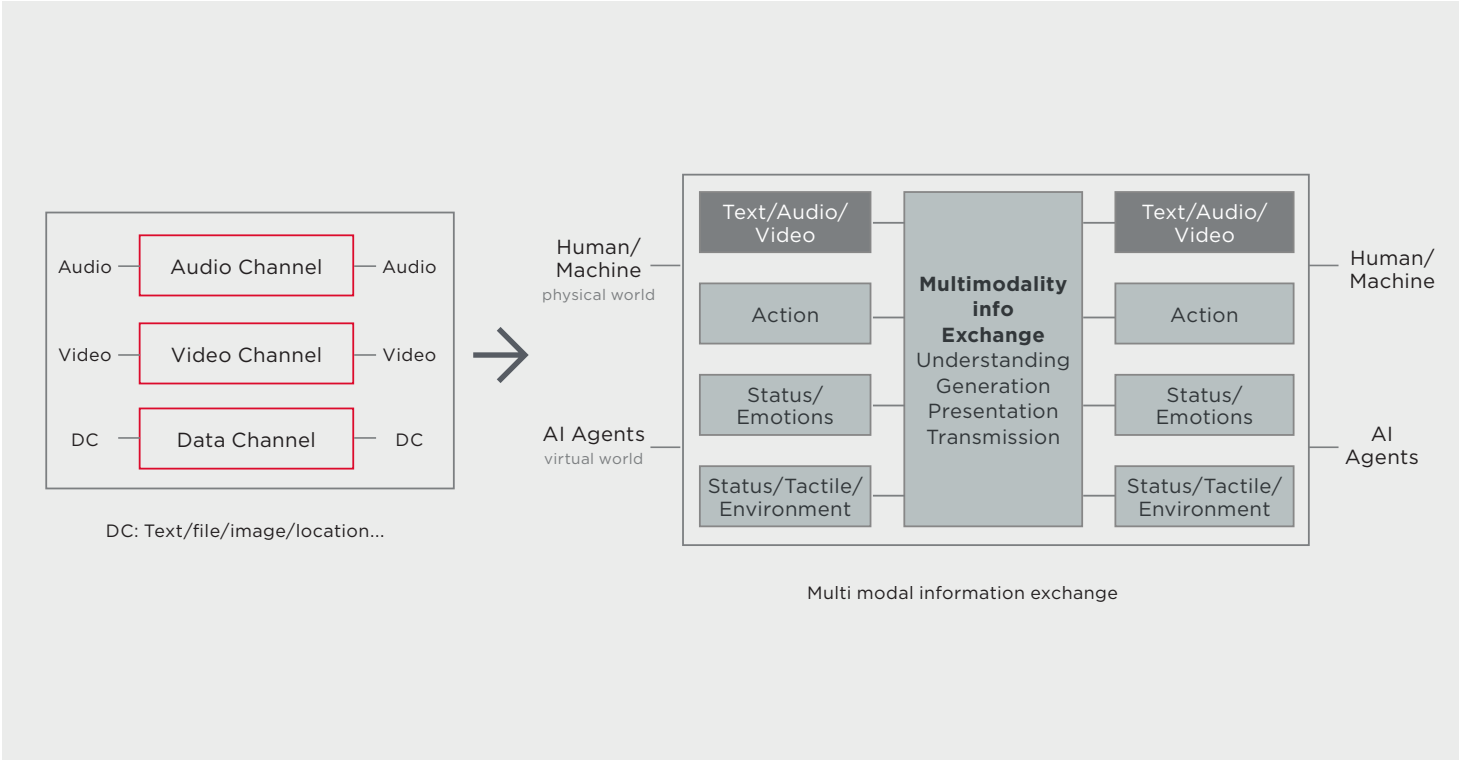

Figure 1, Enhanced IMS network to support multi modal communications

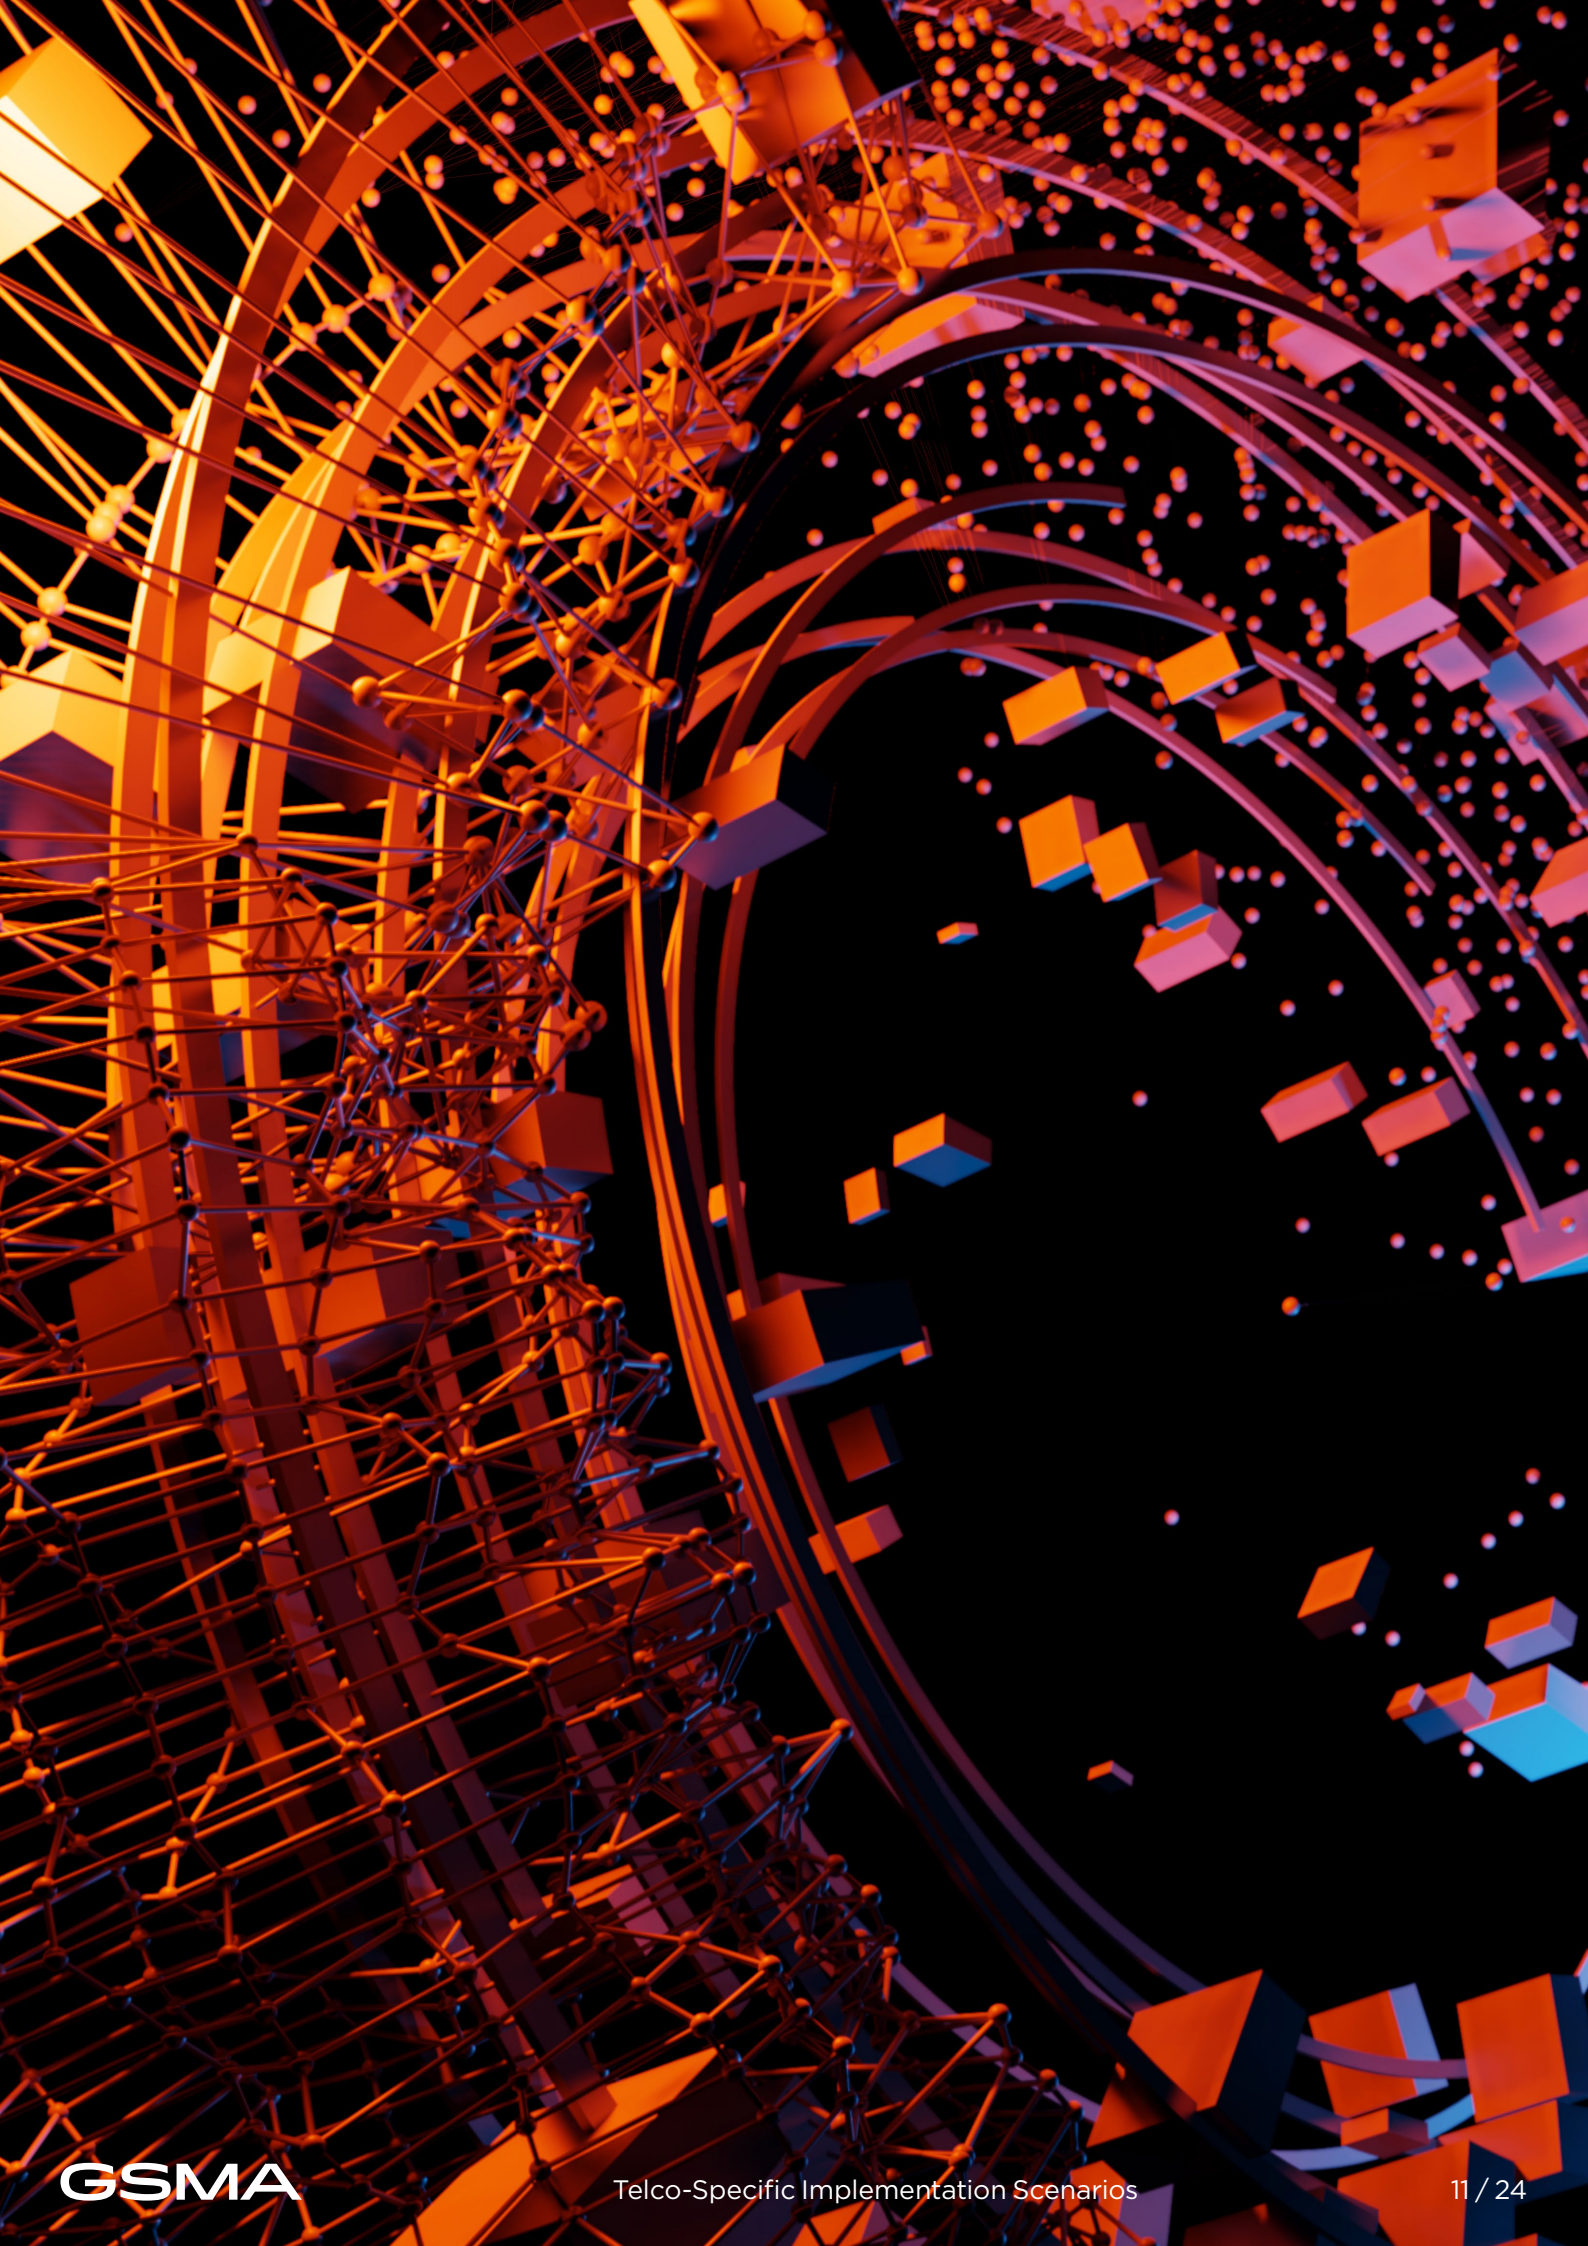

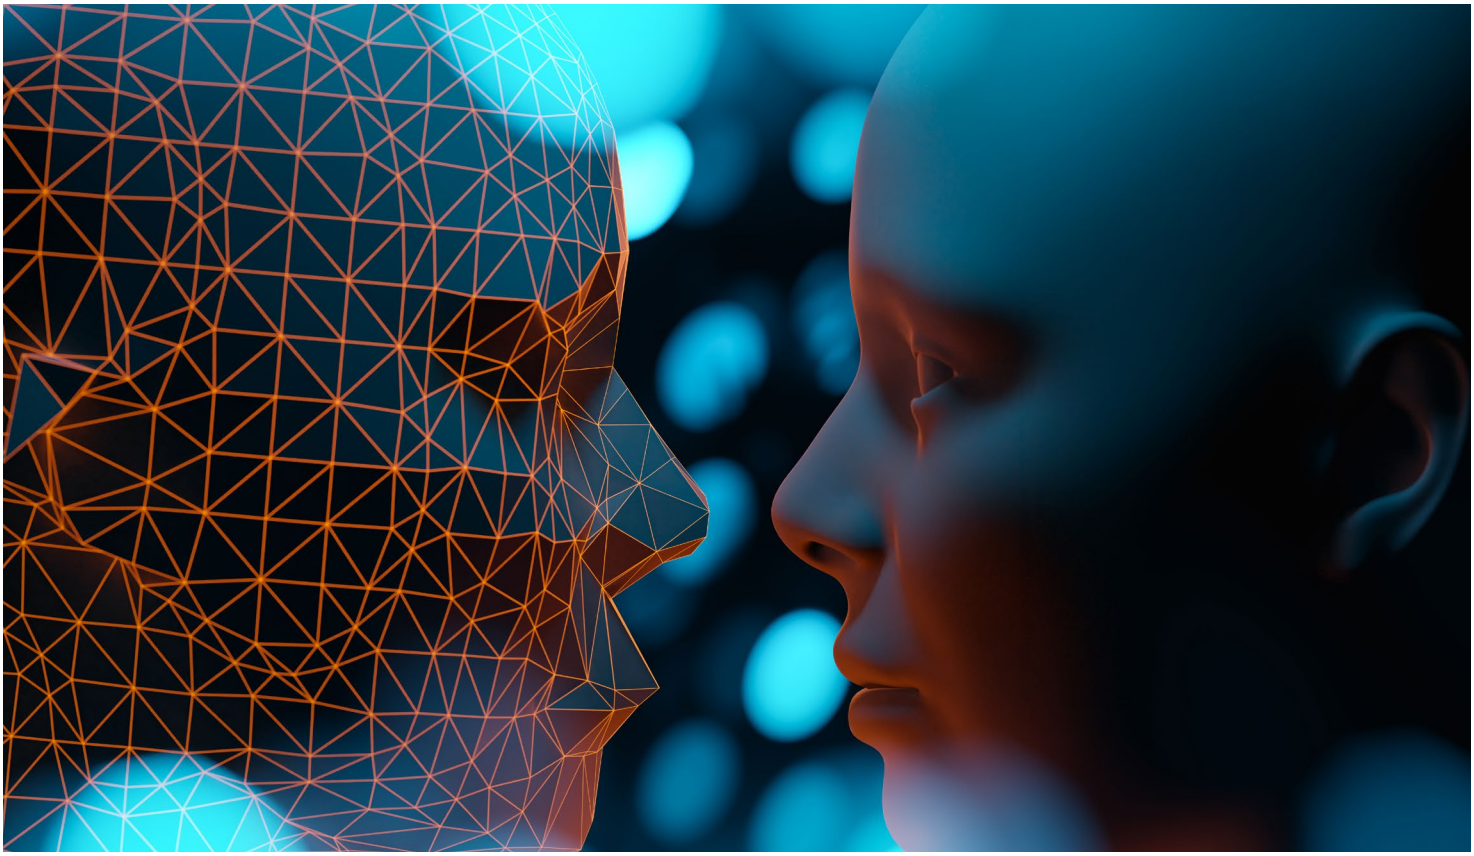

4.1.2 Interoperable 3rd party Agent integration ecosystem

As AI-based services continue to diversify, enterprises and CSPs are experiencing growing demand to seamlessly integrate specialised third-party AI a gents for translation, summarisation, customer support, voice analytics, and more. To meet this need, an open integration ecosystem is required, enabling AI agents from various vendors and technology stacks to work together within IMS, communication platforms, and customer service environments through standardised and interoperable methods.

The key components of this ecosystem include:

- **Agent Communication Protocol:** To establish common formats for intent, emotion, and action exchange across heterogeneous agents.
- **Agent Integration Gateway:** Ensures interoperability by mediating between various interfaces.
- **Agent Orchestration Engine:** Dynamically assigns tasks, synchronises sessions, and manages agent priority based on intent and context.

By adopting this ecosystem, CSPs can avoid vendor lock-in and leverage a wide range of AI capabilities in a flexible and scalable manner, delivering enriched, personalised user experiences. Ultimately, this approach enhances the agility, extensibility, and long-term sustainability of AI-driven services.

4.2 Home AI Agents: Smart Home Assistant

4.2.1 AI Agents at Home: Smart Home Assistant

In the era of Agentic AI, the home represents an important customer touchpoint for CSPs, offering distinct opportunities from several key perspectives. Firstly, from the users' perspective, the home is where they spend a considerable portion of their day, engaging in entertainment and performing numerous tasks involving home appliances, computers, media devices, and Internet of Things (IoT) gadgets, all interconnected through various forms of connectivity. Secondly, the home provides a comfortable environment that encourages users to readily express their intent, making it easier to collect valuable personal data and intent signals. Thirdly, from the network providers' perspective, the home is an established territory where they deliver strong and reliable fixed and/or mobile network access supported by customer premise equipment (CPE) like access points and set-top boxes. This existing infrastructure provides substantial leverage, allowing home AI agents to naturally position as key enablers in delivering AI-driven capabilities and services that span the full spectrum of computational requirements—from low-power on-device processing to high-performance edge and cloud-based AI compute.

4.2.2 Home AI Agent: Scenarios

CSPs can leverage cloud phones and home compute hosts as service entry points to deliver AI service functions to end-users. Through multi-tool and multi-agent collaboration, closed-loop execution of user intents can be achieved. Cloud phones and home compute hosts cover both consumer (B2C) and household scenarios, delivering consistent AI experiences across user environments.

In the context of household and personal life needs, CSPs can utilise AI models to create intelligent agent assistants, delivering more convenient, seamless, and enriched smart living experiences for users. Users can interact with the intelligent assistant via voice commands - for example, by saying, "Help me plan my weekend activities." The intelligent assistant performs user intent recognition and leverages historical user data to formulate a task plan. The user's historical information can help the intelligent assistant to build a unique 'user profile'. For instance, the user previously signed up for a Saturday evening Latin dance class via the assistant. The agent may also infer implicit preferences (e.g., remaining fridge space or eating habits) even if not explicitly stated in the user's commands.

The assistant translates the user's intent into tasks and decomposes them into executable actions. Each action can be executed by invoking tools (e.g., mobile apps). In cases where an app does not provide an open API, the assistant can leverage Graphical User Interface (GUI) image recognition to simulate user interactions.

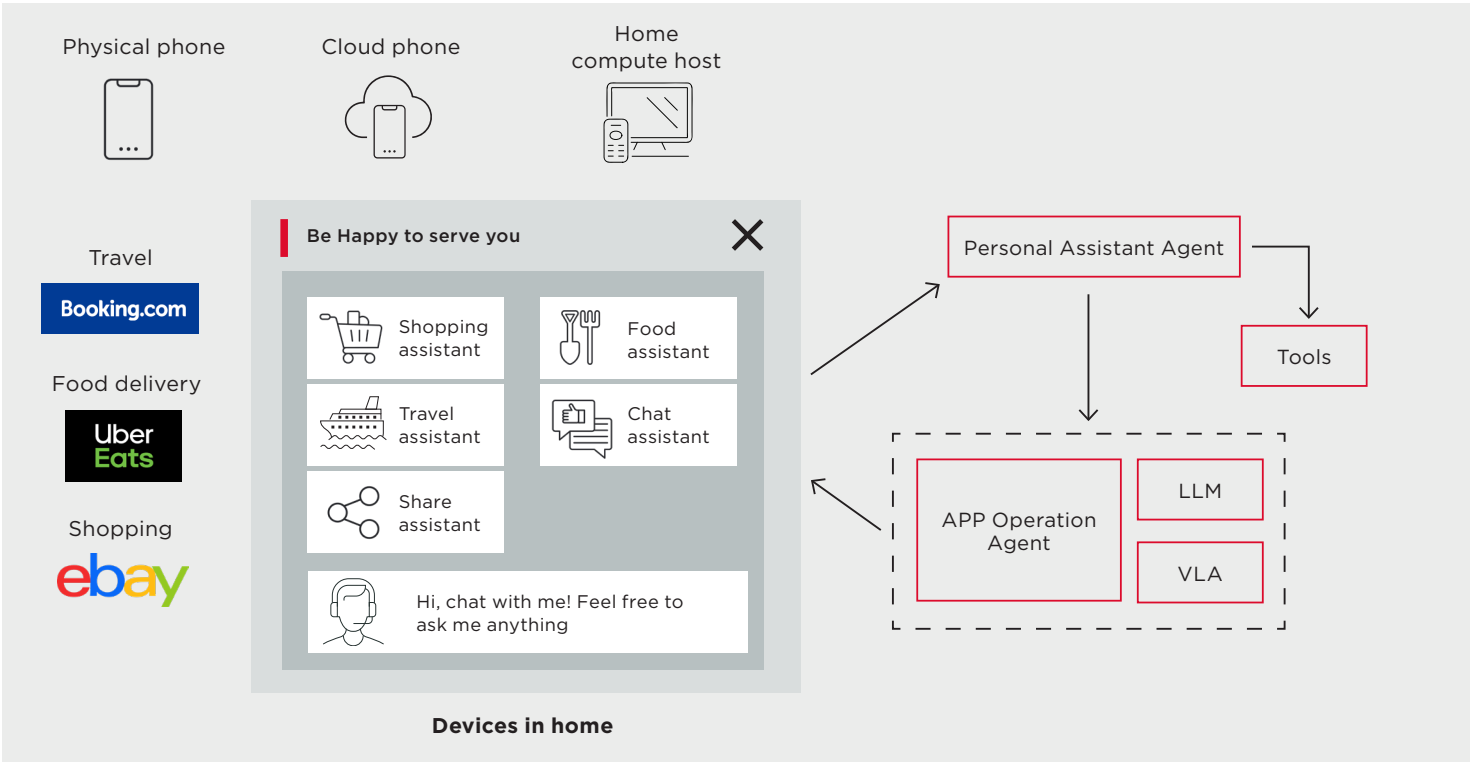

Figure 2, Example of Home AI Agent Scenarios

4.2.3 Home AI Agent: Key Value Propositions

- Expanding Operator Service Capabilities: Beyond connectivity services (e.g. home broadband), operators can offer low-latency AI services like voice assistants, serving multiple household members.
- Personalised Experiences via User Profiling: Long/short-term memory of user preferences enables hyper-personalised recommendations (e.g. suggesting recipes based on eating habits).
- Secure Identity Authentication: Leveraging SIM-based authentication, the assistant ensures secure yet frictionless access to sensitive actions (e.g. payments).

4.2.4 Core Requirements for Home AI Agent Realisation

Home AI Agent services must operate seamlessly across various devices through a unified network. A network-integrated AI service function enables intelligent coordination between agents, contextual awareness, and real-time data processing. This integration ensures that AI services are not isolated but work collaboratively across home gateways, media set-top boxes, IoT hubs, and cloud systems to deliver uninterrupted and intelligent user experiences.

- **Agent Collaboration Interface:** Standard APIs and messaging protocols to enable cooperation and task-sharing among multiple AI agents.
- **Context Awareness & User Intent Inference:** Real-time understanding of user context based on location, time, and behavioural patterns.
- **Distributed AI Processing Based on Network QoS:** Dynamic allocation of AI tasks between local devices and the cloud for optimal performance.
- **Security & Privacy Protection:** Encryption, authentication, and access control to safeguard user data.

To avoid 'vendor lock-in' and promote interoperability, a vendor-neutral integration platform is essential for managing diverse smart home devices. This platform should unify the user experience across different manufacturers and ensure scalability and compatibility through standard protocols and modular architecture. To achieve this goal, support for key elements, e.g., standard protocols and unified control interface, should be defined and broadly supported.

4.3 Cognitive Network Evolution

Agentic AI is expected to cast profound influence in network evolution. To start with, the increasing adoption of agentic AI technologies, terminal devices (e.g. robots) and applications (e.g. personal assistant) may become AI agents themselves. The network is expected to be able to provide connectivity service (for example, agent communication network (ACN)) to these new "users", including authentication and authorisation, routing, and specialised traffic optimisation, etc.

The cognitive evolution of CSP networks is being redefined by agentic AI, leading to an era of autonomous, intelligent connectivity. As CSPs seek greater agility, scalability, and efficiency, agentic AI stands at the forefront of transforming legacy infrastructures into future-ready, cognition-enabled ecosystem, where intent-based automation, function generation, and intelligent orchestration converge to deliver truly personalised and responsive network services.

4.3.1 Generative Network

The emergence of agentic AI marks a significant leap forward. These AI-driven agents are designed to perceive their environment, reason about goals, plan actions, and execute decisions independently across distributed network domains. In a CSP context, this means deploying intelligent agents at various levels (from core to edge) to autonomously manage resource allocation, optimise routing, enforce security policies, and personalise customer services based on user or business stakeholder intent.

The network can recognise user intent, generate the necessary network functions and services, and dynamically orchestrate them into workflows that align with real-time demands.

For example, a CSP could request, "I need a secure, low-latency network for a smart factory deployment," and the cognitive network would automatically translate that intent into a fully configured, optimised network setup. This eliminates the need for demand-side expertise in network design and significantly accelerates service provisioning.

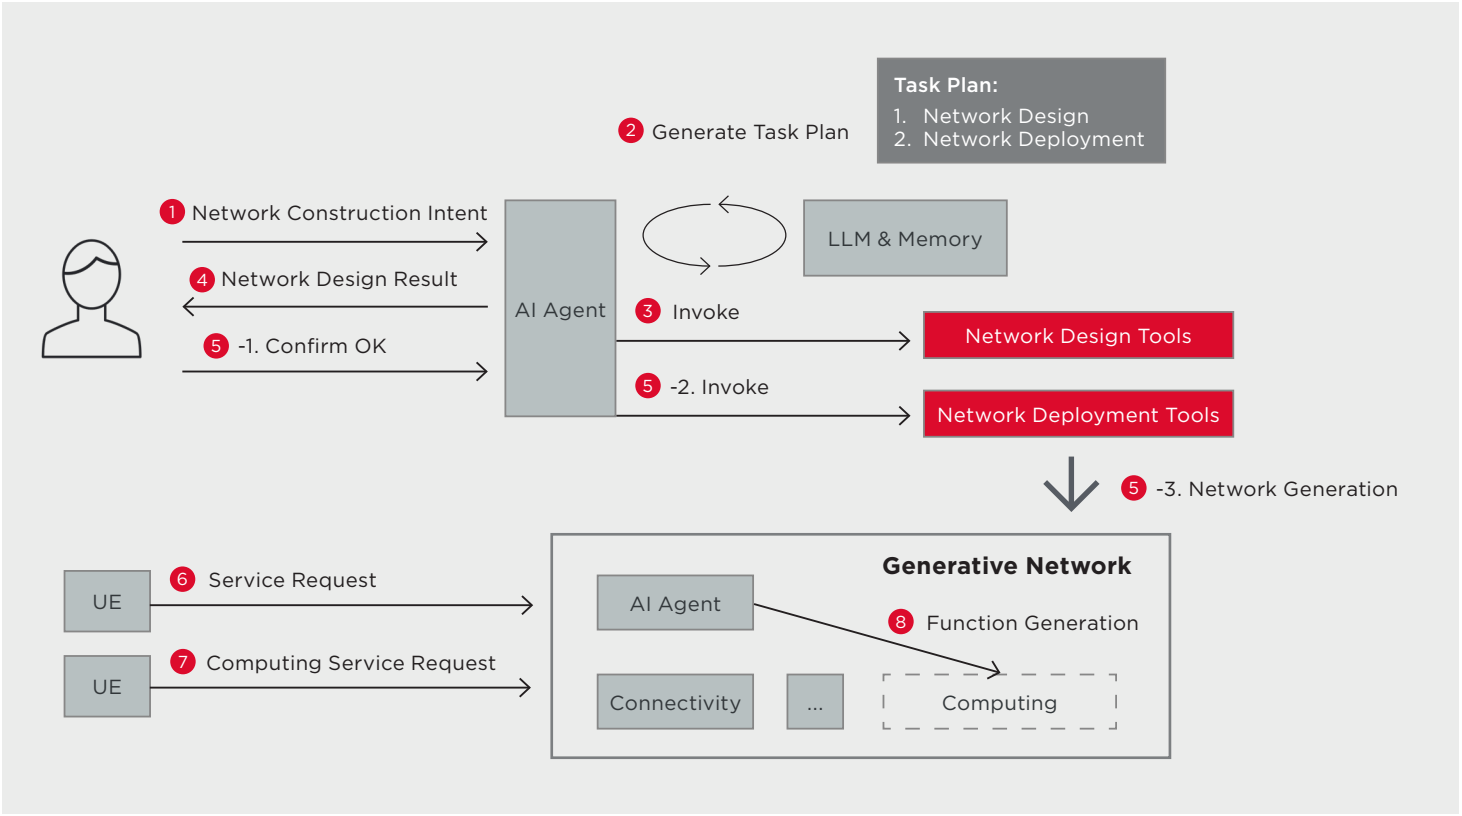

Figure 3, Intent-Based Secure Network Configuration Agent

These capabilities are made possible thanks to advanced technologies such as reinforcement learning, federated learning, and semantic parsing of intents. By leveraging these techniques, the network not only fulfills current requests but also learns from past interactions to improve future responses, creating a continuous cycle of optimisation and adaptation. This eliminates the need for demand-side expertise in network design and significantly accelerates flexible service provisioning with low threshold for customers.

Furthermore, for a real personalised experience, the networks will have to react to dynamic environmental and contextual information. Integrated Sensing and Communication (ISAC) will be the key technology enabler that allows the CSP network to perceive the physical environment (e.g., object detection, motion tracking, channel state estimation) and will support agentic AI with real-time situational awareness (e.g. for environment-aware beamforming, superior vehicular safety). The networks and applications are mutually beneficial.

4.3.2 Personalised experience

To provide users with a personalised service experience, AI agents will enable the network to intelligently perceive specific user requests to tailor customised responses and provide personalised recommendations to the users. By analysing user data and interaction history, they can even anticipate user needs and offer proactive solutions, enhancing the overall user experience.

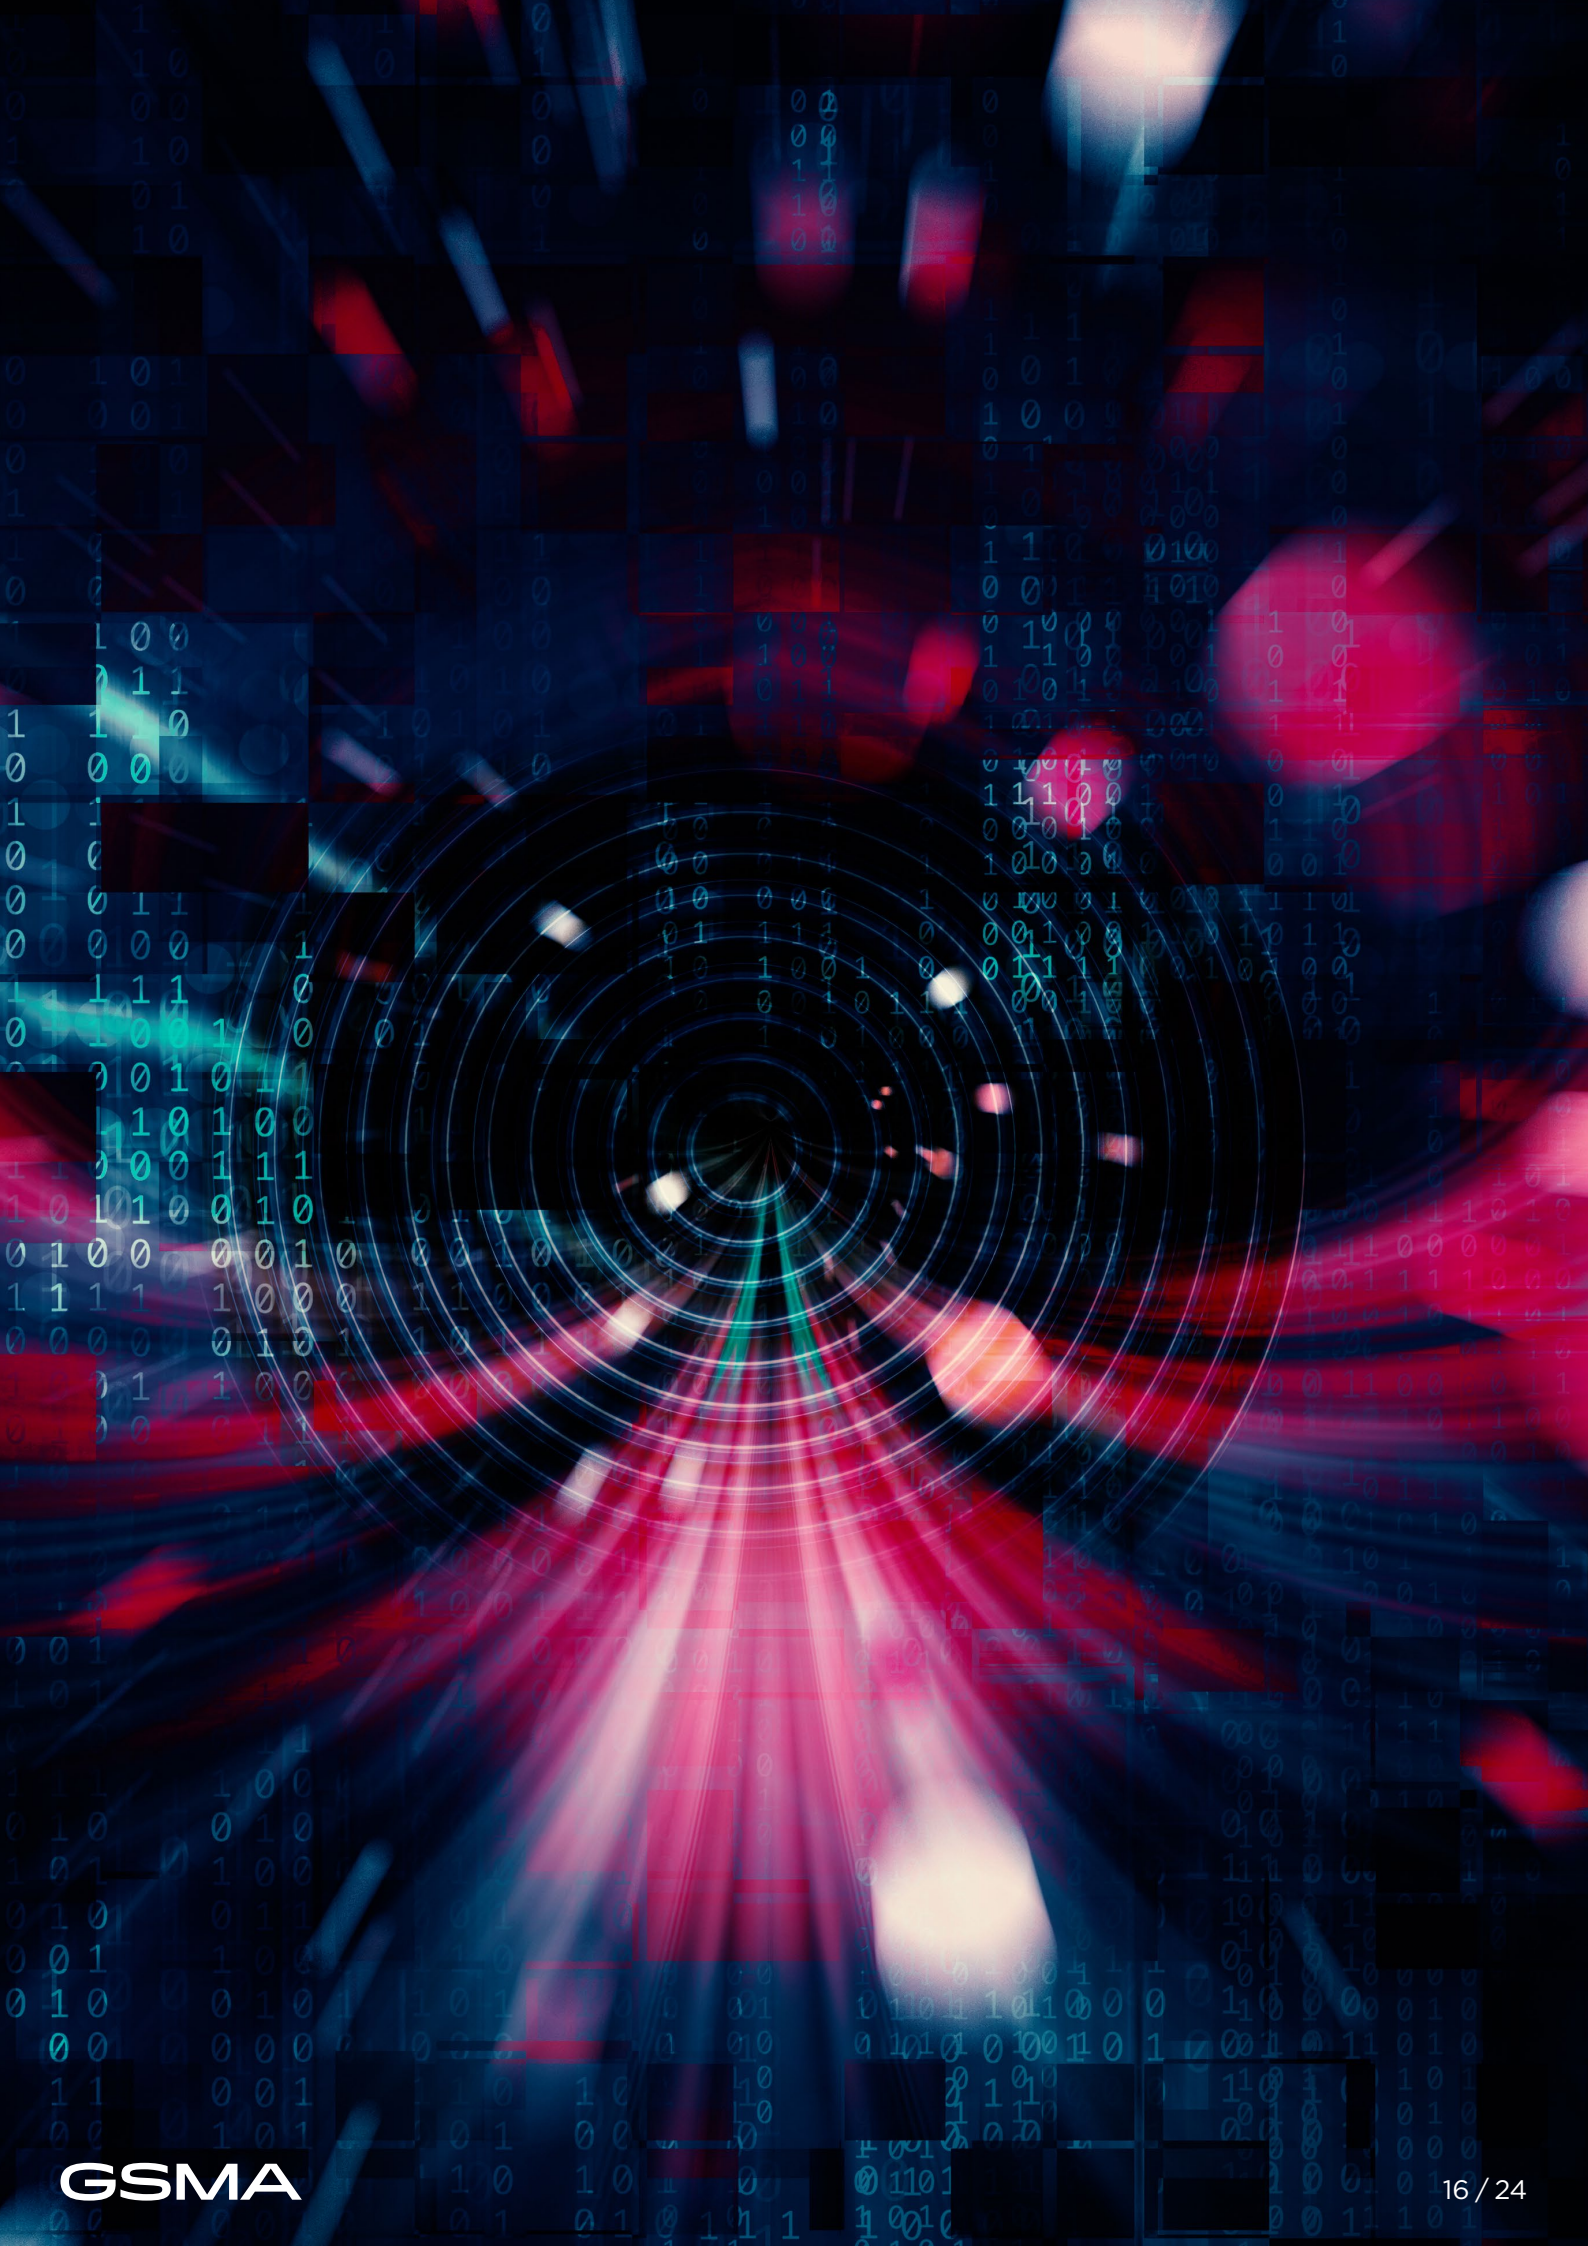

4.4 Smart and Digital Life Companion

Personal assistants for entering the Smart and Digital Life Sector

The utilisation of AI agents as personal assistants in the 'Smart Life Business' by mobile communication companies can significantly contribute to enhancing customer experience and efficient service delivery, potentially establishing a strong market position in the expansion of the Smart Life Business, in the following ways.

Enhanced agentic service Integrating with internal Telco capabilities:

- **Real-Time Location-Based Services**
  By collaborating with communication infrastructure, real-time location information can be utilised to provide information about nearby stores and events. AI agents can understand individual preferences for personalised information presentation.
  For example, they can timely notify users about discount information at specific stores or events happening nearby, making the user's movement more convenient and valuable.
- **Application to Financial Services**
  By integrating with the financial and payment services possessed by telecom operators, AI agents can analyse user behavior histories and preferences to provide personalised financial services.
  For example, they can analyse user behavior and spending patterns to offer savings plans and investment advice.

Unique Personal Digital Companion integrating with external agentic ecosystem

- The seamless integration of related services such as finance, payments, entertainment, and shopping within the telecom operator's economic ecosystem presents the potential for AI agents to support a planned approach. For instance, they can provide useful information and guidelines or steps for executing plans for events such as marriage, career change, and children's education.

5. Technology Enablers Landscape (Open Ecosystem Components)

By aligning agentic services with localised advantages and cross-industry ecosystems, telcos can transform from pipe providers to AI-driven value architects as "Telco Agentic AI Service (TAaaS) Providers". The future belongs to those who master the telco agentic stack, as shown in Figure 4, an interoperable, standardised open Agentic AI stack is essential, spanning various open ecosystem components which would be key to enable different Telco Agentic AI strategies, including tooling, orchestration, collaboration, and agent market/ capability exposure layers.

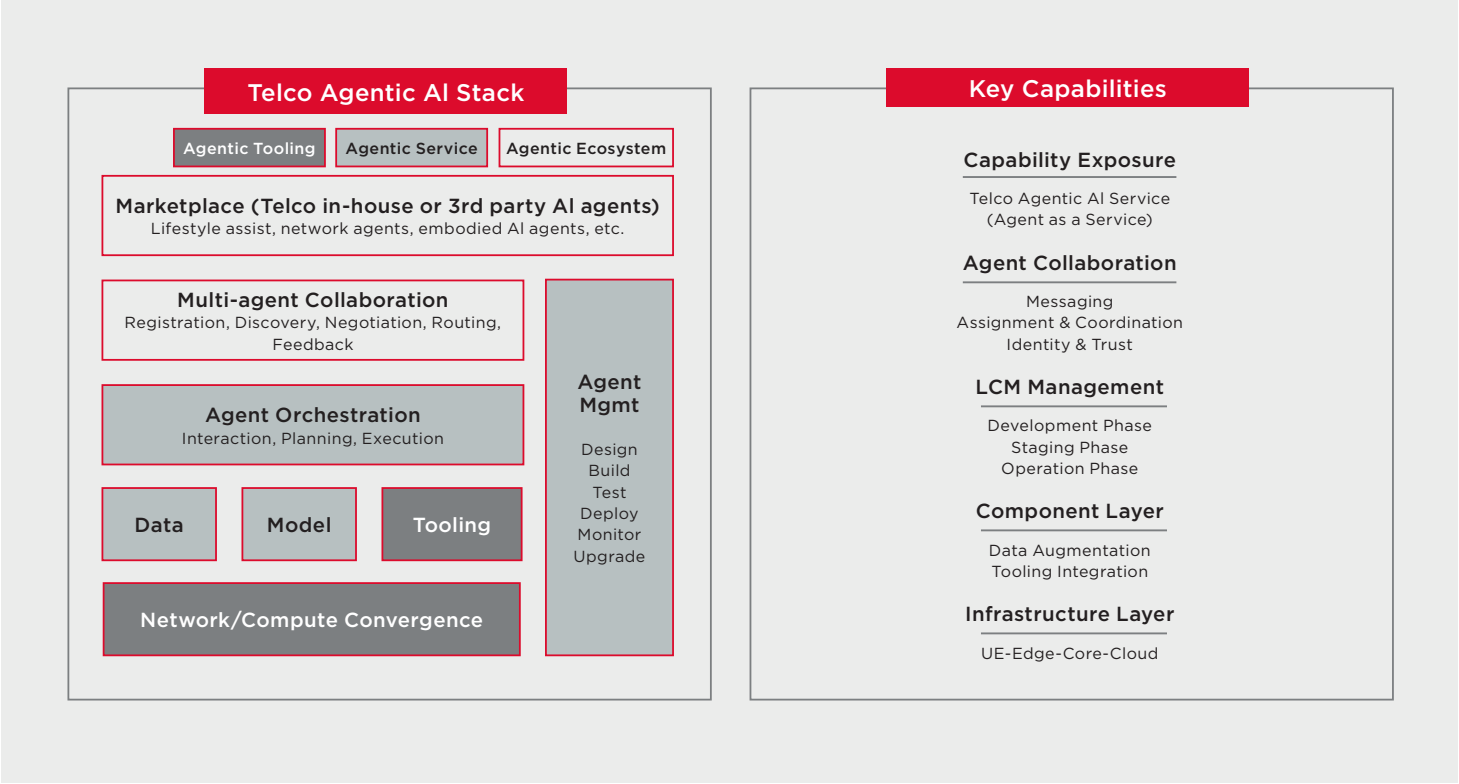

Figure 4, Telco Agentic AI Reference Stack and Capabilities

5.1 Network/Computing coordination

As the demand for computational accelerates (driven by the rise of Agentic AI), CSPs face a pivotal inflection point. Traditional models of centralised cloud computing and isolated on-device processing are nearing their practical boundaries, constrained by limited bandwidth and the ultra-low latency requirements of next-generation AI applications. Agentic AI systems, which must perceive, interpret, and act upon their environments in real time, generate immense volumes of data—often from high-bandwidth sources like video or LIDAR (Light Detection and Ranging). This renders centralised processing increasingly impractical.

To meet these demands, the network edge (e.g., base stations or aggregation nodes) must evolve into a distributed computing network. Edge computing offers a critical middle ground: it alleviates the latency and bandwidth bottlenecks of the centralised cloud while overcoming the processing limitations of on-device AI. For CSPs, this shift is both a challenge and a strategic opportunity. Deploying agentic AI at the edge requires orchestrating a highly distributed, resource-constrained environment while ensuring reliability, security, and seamless interoperability. Success will hinge on the ability to abstract complexity, optimise workload placement, and build intelligent, adaptive networks that can support AI at scale.

5.2 Foundation models and Model as a Service (MaaS)

While general-purpose LLMs struggle with telecom complexities, facing accuracy and compliance gaps, Telecom-specific LLMs are domain-specific models enhanced by RAG (Retrieval-Augmented Generation) and KAG (Knowledge-Graph-Augmented Generation) architectures. Tuned to handle telecom-specific tasks like network troubleshooting and SLA management, via further fine-tuning embeddings and base models, these approaches achieve most task coverage with minimal compute overhead, enabling edge deployment. The sharing of telecom domain data and fine-tuning know-how across operators for these purposes will significantly accelerate such quality improvements.

5.3 Tooling integration and Agent as a Service (AaaS)

A Model Context Protocol (MCP) facilitates standardised connections between AI agents and heterogeneous data sources/tools within a single local host. It defines three service categories: Resources (e.g., files, databases), Prompts (AI Prompts), and Tools (APIs). Designed primarily for intra-host operations, MCP currently overlooks critical security considerations for cross-domain communication. For example, NLWEB, short for Natural Language Web, is also a Model Context Protocol (MCP) server, which leverages semi-structured formats like Schema.org, RSS and other data that websites already publish, combining them with LLM-powered tools to create natural language interfaces usable by both humans and AI agents. Furthermore, it would be interesting to consider telecom-specific extensions to MCP, such as MCP-T, which could use persistent sessions enabling new monetisation strategies.

5.4 Multi-agent collaboration

Multi-agent architectures, where collaborative agents coordinate to accomplish complex tasks, have been gaining increasing attention across the IT industry, with the standardisation of agent communication interfaces becoming mainstream. Examples of such initiatives include Google's Agent-to-Agent (A2A) Protocol and Cisco's AGNTCY Framework.

The Agent-to-Agent (A2A) protocol, introduced by Google, standardises inter-agent communication. A2A tackles horizontal interoperability by standardising how agents from various vendors or runtimes share capabilities and coordinate workflows across the open web. A server agent advertises its capabilities via a JSON-formatted Agent Card, enabling clients to identify optimal service providers. A2A governs task lifecycle management, including completion status and artifact handover—while supporting synchronous/asynchronous communication for multimodal content. However, it currently lacks provisions for user confirmation mechanisms, routing protocols, prompt injection attack mitigation, and channel capability negotiation.

The AGNTCY Framework, introduced by Cisco, provides an infrastructure framework for large-scale agent interoperability. It features the Open Agentic Schema Framework (OASF) for vendor-agnostic agent identity/capability descriptions; the Agent Directory Protocol (ADP), a Distributed Hash Table (DHT)-based agent discovery mechanism; the Agent Connect Protocol (ACP) for cross-platform communication with authentication, context sharing, and output retrieval; and the Agent Gateway Protocol (AGP) for efficient message distribution (unicast, multicast, P2P, PUB/SUB). This framework enables agents to be integrated into unified workflows, independent of their original platforms or environments. Despite its versatility, AGNTCY currently lacks built-in identity authorisation and user confirmation mechanisms.

6. Ecosystem Development Challenges

6.1 Standardisation gaps in agent-network interfaces

Analysis of existing protocols reveals the following essential components of Agent Communication Protocols: Capability Registration & Discovery, standardised agent description and discovery mechanisms; Task Lifecycle Management, unified definitions for short- and long-term task execution; Authentication & Authorisation, robust security frameworks for agent interactions; Contextual Collaboration, context-aware messaging for task coordination; Protocol Negotiation, support for semi-structured data exchange and multi-modal content handling; and Low-Latency Forwarding, real-time communication optimised for interactive use cases.

While current protocols in the industry focus on IT-centric scenarios, in which the communication characteristics of telecommunication networks are not fully considered, adapting them to AI-driven telecommunication networks still faces challenges. These include:

1. **Standardisation Rhythm:** To adapt to the rapid development of AI technology, protocol design needs to be future-oriented and scalable.
2. **High Reliability Requirements:** To support error detection, congestion control, retransmission mechanisms, and redundant paths.
3. **Security Vulnerabilities:** To address cross-domain threats (e.g., prompt injection, unauthorised access).

6.2 Threat Modelling in Agentic AI Ecosystems

Agentic systems introduce failure modes that extend beyond adversarial behaviour. While Adversarial Agent Injection (AAI) is a representative example, similar breakdowns occur through emergent misalignment and execution context flaws introduced by automated decision layers that embed probabilistic elements. These cannot be fully constrained by static policy. Agents may act unpredictably without being malicious, especially when probabilistic outputs are treated as deterministic inputs by downstream systems (boundary collapse). This can lead to unintended state changes in areas such as provisioning, routing, or billing due to hallucination, ambiguous intent, or excessive privileges. Risks also arise in pre-deployment stages: tampered model weights, stale ACLs, or outdated artifacts can degrade integrity before or during production.

Mitigation requires treating any model, its outputs, and its environment as untrusted. Deny-by-default every inference-triggered action, require signed artefacts and reproducible builds, and enforce agent-to-agent ACLs that enumerate exactly which actors may call which tools. Any data presented to a model via tool call must be pre-filtered or sharded by user clearance, with deterministic ACLs enforced at the storage layer before the information crosses any boundary. Any attempt to modify the control plane should correlate to an out-of-band approval. Lifecycle protections, like SBOM delta scanning and privilege-scoped sandboxing for tooling, must be applied across CI/CD, not just in production. By designing for systemic containment rather than just-in-time defence, agentic platforms shift risk from silent compromise to detectable, staged failure, enabling telco operators to maintain trust even in probabilistic, autonomous workflows.

6.3 Telecom-specific LLM development challenge

Key difficulties in developing effective telecom-specific LLMs include:

1. **Data Fragmentation:** The lack of well-balanced cross-operator data sharing to accumulate sufficient corpus for model pre-training and fine-tuning, which would accelerate quality improvements.
2. **Lack of Lightweight Optimisation Techniques:** Parameter-efficient tuning methods (e.g., LoRA) are needed to reduce training costs significantly while preserving model generality.
3. **Evaluation Gaps:** Well-established telecom-specific benchmarks are critical to measure intent understanding and decision-making accuracy. The GSMA's Open Telco LLM benchmarks should be accelerated to improve telecom-specific model evaluation

6.4 Standardisation gap in adaptive agent orchestration

Agentic AI will expand into CSP Level 5 (L5) autonomy (e.g., for network upgrades, spectrum sharing, security optimisation). Static orchestration cannot handle dynamic network behaviours. This necessitates solutions for orchestration pattern selection, task execution sequencing, tool/agent allocation, and information source invocation, as each decision requires unique data pipelines. An adaptive orchestration framework is therefore critical to construct/optimise these pipelines in real-time. As multi-agent collaboration frameworks currently lack industry standards, we urge academia-provider partnerships to validate implementations toward standardisation.

6.5 Balancing AI agent autonomy and regulatory human-in-the-loop

Autonomous networks demand globally aligned human oversight boundaries. The EU AI Act classifies telecommunication networks (critical infrastructure) and high-compute LLMs as high-risk, mandating human oversight from August 2026: CSPs must be able to override AI decisions. While CSP high-risk AI systems currently lack official listing, network security monitoring, predictive maintenance, and autonomous incident response (including Agentic AI) will likely qualify. Providers must harmonise agent autonomy with regulatory oversight globally, establishing a level field for **CSPs** under varying AI regulations.

6.6 Competition with cloud AI providers and the role of telecom infrastructure

However, compared to the hyperscalers, telecom infrastructure is expected to possess particular advantages, in the following perspectives:

- **Providing Edge AI Infrastructure:** For use cases requiring real-time processing and low latency, processing at the edge (e.g., cell towers, data centres, customer premises) where data originates becomes critical. CSPs can offer business models that provide an edge computing environment integrated with 5G-Advanced and 6G networks, serving as a foundation for AI agents to operate efficiently,
- **Offering Dedicated Network Slices:** By providing low-latency, high-reliability, and high-bandwidth dedicated network slices for specific AI agent applications (e.g., autonomous driving, industrial IoT), CSPs can guarantee high-quality services with accompanying Service Level Agreements (SLAs).

Therefore, it is clear that accelerating interoperability readiness for tool-agent, agent-agent, and agent-network standardisation is crucial. This will efficiently help CSPs embrace the Agentic AI era by redefining their roles—through offerings such as AI Infrastructure Provision, Smart Network Orchestration, and Agentic Services—and by addressing the previously mentioned challenges. Through joint innovation, industrial collaboration, and balancing co-opetition with hyperscalers (subject to compliance with applicable antitrust law), CSPs can avoid commoditisation while monetising edge intelligence and network-aware AI agent services.

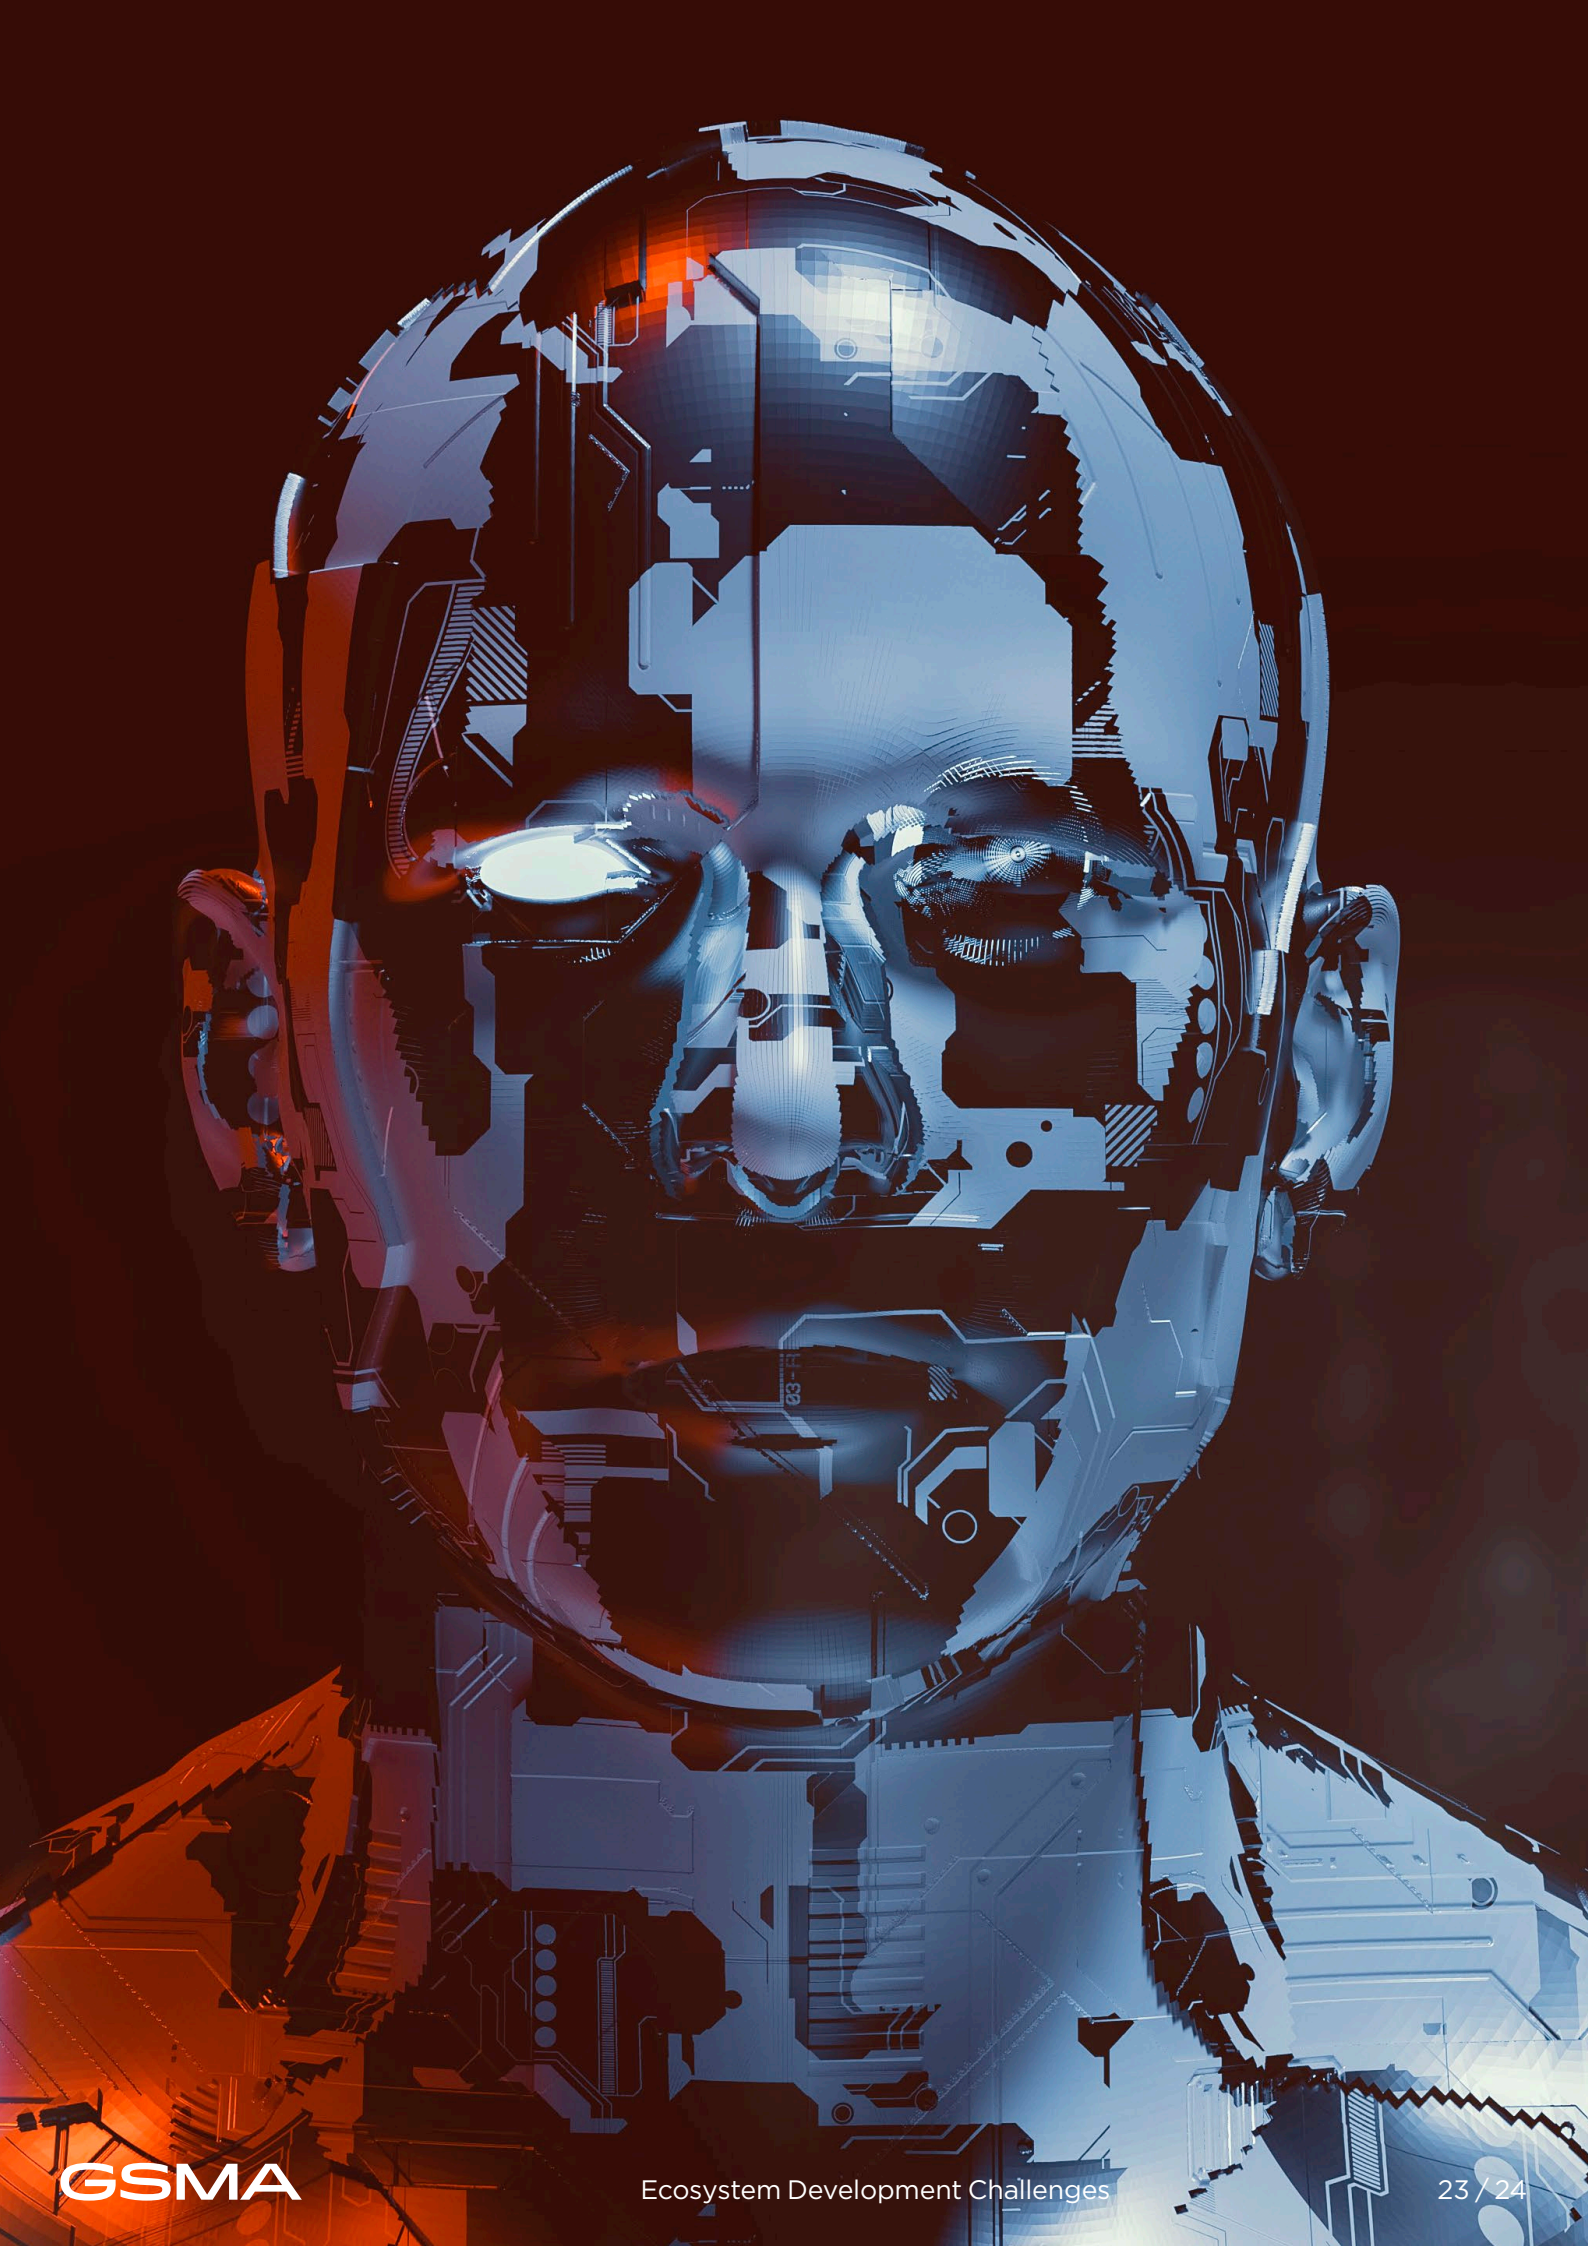

7. Conclusion and Call to Action

The advent of Agentic Artificial Intelligence heralds a transformative era for the telecommunication industry. As this white paper has outlined, Agentic AI offers unprecedented opportunities for Communication Service Providers to redefine their value propositions, enhance operational efficiency, and deliver truly intelligent, personalised experiences. From adopting tool-oriented approaches to embracing full ecosystem integration, the strategic pathways are varied, yet all necessitate a proactive stance on innovation and adaptation. However, successfully navigating this landscape requires addressing significant challenges, including standardisation gaps, the complexities of developing telecom-specific LLMs, ensuring adaptive orchestration, balancing autonomy with regulatory oversight, and strategically engaging in collaboration with cloud providers. The journey towards a thriving Agentic AI ecosystem for telecom is one of shared responsibility and opportunity, demanding concerted effort and industry-wide collaboration.

7.1 Proposed reference architecture development

This white paper posits that the establishment of a standardised framework for TAaaS providers is key for traditional CSPs to become key players in the trillion-dollar Agentic AI market. Such a framework should concentrate on scenario-based applications, general management structure, and external capability exposure, while encompassing key technical components such as system architecture design, interaction interface specifications, and communication protocol standards.

7.2 Joint telecom Agentic AI innovation outlook

Furthermore, as founding members and active participants of GSMA's Telecom Agentic AI initiative, we cordially invite our fellow CSPs to collaborate across markets to tackle scenario challenges, understand standards barriers, and cooperatively build the TAaaS ecosystem.

GSMA Head Office
1 Angel Lane
London
EC4R 3AB
UK

Email: info@gsma.com