Charlotte Patrick
Research

Developing an Optimal Agentic Architecture

Charlotte Patrick, Independent Analyst

December 2025

Abstract: As telcos take their first steps in developing simple, hierarchical agent systems, a range of decisions lies between them and the development of more complex systems that enable them to scale autonomy in their operations. This research looks at how they can create a solid agentic architecture to support their journey.

The need for an optimal agentic architecture

Telcos are at the start of a 3-stage journey to develop agentic systems capable of handling processes that require more complex and autonomous decision-making:

Stage 1: Simple agentic hierarchies enable humans to describe the action they would like to be taken to a copilot, which interprets the request, decides which data and tools it needs to resolve the request and invokes additional sub agents as needed in the hierarchy.

Stage 2: A structured hierarchy remains in place, but agents in different domains coordinate to perform more complex tasks. A "team manager" is also developed at the top of each hierarchy, which works to ensure that top-level business goals are fulfilled. They are given new abilities to make changes within the hierarchy (e.g. developing new agent workflows or requesting new agents to be created).

Stage 3: Telcos move away from hierarchies towards peer-interactions between agents. This is the beginning of "distributed intelligence" where agents deployed in different parts of a process are given decisioning capabilities. The team manager continues to understand human intent and gives chosen agents a set of goals or objectives. This allows the agents to decide how best to achieve their given task and differs from Stage 1 and 2, where agents are given specific tasks to complete.

The agents then operate independently as peers, interacting between each other rather than up and down a strict hierarchy, as seen in Stage 1 and 2. These agents are specialists in their given domain and may have a bespoke data and a memory that they can use in decision making; they will also learn from the actions they take. Unlike simpler hierarchical agents, the agent will have its own high-level business or customer objectives that may be different from the other agents around them, leading to a new level of complexity in controlling their behaviours. New capabilities in orchestration, error correction and conflict resolutions will be needed – significantly increasing the complexity of these systems.

Figure 1: The three-stage journey to autonomy

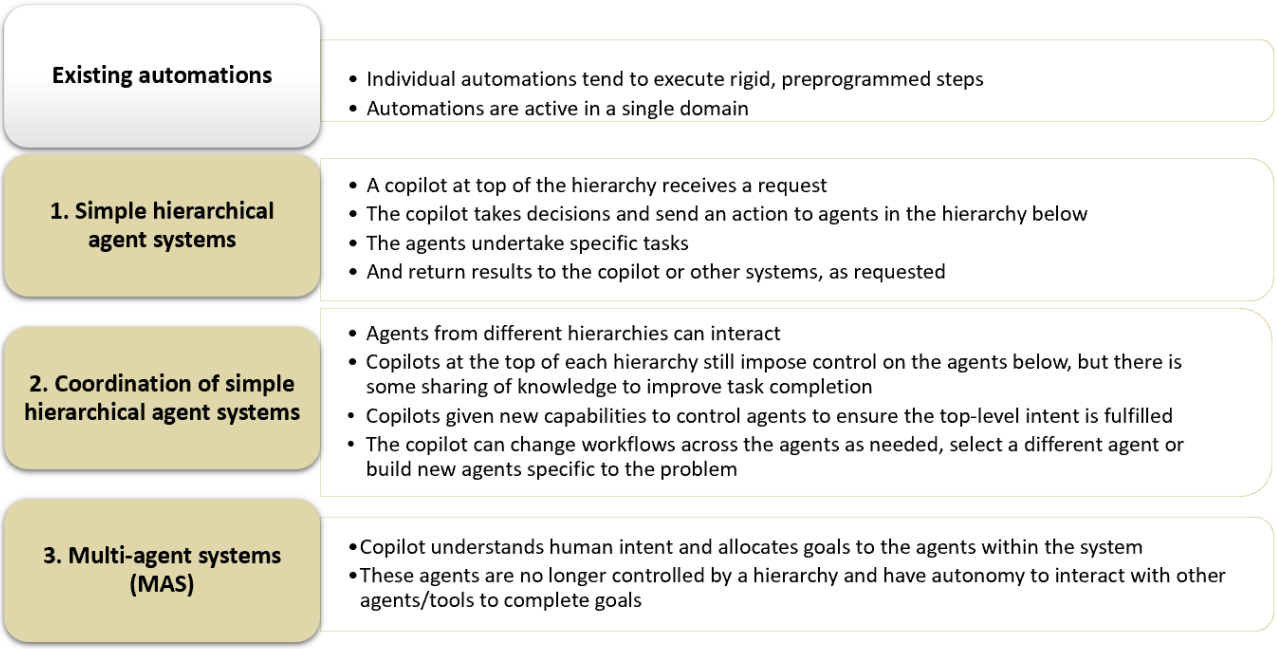

Source: Charlotte Patrick Research

Developing an optimal agentic architecture

The three maturity stages set out above need increasing sophistication in the different layers of an "agentic architecture." This architecture sets out to articulate the layers of technology and supporting processes required to build agentic systems:

Human: the architecture places the needs of users at the top of an agentic system to ensure focus on building trust and good user experience

IT systems and applications: this layer includes all systems with which the agents interact. The systems may be internal to the telco (e.g. CRM system) or external (e.g. customer or partner systems)

Agents: deployment of simple agent systems (which are controlled from the top of a hierarchy) and, in time, multi-agent systems (MAS), where agents obtain an amount of autonomy

Intelligence: deployment of appropriate LLM/SLM and AI/ML models for use by agents and other decisioning tools

Knowledge: the federation and organising of knowledge to support model decision-making with suitable context

Data: the gathering, federation, transportation and application of data for training and use in the models

Compute and storage: supporting hardware, software and cloud resources

Operations layer: processes and technology that support the lifecycle of agents, models, knowledge and data. Security, privacy, governance, measurement, compliance and risk are also considerations here

Figure 2 describes this architecture in more detail, providing details of the goals of each layer and identifying the new technologies and processes that will be required to reach these goals.

An open platform to scale agentic use cases

Netcracker's Agentic AI platform creates, runs and orchestrates AI Agents and is designed for openness and flexibility:

- *any agent – Netcracker or third party*
- *any LLM – freedom to choose the best model for each use case*
- *any data and tool – real-time access to trusted data from any source and dynamic access to the right tools – from code generators to ontologies to APIs – to execute tasks and achieve accurate outcomes*
- *any channel – seamless agent integration into mobile apps, B2B portals and internal workstations*

The platform includes an Agent Design Studio to create and test agentic use cases with speed and confidence and an AI Trust & Control layer for enterprise-grade security, governance, ROI tracking and performance optimisation. Agents can collaborate in different ways – hierarchical, peer-to-peer or hybrid – while sharing persistent knowledge about preferences, history and prior actions to deliver richer outcomes.

Figure 2: The technology needs of an agentic architecture

Layer	Overarching goals of architecture	Long-term technology needs
Human	Can easily provide goals/intent and are in the loop as needed Have trust in the system and receive good user experience	Human-to-agent interfaces with NLP, dashboards, explainability tools, validation checkpoints and feedback mechanisms
IT systems and applications	Provides good interactions with humans Receives usable outputs from agents Receive clear (and achievable) requests for input from agents	AI-integration tools such as MCP and other mechanisms such as semantic schemas to help apps and agents communicate
Agents	Copilot at the top of a hierarchy interprets user requirements and issue suitable instructions Agents and tools below adequately perform the task assigned and return results In more complex agentic systems, agents work towards specific goals set for them	Frameworks and platforms for creating and running agents Coordination and orchestration mechanisms Error management and conflict resolution tools
Intelligence	Models provide suitable inference or symbolic reasoning to support the agent/tools Able to ingest and understand prompts Models are trustworthy and improve with experience They offer transparency and other good practices	Full range of large and small models, symbolic AI/ML, time-series models, digital twins, and reinforcement learning systems Models to support coordination, error management etc
Knowledge	Offers suitable context to allow the agent to make good-quality decisions Including simple views of catalogue or inventory, as well as more sophisticated knowledge graphs or vector databases	Knowledge graphs, ontology frameworks, vector databases Semantic models and schemas Metadata management
Data	Provides enough data of suitable quality to train and support agents Data is available within a suitable timeframe for real-time usage, if needed	Data storage and lakehouses, modelling and transformation tools, and real-time processing engines. Data catalogues and quality valuation systems
Computing, storage	Stable, scalable resources for high-speed models Increasingly large storage capacity to support data and knowledge	Cloud and distributed/edge compute. High-performance and parallel compute (GPUs), storage and caching systems
Ops layer	Agents, models, knowledge and data are designed and managed according to well-defined rules and processes Robust management of governance, compliance, monitoring, risk, privacy, and security	AgentOps, MLOps, data and knowledgeOps Governance and policy management, risk assessment tools, privacy and security management, ethical and regulatory compliance capabilities, audit tools

Source: Charlotte Patrick Research

Actions needed in each layer of the agentic architecture

Figure 3 sets out some of the main decisioning that will be required when developing a robust agentic architecture.

Figure 3: Action areas within an agentic architecture

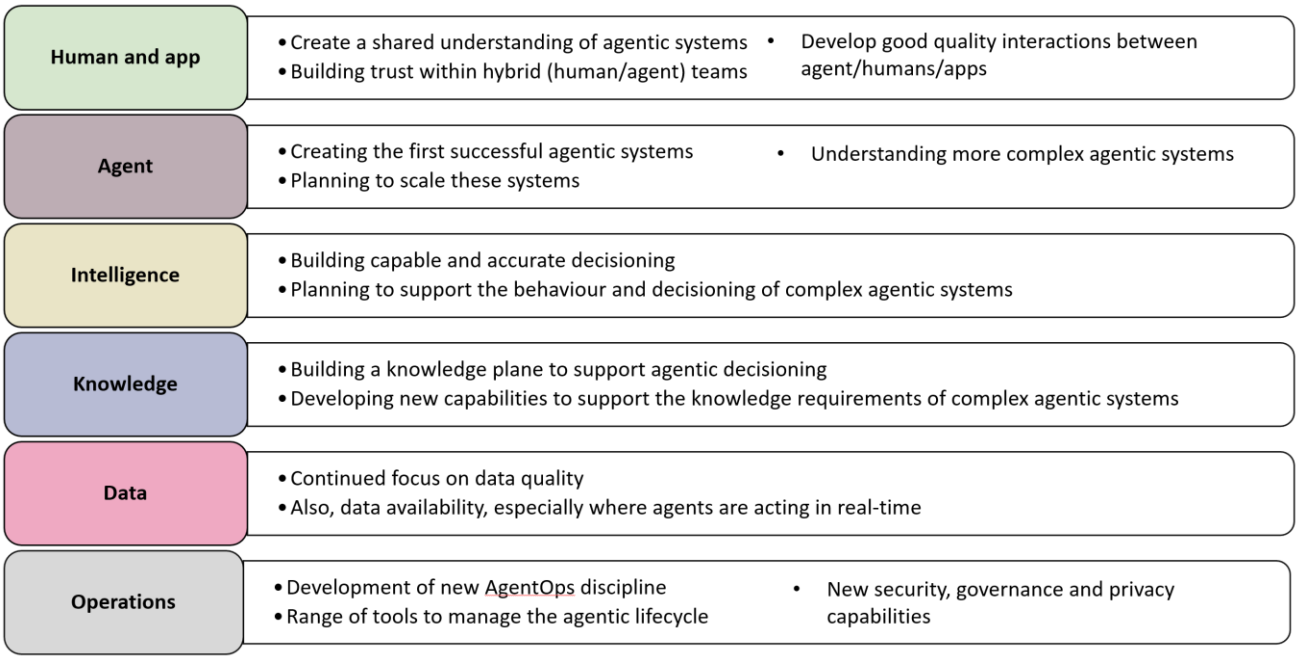

Source: Charlotte Patrick Research

Humans and apps

In the early days of agentic deployments, those interacting with agents need a working understanding of what an agentic system is and what it can and cannot (yet?) do. This will be part of building trust, along with new actions that take existing concerns around AI, such as building explainability, and extend them to agentic systems. For example, the need to build explainability into an AI/ML model will become the need to build explainability across a chain of agents, and the need for humans to be “in-the-loop” (approving model outputs where needed) will need to evolve to become about humans "on-the-loop” (where they supervise, calibrate, and redirect agents). Ethical work practices will also need to be developed for hybrid teams of humans and agents, including new policies for topics such as shared accountability (e.g. override mechanisms that ensure that a human can always pause, revoke or modify agent actions).

Agents

Early days action is around the deployment of the first agentic systems. The scope of these systems needs to be narrowly defined so that intelligence and workflows can handle the tasks given without making too many errors. It is also the time to focus on creating cross-organization capabilities that will support the ramp-up of agentic systems; for example, an agent factory and agent access management policies. Creating a plan to scale the next generation of agentic systems should also be in progress and include standard IT practices such as the development of reference agents and architectures. Also, looking at how agents doing similar tasks in different processes (e.g. anomaly detection) can be standardized and at opportunities for agents used in industries further forward in agentic systems to be brought into the telco. Planning should also include a “long-term future” stream, considering Stage 3 requirements such as new orchestration, conflict resolution, and error detection/correction capabilities to support distributed decisioning amongst agents. Many of these new technologies will be available as industries approach deployment at Stage 3. However, there are also

requirements for more distributed data and knowledge, and these may impact early-day decisions on the development of data/knowledge architectures.

Intelligence

Building capable and accurate decisioning to support agents, especially when deploying large models, requires a mix of design decisions (e.g. changing the structure and size of a model or using new AI techniques such as Mixture of Experts), a range of model development techniques (e.g. fine-tuning) and the development of models for specific tasks (e.g. a single RAN-specific action). Building Stage 3 agentic systems also requires new mathematical techniques to improve the behaviour and decisioning of more autonomous agents (e.g. game theory to balance cooperation and competition between agents).

Knowledge

Telcos need to build a pervasive system that maintains relevant information for use in the decisioning of agents (and related tools). This requires a range of new assets: knowledge graphs to create understanding of entities in a system (e.g. routers in a network) and the relationship between them, vector databases which store unstructured knowledge (e.g. recent network alarms) and improvements in existing knowledge stores (e.g. product catalogues). Decisions around how best to structure the knowledge assets – for example, domain-specific knowledge graphs that federate their knowledge using a "mesh" system will probably be more practical than a single larger graph. There are also future-focused topics such as how agents can feed back learnings into the knowledge assets, the deployment of technologies that can maintain and improve knowledge assets, and the development of robust methods for the knowledge assets to provide context successfully to agent decisioning.

Data

Agentic systems require telcos to continue to improve data quality, adding new capabilities such as data quality scoring (using ML, and possibly AI in future, to identify and prioritize data quality issues in real time) and feedback loops which allow agents to flag poor quality inputs. Data availability, particularly where agents are acting in real-time, also require new solutions:

- Although most data will come from a telco's existing IT and network systems, there will also be a need for new data sources (such as NWDAF or eBPF in the network) including external sources (such as customer systems). These new data sources require new data architecture; for example, local storage where agents are geographically dispersed at the edge or deployment of agents which interface with external data sources
- Where data is not available, new techniques allow inference on limited data sets or allowing agents to take urgent decisions based on the last known state (using the last data that it received) will be available.

Operations

The agentic architecture will need to be supported by a range of new capabilities, often developed specifically for the new risks that agents pose:

- New security risks arise across the agentic architecture: at the top level, attackers may manipulate agent prompts to induce harmful actions, there is then an expansion of the attack surface from new inter-agent and agent-tool interfaces, knowledge bases could be manipulated or data stores attacked. A variety of new security measures will be needed including identity enforcement of each agent, an agent registry, encryption of all storage, new protocols for interactions and end-to-end observability and auditing (where every orchestration or agent call is logged with metadata
- AgentOps creates a robust operational foundation to support agentic systems by focusing on the end-to-end lifecycle of large-model-powered agents, including planning, reasoning and interactions. It draws on principles from DevOps and MLOps to cover the agent as a complete entity (managing goals' planning, reasoning, actions and interactions with external systems)

- Governance across the architecture will include all existing governance exercises already in place around AI, data and knowledge. Although, regulations and good practise are still being developed specifically for agents. A range of issues will need to be covered including explainability and accountability in a hybrid environment where more than one agent, and possibly more than one human, have taken action. Current best practise advice includes human overview of all agent tasks, constraint of the actions that an agent can do without human approval and ensuring that agent tasks provide graceful shutdown so that they can be interrupted if needed.

In conclusion

Five actions that telcos should take today to progress agentic architecture:

1. Create a clear set of goals for your agentic deployment ensuring that they include an understanding of which top-level business goals they can best support.
2. Focus on developing a set of data and knowledge assets that will support the company's agentic vision. These are the most significant set of work items facing the telco and extracting the maximum financial value will take time.
3. Focus on educating the wider organization on agentic systems: what they are, what they can (and can't) do, what sort of processes offer strong business cases, how an operational team can get started, and where they find help with choosing and designing solutions to meet company standards.
4. Set up a long-term planning stream which creates a shared understanding of where more complex agentic systems might be needed and an idea of what will be needed to successfully deploy them.
5. Work by those leading the development of agentic systems should include understanding beyond the telecoms industry. Many agentic processes have directly analogous processes in other industries, and there may already be technologies and skillsets that can be bought into the telco.

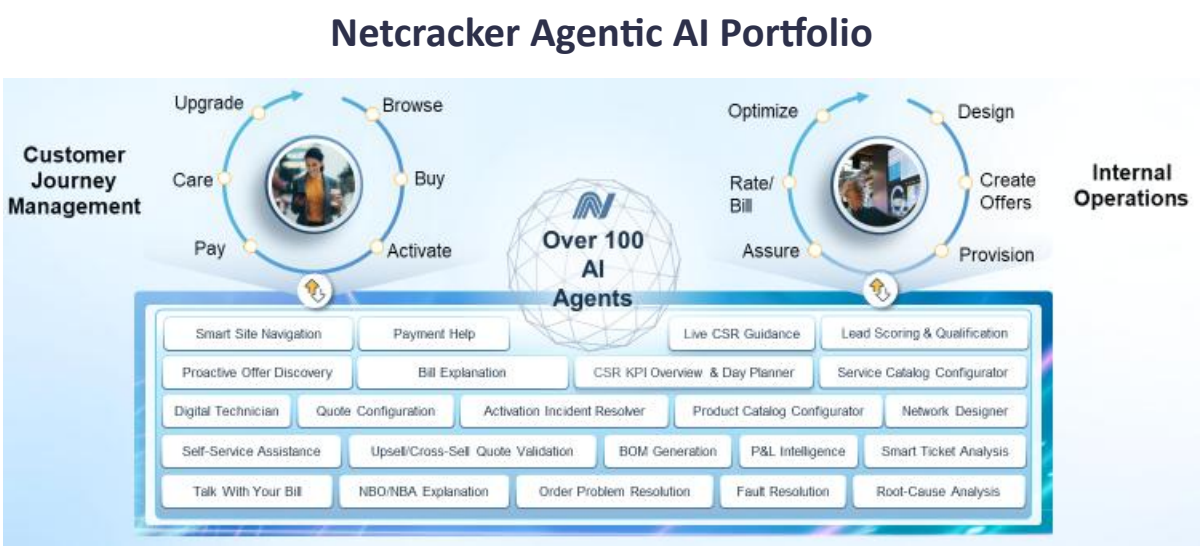

About Netcracker Technology

Netcracker Technology, a wholly-owned subsidiary of NEC Corporation, has the expertise, culture and resources to help service providers around the world transform their businesses to thrive in the digital economy. Our innovative solutions, value-driven services and unbroken delivery track record have enabled our customers to grow and succeed for more than three decades. With the latest technological advancements in key areas including 5G monetization, AI, automation and vertical industries, we help service providers to reach their transformation goals, advance their telco to techco evolution and realize business growth and profitability. To learn more, visit www.netcracker.com.