Network Innovation

EXECUTIVE BRIEFING

# Supporting the Telecom Agentic Journey

As agentic systems reach the top of the hype cycle, European telcos are beginning to run pilots in processes where there is return on investment. This research note investigates the need for an agentic architecture to support the ramp-up of agents in the network with both AI and knowledge.

Charlotte Patrick, Independent Analyst | October 2025

# Table of Contents


Foreword: From proactive analytics to the agentic journey with Vitria ..... 3

Executive summary ..... 4

The agentic journey ..... 4

What is an agentic architecture? ..... 4

How to establish a solid agentic architecture strategy ..... 5

The agentic journey ..... 7

Benefits of an agentic architecture ..... 9

The multivariate challenge of creating agentic systems ..... 12

1. How can a telco successfully build an agentic architecture? ..... 14

2. What are the financial benefits of agentic systems? ..... 14

3. How do you supply models with context? ..... 15

4. Can available data support training and decisioning? ..... 15

5. What are the benefits of a multi-agent system? ..... 16

6. What is required to build a multi-agent system? ..... 16

Conclusion ..... 17

## FOREWORD: FROM PROACTIVE ANALYTICS TO THE AGENTIC JOURNEY WITH VITRIA

This journey toward agentic systems has a deep historical foundation.

For decades, companies have worked to leverage analytics and automation to transform telecom operations, and as a pioneer in this field, Vitria Technology demonstrates a long history of success in streaming analytics and operational intelligence.

Vitria has laid critical groundwork by focusing on turning massive volumes of cross-domain data into actionable insights for automated remediation of network problems ("self-healing"). Its **VIA AIOps platform** demonstrates a sustained commitment to helping telcos overcome long-standing barriers to service quality and assurance by enabling them to apply AI, machine learning, and structured knowledge to detect issues and identify root causes across multi-vendor, multi-domain environments.

The transition to agent systems, which moves beyond scripting to enable hierarchies of agents (and, potentially, fully autonomous agents in future), is the next logical step in the journey towards autonomous networks.

Vitria's track record of driving significant, measurable improvements in service availability and resolution time provides an informed perspective on what this phase will entail:

- telcos will continue to require clean, usable data both for model training and, increasingly, for supporting real-time dynamic automations
- they will also need to build a knowledge plane, which provides solid context to improve agent understanding and support all other machine learning and AI in the network.

Vitria's unique **AI with Knowledge**™ approach utilizes a "knowledge plane" grounded in knowledge graphs to store and leverage contextual, topological, and diagnostic information. This approach is particularly effective in addressing critical needs across the network where the ability to derive accurate insights from high-volume data streams across increasingly complex service delivery layers is paramount.

By embracing this new era of intelligent automation, telcos can transform their reactive systems into autonomous networks that can anticipate and prevent service disruptions. This marks a shift toward a truly adaptive and customer-centric future for the telecom industry—a journey for which Vitria and its peers have been preparing the industry for decades.

# EXECUTIVE SUMMARY

## The agentic journey

The concept of agentic systems has risen rapidly up the hype curve over the last 18 months. It is currently beginning its descent towards a more pragmatic view of achieving complex goals, such as full network autonomy. This research noted a significant amount of vendor activity, including some "agent washing" (the rebranding of existing products, such as AI assistants and RPA, as agents with very limited agentic characteristics). Then, during interviews with European telcos, agents were observed in pilots, with telcos considering how best to scale up and support their agentic systems with knowledge and data.

## What is an agentic architecture?

The architecture articulates the layers of technology and supporting processes required to build agentic systems. Figure 1 describes these layers as:

**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

**Apps:** apps interact with both humans and agents. These may be tools 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 models to power agents and other tools

**Knowledge:** the federation and organizing of knowledge to support model decision-making with suitable context, whether the models are part of powering an agent or deployed independently as tools

**Data:** the gathering, federation, transportation and use 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, compliance and risk are also considerations here.

**FIGURE 1: GOALS OF AN AGENTIC ARCHITECTURE**

| | |
|---|---|
| **Human** | Can easily provide goals/intent and give feedback on outcomes from the system<br>Are in the loop as needed<br>Have trust in the system and receive good user experience |
| **Apps**<br>(internal or external) | Provides good interactions with humans<br>Receives usable outputs from agents<br>Receive clear (and achievable) requests for input from agents |
| **Agents** | In early agentic systems, there is a copilot at the top of a hierarchy that can interpret user requirements and send suitable instructions into the hierarchy<br>Agents and tools below adequately perform the task assigned and return results<br>In more complex agentic systems, the copilot sends tasks to more autonomous agents which work together on the task |
| **Intelligence** | Models provide suitable inference or symbolic reasoning to support the agent/tools<br>Able to ingest and understand prompts<br>Models are trustworthy and improve with experience<br>They offer transparency and other good practices |
| **Knowledge** | Offers suitable context to allow the agent to make good-quality decisions<br>In the early days, this might be a simple view of catalogue or inventory<br>But will develop into more sophisticated knowledge graphs or vector databases; to ground decisions<br>Can be used by a large model or a digital twin |
| **Data** | Provides enough data of suitable quality to train and support agents<br>Data is available within a suitable timeframe for real-time usage, if needed |
| **Computing, storage** | Stable, scalable resources for high-speed models<br>Increasingly large storage capacity to support data and knowledge |
| **Ops layer** | Agents, models, knowledge and data are designed and managed according to well-defined rules and processes<br>Day 1 capabilities include an agent factory that is used across the organisation<br>Processes and, where necessary, new systems are in place to manage governance, compliance, monitoring, risk, privacy and security |

*Source: Charlotte Patrick Research*

## How to establish a solid agentic architecture strategy

Figure 1 illustrates how many different elements are required to function well to support an agentic system. These agentic systems must then be deployed across every network process, domain, and layer if the telco's ultimate goal is an autonomous, intent-based network.

Creating a strong start to this journey should include:

1. Big picture, strategy thinking
    - Given the amount of hype around agentic systems, discussions should be brought back, as always, to the customer. How can we enhance their experience and provide products that meet their needs?
    - When priorities are understood, strategic thinking should turn to the long-term architectural requirements to deliver these products and experiences, attempting to avoid technical debt from implementing immature agentic technologies
2. Focus on a sensible set of first use cases. These could include:
    - simple monitoring or configuration tasks that monitor for events or anomalies and take simple decisions. Typical starting places are assurance, security and fraud
    - the creation of agents that support successful gen AI deployment around improving user interfaces, natural language processing, or working with documents. Typical starting points include the creation of technical documentation at the end of a workflow on the network

- automations where there is already progress in developing quality data and processes. Trouble ticketing, as an example

3. Architectural decision-making to support the agentic system, both now and in the future, across all layers. Key questions should include:
    - How can we build trust amongst our teams?
    - How can we develop suitable models – potentially using peer collaboration or working with vendors?
    - When considering new models, it should also be asked whether any existing models are providing a good enough solution to remain in place into the longer term?
    - What set of new techniques and technologies is needed to ensure the provision of quality data into the longer term?
    - How do we develop a solid foundation for the delivery of agents across the organization (e.g. agent factories, governance and security)?
    - How can we instantiate human and machine knowledge to create a suitable layer of context for the agents?

**Vitria Technology's pioneering journey in operational intelligence has positioned its VIA AIOps platform to directly address the hype surrounding today's agentic systems**

As the industry moves toward a more pragmatic view of autonomous networks and grapples with "agent washing," Vitria grounds its approach in its unique "AI with Knowledge™" methodology.

Unlike limited agents or simple RPA, Vitria's architecture leverages a robust knowledge layer by utilizing structured, human- and machine-readable knowledge graphs to provide critical context for model decision-making.

This allows Vitria's agents to move beyond basic monitoring tasks toward explainable autonomous resolution and self-healing. By focusing on customer experience and delivering solutions with measurable ROI, Vitria offers a clear path through the complexities of autonomous systems.

## THE AGENTIC JOURNEY

Transitioning from existing automations in the telco to using agents to create more autonomy is described as a "journey" because many underlying capabilities need to be established, and some technologies required to power more complex agentic systems are not yet developed.

This journey is typically described in three stages:

**Stage 1** enables humans to describe the action they would like the automation to take. The copilot at the top of the agent hierarchy interprets this action and can make some (fairly simple) decisions on which agent or tool should be given a task. The agent/tool then brings back information or confirmation that the task is complete

**Stage 2** sees agents still operating within a structured hierarchy – but agents in different hierarchies coordinate to perform more complex tasks (such as service fulfillment) or to gather a fuller picture of the performance of services between domains (e.g., RAN vs. core). The concept of a "team manager" is also developed; the copilot is given new capabilities to adjust workflows across agents and ensure top-level business goals and intents are fulfilled

**Stage 3** becomes significantly more complex as agents are granted autonomy to complete tasks. This allows them to operate in highly dynamic solutions, bringing fully "self-x" deployments (e.g., self-healing), but their business model and technical feasibility remain uncertain.

Agentic systems are typically described as having 5 key types of functionality, which become more sophisticated as the telco progresses through the three stages described above:

**Adaptability**: the agentic system will need to become more adaptable to dynamic and unpredictable environments

**Perception**: it will need to become more adept at finding and using data from a variety of sources

**Reasoning**: simpler agentic systems use more deterministic, rules-based decisioning, while a multi-agent system requires an increasing number of agents to reason using large models

**Learning**: simple systems have some closed-loop learning, but agents in more complex systems will learn together

**Action**: the types of action that can be taken by more sophisticated systems will, obviously, become more sophisticated at the later stages of agentic systems.

Figure 2 provides more detail around the progression across the 3 stages of the agentic system (in the columns) and these 5 key functional areas (in the rows):

**FIGURE 2: INCREASED SOPHISTICATION IN THE AGENTIC JOURNEY**

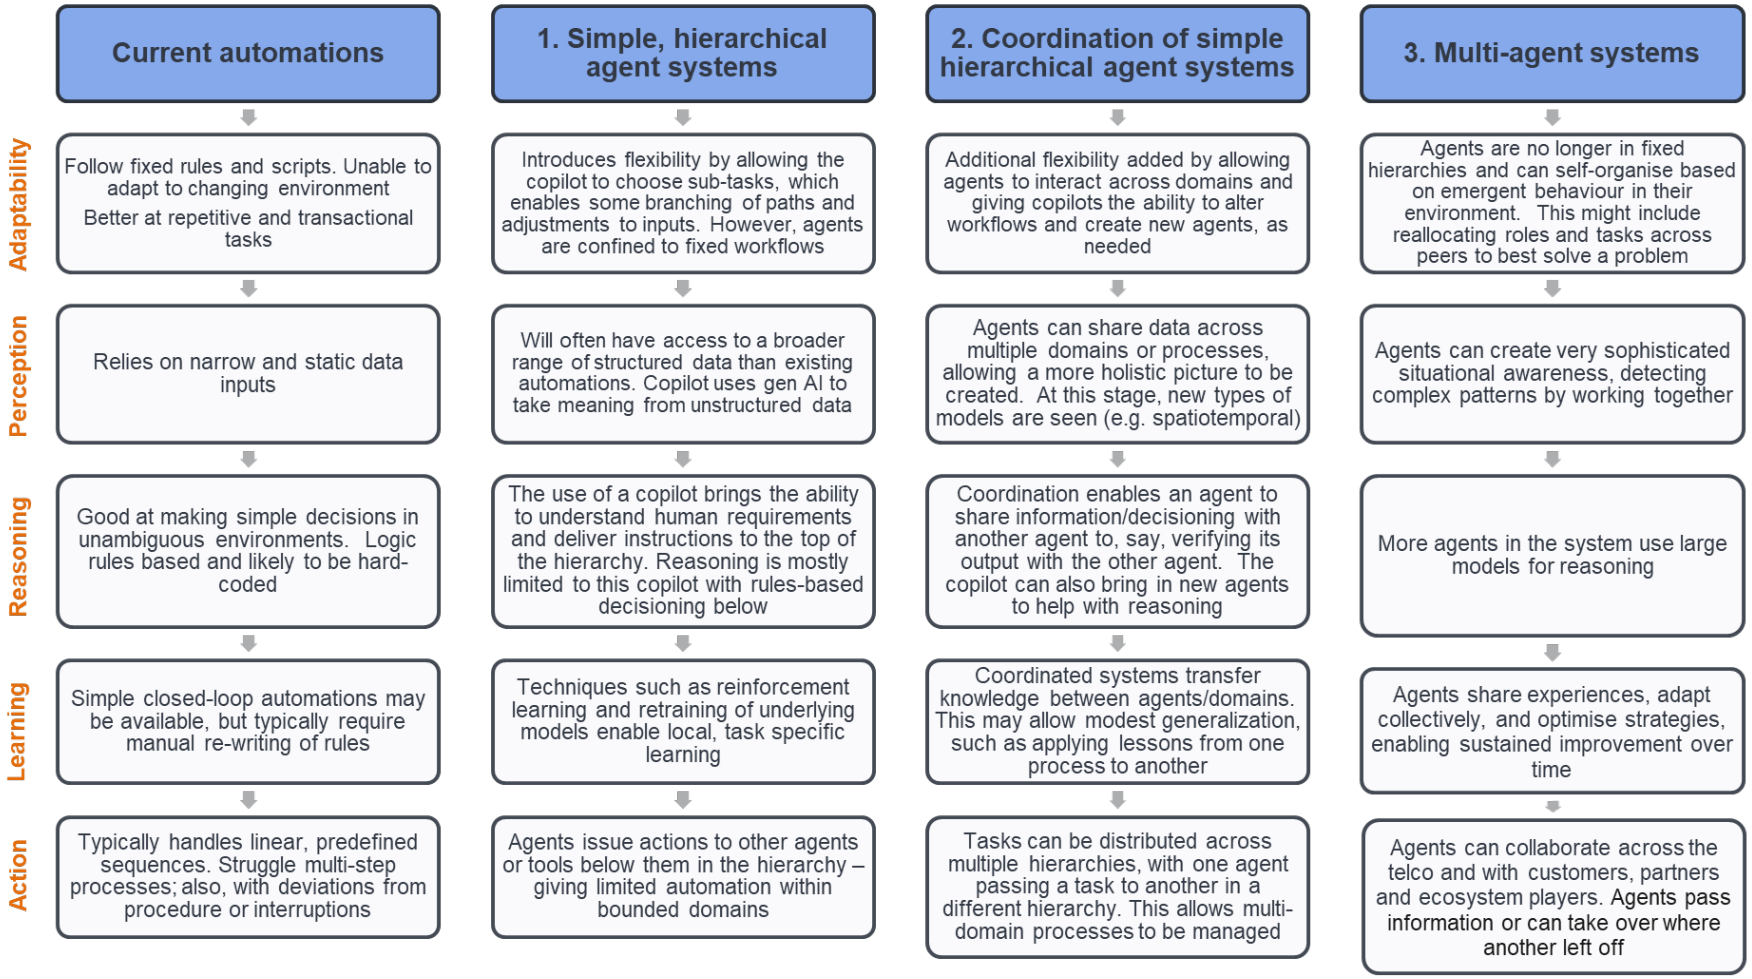


*Source: Charlotte Patrick Research*

## BENEFITS OF AN AGENTIC ARCHITECTURE

Figure 3 sets out 8 areas of benefit from deploying an agentic architecture. Each is placed on a scale of:

- **Expected benefits** (x-axis) - where the expected financial benefits are ranked from higher on the right-hand side of the graphic, to lower
- **Availability** (y-axis) – with near-term benefits at the top of the diagram and those unlikely to be available for multiple years at the bottom.

Those areas above the x-axis are expected to be mostly achievable with the first, simpler hierarchical agents. With those at the bottom increasingly requiring multi-agent systems.

The "support for future products" box ranks low on expected benefits, mainly because it is still unclear whether agentic systems will truly help telcos generate significant financial benefits from new products. In addition, the "financial benefits" box is shown across time, illustrating the potential difficulty of demonstrating ROI in the early days of agentic system deployment.

More easily realisable benefits like improved customer experience and an incremental improvement on current automations will see telcos creating smaller projects of simple agentic systems in the early days as they build up skill sets that will allow them to access some of the larger benefit projects.

**FIGURE 3: THE BENEFITS OF AN AGENTIC ARCHITECTURE**

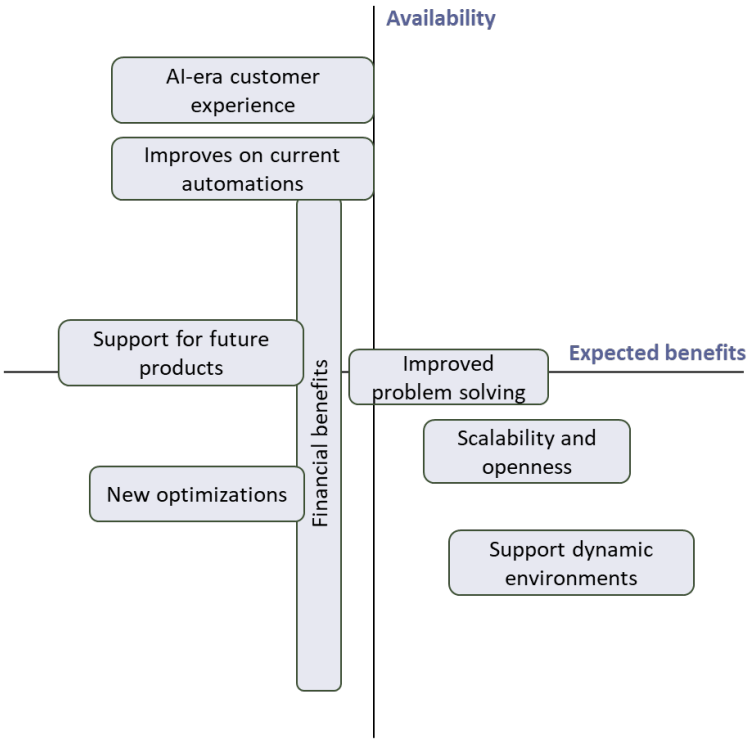


*Source: Charlotte Patrick Research*

**Improves on current automations**

Agentic systems start by enhancing the coordination of automations across network domains, which leads to more sophisticated observability of previously siloed data sets. The building of coordination between teams will begin to eliminate delays that occur due to human handoffs in traditional workflows and agents may, over time, be able to handle a much larger number of tasks in parallel. Agentic systems also bring the potential for various machine-like qualities: they can work for extended periods and are not prone to human errors, improving speed and efficiency.

**AI-era customer experience**

Currently, humans both inside and outside the telco are increasingly interacting with AI-powered natural-language processing, raising their expectations about the support a machine can provide. Copilots bring these new experiences to agentic systems, supported by agents and tools. They promise to enable better quality interactions (than today's chatbots) using natural language, bringing back quality and personalized responses, and perform multi-stage tasks to generate their responses. Also, an agentic system may offer a telco new ways to build company-specific experiences. These experiences utilize data and knowledge (collected from both systems and people within the organization) to create automated and consistent customer experiences that are unique to the company.

**Support for future products**

Agentic systems will support certain processes within new services that require more complex automations, such as end-to-end service orchestration, which would have been difficult to achieve with existing levels of automation. Agents also allow the company to interact more closely with its customers and partners, and there may be particular new products, such as 6G's Integrated Sensing and Communications, which require millisecond decision-making at the edge that can only be undertaken by localized agents. In a broader sense, agent systems may also allow telcos to develop services outside of their traditional skill sets, building new revenue streams in target verticals by adding vertical-specific agents into their products.

**Improved problem-solving**

Agents will bring new ways to tackle problems found on the network: In the early days, an agent can be trained on the data and knowledge of a particular domain, giving it specialist expertise. This expertise can then be called upon when needed by other agents/tools. Adding intelligence and memory to these domain-specific agents starts to create distributed intelligence, offering the ability to solve more complex problems. For example, non-linear problems, which lack a clear beginning or steps leading to a specific resolution, can be tackled by more than one agent. The first agent might specialize in detecting certain patterns, while another agent reasons about the best response. By communicating and sharing their partial solutions, they can adapt to scenarios that don't follow a linear set of rules.

**Scalability and openness**

Agentic systems are designed to be open and flexible, allowing capacity to be added/removed as needed, and new technologies to be swapped in as they become available, reducing the opportunity for traditional challenges such as monolithic stacks and vendor lock-in. More sophisticated, decentralized, agentic systems also reduce the opportunity for single points of failure, thereby increasing resiliency.

**Better optimization**

Having agents that sit in one domain or at the edge of the network, controlling their local environment and interacting with other agents as needed, allows for new optimizations. The most financially impactful are likely to be those that improve load balancing on the network or cloud, and those that turn off resources when not required. However, others will monitor for disruptions, reroute operations, and escalate to the NOC/SOC only when necessary, allowing a range of processes to run more smoothly and reducing the number of events the NOC/SOC need to handle.

**Support dynamic environments**

The capabilities described already contribute to the final goal of agentic solutions, which is to support more complex and dynamic environments, such as disaggregated 5G SA or 6G networks, where every part of the network is in a state of change.

Agents with distributed intelligence will:

- adapt to the needs of the network in their local domain
- adjust their activities in that local domain in real-time
- make decisions at all levels of the network more quickly than the current mix of humans and machines
- whilst bringing more detailed insight to a human for certain decisions than they have previously had access to.

**Financial benefits?**

The question of total financial benefits from agentic solutions is a complex one. Previous calculations suggest that 1/3 of the total financial benefit from deploying AI, ML and automation in the network will come from Level 4 and 5 processes, which look set to be enabled by agentic solutions. This is not a huge amount, calculated at around $300 million a year for an average telco with revenues of $15 billion – leaving a question as to the ROI available for the significant amount of new technology needed to run the agentic system and supply it with suitable data and knowledge. The likely outcome of this will be the need for individual agentic projects to prove their ROI.

Other longer-term financial benefits include the ability of open agentic solutions to reduce the size of a telco's yearly bill from its systems integrators. Additionally, really intelligent agentic solutions could continually improve their own performance by designing new agents and modifying the capabilities of existing ones to deliver enhanced results.

## THE MULTIVARIATE CHALLENGE OF CREATING AGENTIC SYSTEMS

Moving from a world of rule-based automations, machine learning, and simple closed-loop systems towards more autonomous decision-making presents numerous challenges. This section focuses on the 6 major biggest questions that telcos will have to answer as they move from simpler architectures towards a MAS. Figure 4 sets out these questions.

FIGURE 4: THE SIX IMPORTANT QUESTIONS SURROUNDING AGENTIC ARCHITECTURE

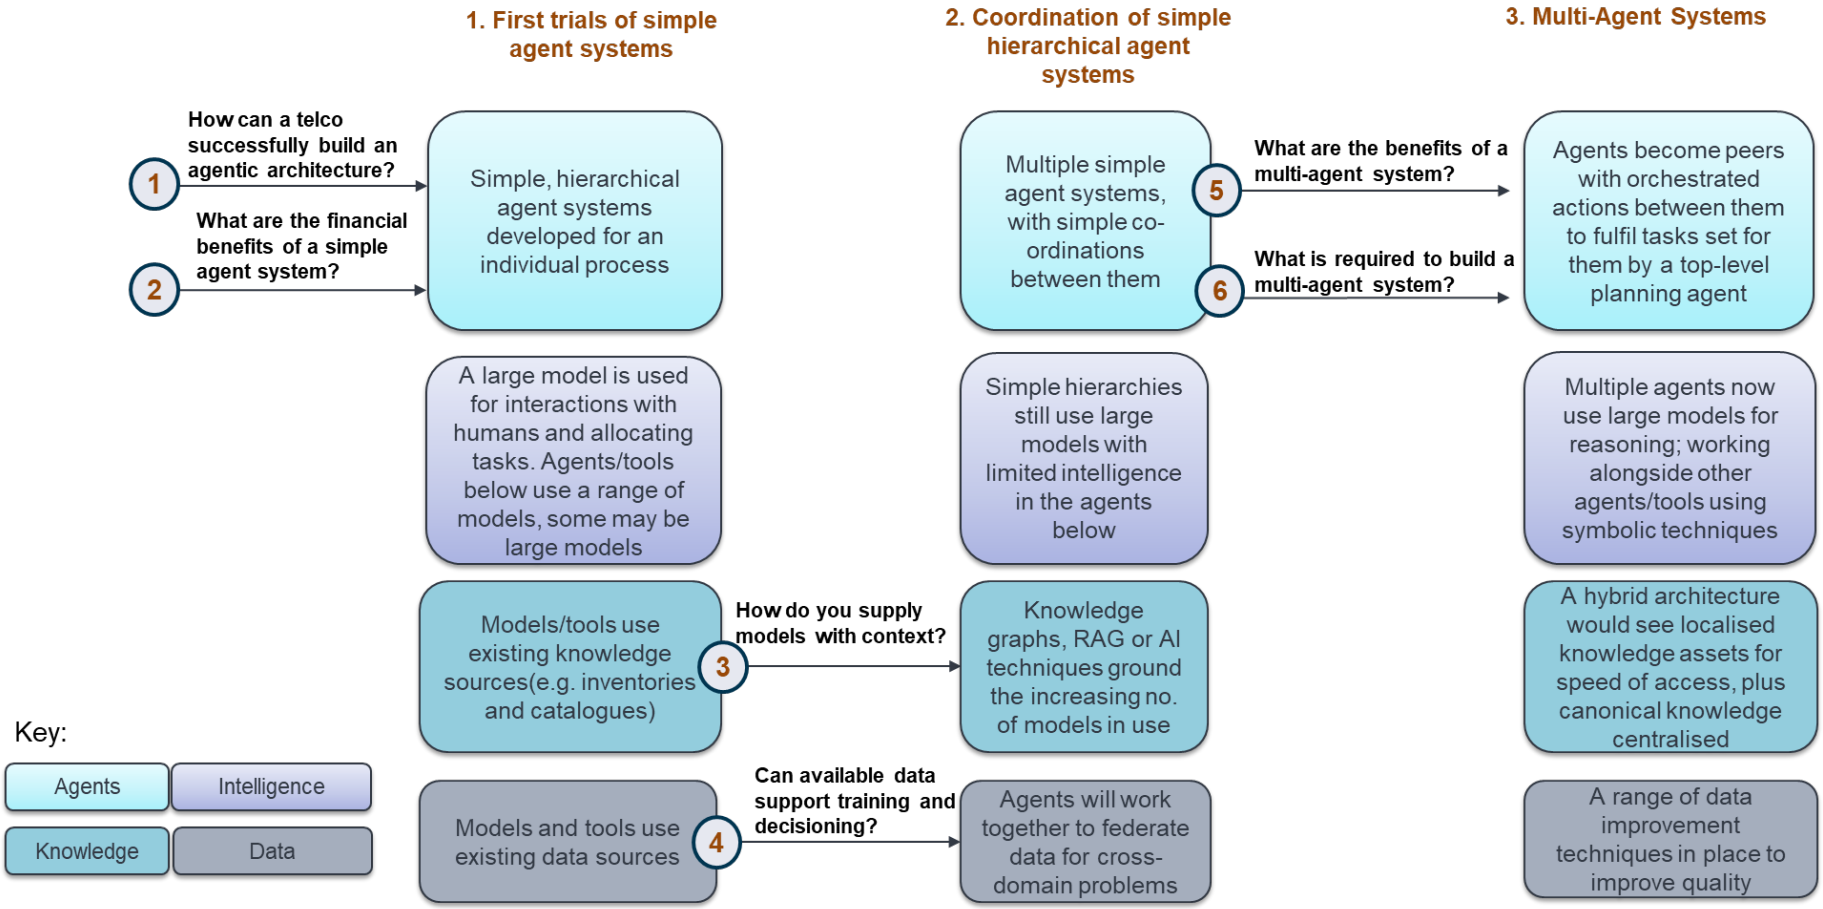


*Source: Charlotte Patrick Research*

## 1. How can a telco successfully build an agentic architecture?

Telcos beginning their agentic journey need to focus on:

1. Ensuring that the first agentic projects selected are a success. This means selecting small projects with simple, hierarchical agents. In the early days, any agent decisioning should be narrowly defined, limiting its complexity; and any orchestrations between agents or tools should be simple rules-based handoffs or workflows
2. Developing a solid lifecycle management for the agents, their models, data and knowledge to ensure that there are no issues as agents pop up across the organization
3. The cross-organizational steering groups already involved in AI should turn their attentions towards agentic systems. The group needs to understand the range of new challenges added to their existing work items, particularly those related to risk management and security.

## 2. What are the financial benefits of agentic systems?

The business case for agentic systems will be at its strongest if a significant figure can be articulated for the financial upside of delivering a Level 4 autonomous network. In recent research, an annual upside of $242m was calculated for an average telco (with yearly revenues of $15.6bn) in reaching Level 4. This is not huge because much of the cost reduction looks set to be achievable using simpler automations in Levels 1-3; and placing large revenue uplifts against the delivery of 5G and 5GA seems imprudent currently.

Moving existing automations into simple agent hierarchies without the need to invest much in data, knowledge assets, or more complicated MAS will likely be financially viable. But then, more sophisticated MAS deployments will only have a positive ROI if the use cases are:

- in processes that require slower and more sophisticated thinking, achieved via the coordination of agents with intelligence and memory
- in processes that are so dynamic that they require distributed intelligence to make multiple, speedy, and reasonably complex decisions.

It is likely that deploying more sophisticated, distributed MAS will take 5 to 10 years. In this timeframe, ROI for their deployment will improve because of the increased maturity of open architectures, large and small models, data mesh and knowledge planes. These are all required for the creation of a MAS's distributed architecture, and their maturity will help drive ROI.

## 3. How do you supply models with context?

The grounding of models with knowledge will help increase their accuracy going forward, with telcos currently using existing stores of information, such as catalogs and inventories. The development of knowledge graphs will create structured, shareable knowledge that can be used and improved by both humans and AI; and the more sophisticated telcos have pilots and first deployments of these graphs in processes such as the creation of trouble tickets, where the graphs bring together knowledge from multiple systems and provide links between graph entities to help with root-cause analysis.

Knowledge graphs require considerable effort to develop, and questions still to be answered include the need for development of a single graph (versus multiple graphs that represent particular data sets); also, the need for new capabilities, such as neurosymbolic AI, to enable large models to take information from the graphs.

**Think of Knowledge as the backbone of context for AI**

A smarter, richer knowledge base allows the AI to spot issues before customers even notice, recommend real-time fixes, and learn from every situation. This shifts the focus from constant firefighting to proactive problem-solving.

This isn't just theory. Vitria has seen this in action with major telecom providers:

- an internet service provider struggling with mystery outages and unhappy customers slashed incidents by 65% in 90 days after deploying Vitria VIA AI platform. They began detecting 90% of issues before customers reported them
- one of the largest U.S. ISP eliminated a quarter-million unnecessary technical visits annually, saving an estimated $20 million per year. They did this by using the VIA AI to correlate customer care issues with infrastructure problems
- a Fortune 200 mobile carrier accelerated its 5G rollout by up to three months by building a knowledge-based assurance system from the ground up.

## 4. Can available data support training and decisioning?

The availability of sufficient data in each telco process to support the training and continuous operation of larger models remains an issue when considering the viability of deploying agentic systems. Solutions to the lack of data include continued efforts to federate data and telcos working together to build models using pooled training data. There are also techniques such as the use of synthetic data to mimic real-world scenarios, transfer learning (which uses pre-trained models or datasets from related domains (e.g. other networks or industries) and fine-tunes them on the limited telco data available), and data interpolation techniques, which create new data points by interpolating between existing ones. It is assumed that most

data availability and quality issues will be overcome over time. However, an additional challenge is the availability of data for real-time decisioning, which could be expensive to collect. This might see agentic systems focused on real-time, "slower" problems in the near term.

## 5. What are the benefits of a multi-agent system?

There remains a degree of uncertainty as to whether distributed intelligence is necessary, and if it is, whether it will be cost-effective.

There are, however, reasons for moving in this direction:

**New ways to support complex dynamic environments** – the complexity of a disaggregated 5G SA or 6G network, with every part of the network in a state of change, creates a complex challenge. Distributed intelligence sees agents with specialist local knowledge trained to adapt to their own domain and then working with other agents, as required in these environments

**More complex problem solving** – where problems are not linear (with a clear beginning and a series of steps that lead to a specific resolution) rules-based systems become "brittle" and break down. Distributed problem-solving can tackle these types of problems

**Ability to optimize multiple goals simultaneously -** an intent-based network will throw up competing goals (such as energy efficiency, quality of service, and latency requirements), and distributed intelligence can optimize for more than one goal

**Less chance of a bottleneck** - decentralizing intelligence creates a more robust architecture where one agent can fail without bringing down the whole MAS

**Future ecosystem interfaces** - when AI/automations need to be shared/integrated with partners, customers, and other 3rd parties, a MAS will naturally develop to support the need to integrate

**Provides solution to privacy concerns** – The localization of agents means that data will be less likely to have to move across clouds, domains, applications, and geographic regions.

## 6. What is required to build a multi-agent system?

A range of issues will arise from the deployment of a MAS, related to the complexity of its design and management, the control of agents, and the development of viable agent decision-making. Some of these issues will require more traditional software solutions, such as new monitoring tools, to create robust lifecycle management for the large numbers of agents deployed. But some issues will require new control solutions that are more cutting-edge – for example, deploying mathematical frameworks to determine the best way to resolve conflicts between multiple agents. In general, there appears to be a theoretical solution to every problem identified today, but each of these solutions typically requires deploying a new model. Even if this model could be deployed to support all agents within a MAS, each deployment increases expense and complexity, leading to a view that MAS are not going to be cheap.

## CONCLUSION

When considering the future of agentic systems and how telcos should progress towards this future, it is interesting to consider them in light of other recent technologies deployed within the telco.

For example, agentic systems offer telcos more flexible and intelligent automations, similar to the first RPA systems that provided them with their first opportunity to automate operations widely. Reviewing reported progress from RPA's starting point in around 2016, a Forrester survey (4 years later) in 2020 noted that "just 52% of enterprises have progressed their RPA initiatives beyond their first 10 bots". Many firms found RPA deployments harder to scale than promised. Bots broke when applications changed, maintenance costs were high, and new processes were required to support the bot lifecycle. ROI from RPA offered clear savings on high-volume, repetitive processes (claims processing, billing, reconciliations), but fewer savings were seen when processes were variable, cross-functional, or poorly documented.

Making future assumptions about agentic systems from the telco's experience with RPA, it could well be that many firms fail to scale agents beyond trials in the next 5 years – most likely due to a range of factors not dissimilar to those noted in the paragraph above; for example, a lack of robust lifecycle management. As discussed in this research, agents require a considerable amount of work within the supporting architecture. Although some basic data and process quality improvements will likely be completed as part of earlier automation support, failing to support the complete architecture's development will significantly hinder progress.

Taking a second example of Open RAN, this has parallels with agentic systems as both are based on disaggregated and open systems. Progress towards the full vision of Open RAN is limited, according to an STL Partner's 2025 survey, with "vRAN giving disaggregation without the integration challenges, potential for interoperability failures and degraded performance [of Open RAN]. vRAN also offers a commercially proven solution with a limited set of features compared to a full Open RAN, but provides new value around elasticity of the network and ability to create new use cases, offering more solid ROI in the early days.

It may well be that agentic systems will also find a natural place on their journey where telcos decide to "rest a while". Here, there are tangible benefits from the use cases where agentic has already been deployed. With only the most enthusiastic telcos choosing to engage in the next part of the journey; battling with immature technologies and complexity. The natural spot for this resting place would be between stage 2 to stage 3 in Figure 2. This space would see the telco spending most time on developing the knowledge and data levels of the agentic architecture, which provides them with a unique asset with which to pursue deployment of new services on a truly autonomous network.


VIA AI: Pioneering AI+Knowledge for an Autonomous Future

www.vitria.com