Whitepaper

June 2026

“It helps to know what you’re talking about”

Effective and efficient agentic IP network automation

Author: Grant Lenahan, Principal Analyst

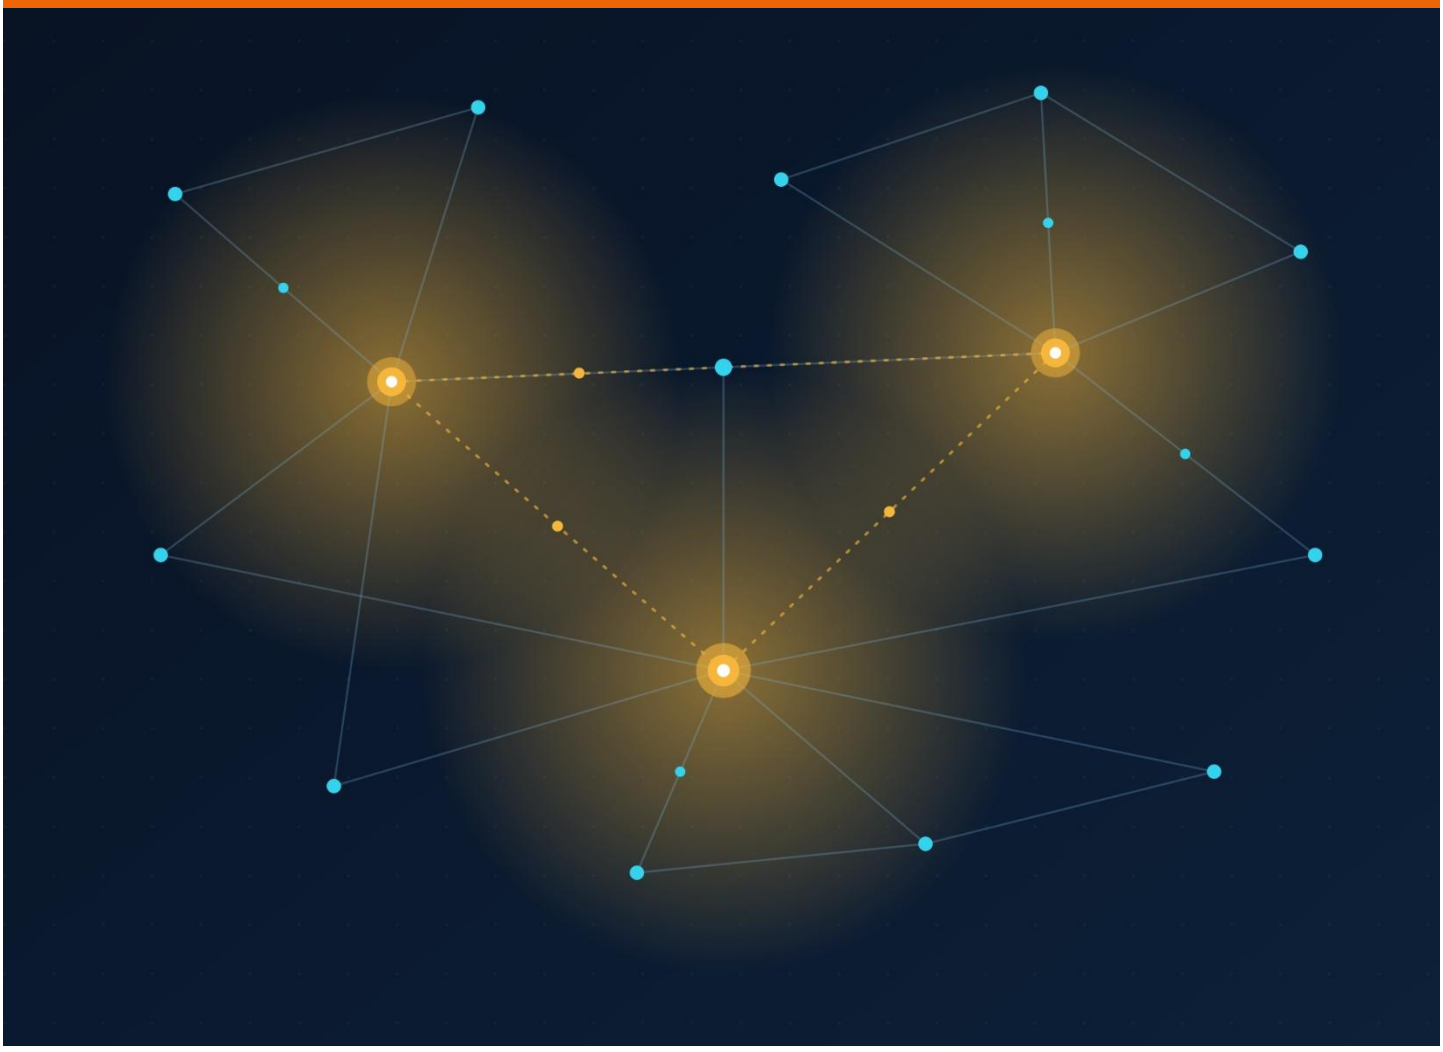

In partnership with

NOKIA

Appledore RESEARCH

Published by Appledore Research LLC • 44 Summer Street Dover, NH. 03820

Tel: +1 603 969 2125 • Email: info@appledorerg.com• www.appledoreresearch.com

Figures and projections contained in this report are based on publicly available information only and are produced by the Research Division of Appledore Research LLC independently of any client-specific work within Appledore Research LLC. The opinions expressed are those of the stated authors only.

Publish date: 10 June 2026

Cover image: Appledore Research

TABLE OF CONTENTS

Foreword

This paper will outline the best practices emerging on how to successfully introduce Agentic AI into network automation, with a specific focus on those characteristics unique to complex network technology domains, such as the IP core and edge.

Based on extensive recent primary market research we will discuss:

- The importance of Control Theory, intent, and autonomy as a foundation for AI agents
- The consequent importance of a solid autonomous domain platform foundation
- The state of the industry, with a focus on network-operations-centric AI
- Consistency with best practices – including academic Control Theory, TM Forum, and Appledore’s own recommendations and best practices – (documented here).
- The emerging industry consensus on platform priorities, architecture, features
- Supplier and sourcing imperatives – including DIY, AI firms, and telecom domain specialists
- … all with a focus on network domain automation in general and the IP core/edge network in particular

Our findings and recommendations are anchored in extensive primary market research that Appledore performed over this past year, including many dozens of deep briefings with CSPs, network suppliers, ISVs and hyperscalers. From this we have identified priorities, as well as a broad consensus about what makes Agentic driven network autonomy successful – with a focus on network domains. In this paper, we will define success criteria, and finally note consistencies with the approach Nokia is building into its Network Services Platform (“NSP”) which is Nokia’s intent-based manager for IP core and edge networks and operations. The net result is that in NSP, Nokia’s approach to network autonomy, Agentic AI, and their application of IP networks in particular is entirely consistent with these research findings. For those interested in diving more deeply into our research and findings we present a Bibliography of linked research at the conclusion of this paper.

Nokia’s new CEO, Justin Hotard, is focusing the company on what he believes will be a two-pronged impact of "The AI Supercycle” – both how it will create huge new demand and challenges for Nokia’s customers; and how it can be leveraged ubiquitously to change the quality, productivity and agility “production possibilities frontier” (my words, not his). The investment in NSP, along with similar investments in many other Nokia management systems begin to provide concrete evidence of that second impact.

Executive Summary

Over the past 12-18 months, AI, and in particular agentic AI, has burst on the scene. Like any highly promising new technology, it immediately entered a form of hype cycle—hailed as the panacea for productivity and automation. Two differences stand out in this cycle however:

1. The promise of AI to accelerate and simplify automation appears to be legitimate and holds unusually great promise
2. A quick industry pivot from vague excitement to an industry building consensus on the path forward, and then building the less sexy but essential infrastructure on which Agentic AI success in the network must be built.

A convenient date for the "coming out party" for network operations centric Agentic AI is the TM Forum's DTW in Copenhagen in 2025. At that conference, while agentic AI was the marketing buzzword, the real discussions behind closed doors were on far more meaningful and practical engineering efforts underway across the industry. In particular, we had myriad discussions about structured data, ontologies, digital twins, foundational database techniques, and how much easier application of AI is when a proper intent-based control loop already exists.

The 12 months following DTW were packed with many dozens of in-depth discussions with CSPs, leading suppliers, and industry experts, augmented by a large survey with CSPs globally (all undertaken independently by Appledore). All the resulting research is linked in the bibliography. In net, the industry has developed a consensus blueprint, and leading network and network operations software suppliers are perusing all or parts of this blueprint.

This promising path forward is based on a solid approach: build the infrastructure that enables efficient, accurate and open Agentic AI reasoning. The magic of such foundational work is that nearly any use case will be made easier and better; and part of it opens up data and processes to innovation – whether by suppliers, CSPs or third parties.

The following summarizes, from our research, the key areas of industry consensus:

- Begin by applying Agentic AI to independent technology domains, before attempting to link and cross domains (which increases complexity)
    - Note: this is entirely consistent with every set of best practices for both Autonomy and Control Theory
- Adhere to the core principles of Control Theory, specifically:
    - Define a control loop based on intent and self-correction
    - Act as locally as possible; then abstract and, only when needed, remit larger problems to the next layer up
- Focus on domain expertise, not "generic AI" expertise. Domain expertise and data includes:
    - Topology, decorated with Ontologically rich structured data
    - Technology specific digital twins
    - Deep understanding of operational parameters and consequences

- Build a foundation that can support and simplify the development of Agentic workflows and use cases for decades – making AI use cases, simpler to build, faster, more accurate and with less risk. Common elements of this infrastructure include:
    - Domain specific topology (Including Physical, Logical and Service topologies)
    - An ontological structuring of data (mapped to topology)
    - The most relevant domain metrics, overlaid on that ontology
    - Open, AI-native interfaces (MCP APIs primarily)
    - Guardrails and rules, to ensure secure, appropriate access to data and constrained operational actions
    - Agentic Studio to efficiently develop use cases, using pre-built capabilities
    - Natural Language Processing to replace traditional UIs and vastly improve human productivity, while simultaneously increasing access to information and features within documentation and next-generation OSS (Domain managers)

Coming out of Copenhagen, and strongly reinforced by our research, we published that focusing on the technology of AI (*which model? Which AI supplier? Massive training?*) was in fact the wrong path. The correct path viewed autonomy and productivity as based on four inter-dependent pillars as illustrated in the graphic below:

Figure 1: The four pillars of successful Network Agentic AI

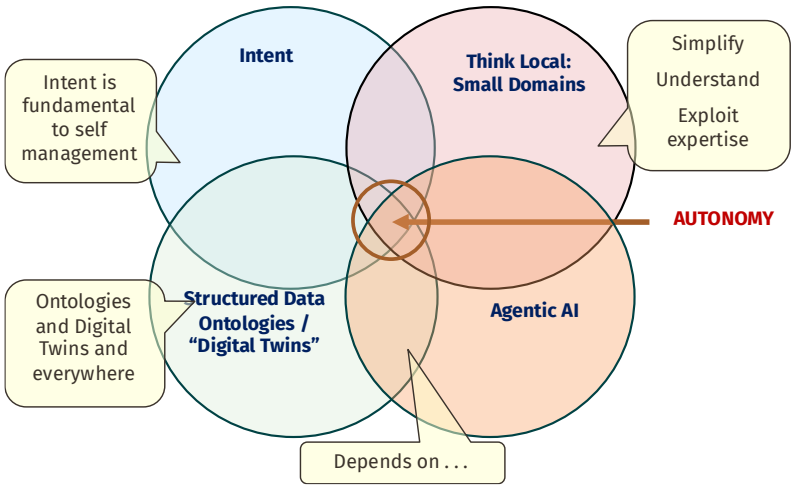

Successful network autonomy rests on four pillars, which are themselves inter-dependent:

1. Intent and control theory provide the foundation, and must be architected into underlying systems
2. Agentic AI provides a powerful new tool for both finding intent and understanding root causes
3. A "domain first" approach anchors operations and problem solving in deep expertise (critical for Ontologies, Intent and Agentic logic)
4. Structured data – including "ontologies"; "Digital Twins" provide Agents a solid foundation of highly relevant data and relationships.

Source: Appledore Research

Collectively these four pillars suggest that Agentic AI infrastructure for network operations – and especially within technology domains – is neither a task for CSP DIY nor for AI generalists, but rather for a relatively handful of players committed to a long-term roadmap centered on network domain expertise. Unsurprisingly this is a nearly identical conclusion to that we published with respect to more generalized Autonomy best practices, here. Such a separation of concerns between CSPs, Domain Technology Experts and AI experts results in a market structure more like other, more mature industries – in which production is dominated by firms with the scale to specialize and invest over the long term; with supply

chains from specialists further down the chain. So, CSPs deploy and customize use cases; domain experts deliver domain autonomy coupled with enabling Agentic infrastructure; and AI firms and open-source projects deliver software componentry.

While outside of the scope of this paper, it's worth noting that one layer up, at the cross-domain or service layers, the picture changes; since expertise in any one domain is less necessary, while the ability to abstract and reason across them becomes paramount. This is the model endorsed by Appledore, the TM Forum, and pursued by several CSPs. Even Nokia, the sponsor of this paper, treats its domains and cross-domain Autonomy and Agentic AI infrastructure as distinct product lines[1].

If we stand back a few hundred feet, we can see that many of the leading players are pursuing similar directions. We take this as validation and consensus of an effective path forward. We also see that it fits strongly with work that Nokia has been undertaking for years—from before the recent rise of AI.

First, this path depends on a proper intent based closed loop, which Nokia, across organizations, has endorsed for many years. Second, it depends on richly decorated topologies that allow performance information to be related to networks and services and service components and nodes to each other. Similarly, Nokia has been building these out for many years. Finally, it builds on experience gained with ML (predictive AI) where Nokia, in its IP and optical businesses, has long employed such algorithms drawing on Bell Labs research, but made effective through deep operational (IP and optical domain) knowledge.

The goal: Autonomy. Autonomy's killer apps: Agents **(and Technology-Fluent Agentic Foundations)**

> ***Network autonomy and Agentic AI are tightly intertwined. CSPs' goals are (broadly) vastly improved productivity, agility and cost – and the method is true, intelligent Autonomy. Agentic AI is the killer app that promises to help deliver on this goal. - Grant Lenahan, Appledore Research***

The critical point is this: **nothing in Agentic AI changes the fundamentals, nor the importance, of autonomous network architectures**. But it is a potentially world-changing jump in algorithmic power to accomplish this goal.

The following diagram illustrates the relationship between Agentic AI and Autonomous Networks, and Control Theory, using a standard "textbook" view of a control loop. One key point is that the two are tightly interlinked, and in fact Agents depend on an intent-based, control-loop software foundation. A second take-away is that Agents are best applied to very specific, well-defined tasks (where tasks and data requirements can be understood and therefore defined, and outputs may be clearly measured for trust). The "system" as textbooks call it, is in our case a telco IP network – with performance data being collected by the "sensor"; irregularities identified and remediations calculated by the "controller"; and

[1] Nokia's cross domain solution is named "Digital Operations". Its input is abstracted information from myriad domains.

Agents performing myriad tasks along the way. Nokia's NSP is an example of a controller (with some sensor capabilities as well) – specific to IP networks.

Figure 2: A control loop, with Agentic AI examples overlaid

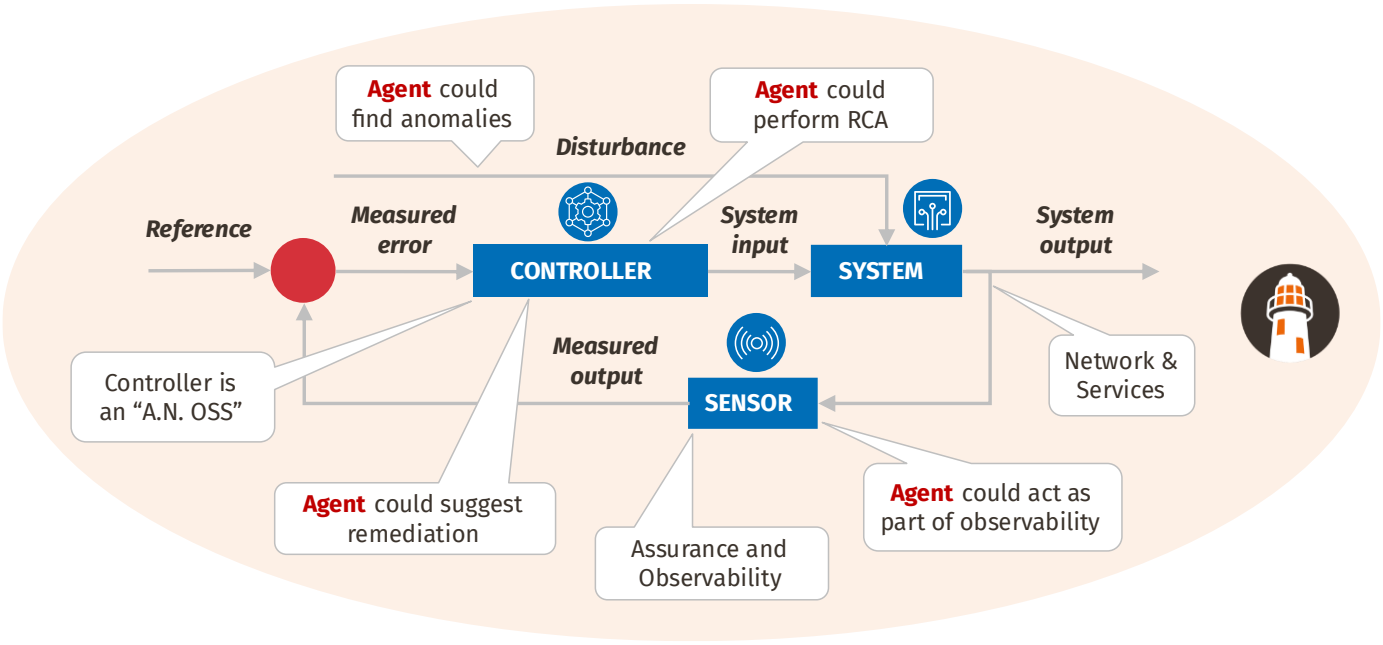

Source: Appledore Research

Control loops are established engineering theory and are how myriad systems from rocket control to robotics works. With the addition of Agentic AI, the goal is to apply similar systems to networks. And IP (or RAN, or…) network experts are our domain's rocket scientists.

Review of recent research: What makes for successful Agentic AI in Network Operations?

Network Domain AI differs significantly from general purpose AI

Agentic AI in highly technical domains differs in one significant way from the highly visible public AI solutions available commercially – whether from OpenAI, Anthropic, Google or myriad others. In the public AI sphere, the domain is unspecified and unknown. The nature of a topic or question is unknown; the user is almost by definition not an expert in what they are asking; and the user does not necessarily have a large corpus of relevant research. Therefore, general purpose AI is trained to be as broadly capable as possible, on what is referred to as "the pile." The quality and vague nature of "the pile" Is one of the chief sources of errors and hallucinations.

Commercial AI deployments by large institutions (banks, government agencies, etc.) modify this significantly; with the majority of training data being internal and often proprietary. Without belaboring the point, using internal vetted data delivers both security, and also far more accurate and relevant data.

As we consider specific technology domains within CSPs, this trend continues and sharpens. Suddenly we're at a point of specialization where it is possible to reduce the volume of data to that which is most

relevant, while at the same time building in important implications and dependency / relationship data. The result is an ontologically structured data set, which in many cases is anchored via a deep knowledge of how a specific domain technology is built- including network topology (both physical and logical); equipment configurations and port/PCB/chassis relationships, service topology, and configurations. For example, in IP it is necessary to understand how services are mapped to LSPs and LAG groups; in RANs we must understand the relationship between moving endpoints and cellular antennas – as well as the complex and uncertain propagation environment. In the end, these are in no way interchangeable, and good models and deep understanding mapped to such models provides a huge and dependable source of inference.

The result is better answers, fewer hallucinations, simpler training, faster operation, lower power consumption, and lower cost – just for starters. And this is one important reason why domain expertise may be the most important characteristic of an agentic AI infrastructure supplier for a particular network domain.

Research Findings and An Emerging Consensus

While AI has become the buzzword *du jour*, Appledore is simultaneously convinced that this buzzword has wings—the potential, if properly understood and deployed, is enormous. Similarly, we are convinced that the best way to apply and integrate AI into network operations is less mysterious than it may at first appear (see above). Consequently, we have spent the last year talking to leaders in the industry on both the supply in demand sides, as well as academics and industry organization leadership (TMForum, ETSI, MIT, Mplify Alliance).

Our first broad-scale, primary research project in this field, completed at the end of last year, aimed to understand the state of agentic AI deployments to support *network operations*, and especially to understand the leading use cases and the levels of productivity and other improvements that were being realized. We actually learned something very different: yes, there were use cases being deployed, but it was very early days within the network domain[2]. While at first we were slightly disappointed, we concluded that this was a very positive finding for the industry overall: there was huge progress being made in terms of enabling infrastructure – and importantly, many organizations, both suppliers and CSPs, were coming to very similar conclusions as to "what makes for effective and efficient Agentic AI in network operations." The pair of major reports that resulted from that work are here and here. We summarize them in the next section.

[2] The situation in NETWORK operations is very unlike that of the customer experience and billing/BSS operations, where AI powered Agents are being widely deployed at scale, both taking on the tasks of representatives, but also often making representatives hugely more productive. In this environment, apparently, the huge piles of information, unstructured questions of consumers, and often open ended "resolution" options make the proven public AI model highly relevant (except, with internal, relevant data). There scale is already meaningful and poised to continue growing.

Success factors for Agentic-augmented domain autonomy

"If I have seen further, it is by standing on the shoulders of giants," – Sir Isaac Newton

Dozens of experts, across suppliers and CSPs, participated, often in depth, in our research this past year. While details varied somewhat; the consistency and commonality of direction was striking. Building this input, we are able to summarize and structure the critical components and success factors for effective and efficient agentic AI in the operations of network domains.

Some will have immediate impact, such as deploying AI-based natural language processing as a vastly improved UI to documentation, data, and systems in next-generation "OSS" (NAS) systems. But the majority are foundational investments that make AI more accurate, more efficient, more effective, and easier to program. Figure 1 above, and the table below expand on these findings. Be aware, all are moving targets: all of them will become more sophisticated with time, and only a handful of suppliers have tackled all.

Figure 1 (above) provides a visual representation including the critical point that these four pillars are interdependent. The point is: "Agentic AI" is not the answer, rather it is a critical part of a bigger system. Example: you can't have effective autonomy without intent. Similarly, without autonomy (and the software platform), Agentic AI has no structured paradigm on which to operate. don't re-invent the wheel: build on the momentum of network autonomy.

Table 1: A Consensus Emerges – The Components of an Effective Agentic Platform for network Domain Operations

Foundational Capability	Commentary / Why? Important Points
Intent Based Control Loop	• A properly designed intent-based control loop provides many of the necessary foundations for Agent to operate – from defining desirable objectives ("intent") to having algorithms to make well constrained changes. Interestingly both depend on a very similar ontology to work effectively and be sufficiently flexible. Build for one, build for the other.
Domain-specific, structured data	• Building blocks of an ontology, and often including inventory and topology data, as well as configuration and service state. • Recognition that "big data lakes" were not the answer, and what is needed is more focused, well structured, and valid data on which to reason. Quality – vetted by domain expertise -- over quantity.
Ontologies	• An Ontology is "a formal framework that organizes knowledge by defining entities, their attributes, and the relationships between them" • More specifically it combines not only network data, performance data, topological data, configuration data – and the dependencies, relationships, and implications between them, so that meaningful causality and therefore resolution may be sped and guided. • Collectively structured data, topology, Ontology and Digital twins are a huge part of why we strongly suggest that **domain specialists** (not, for example, hyperscalers[3]) build domain Agentic foundations.

[3] Another finding of our research, sufficiently tangential to be a footnote here, is that the industry seems to be recognizing the role of hyperscalers and AI giants as infrastructure suppliers, but not experts in telecoms, networks, network data, or the nuanced inter-relationships that support good operations. Just as auto OEMs don't make steel or tires, we are seeing a maturation in both network operations and Agentic AI with a layered supply chain.

Foundational Capability	Commentary / Why? Important Points
Digital Twin(s)	• Digital Twins provide a foundation to ask questions of; simulate and emulate myriad questions in a safe environment that will mimic the real world. While there is no (not yet) 100% digital twin; various twins can provide reliable test environments for new traffic loads, configurations and other questions. • Digital twins have been used for decades under different names. Some of their leading use cases are circuit simulation; radio propagation simulation; and ballistics test and computation (rocket control). • We view the combination of the above data, structure, relationships and topology (a form of both data and relationships) as, when cumulative and structured to support specific types of reasoning, a Digital Twin. Note most digital twins are at least somewhat task specific – intended to simulate configurations, capacity and performance or other specific and limited queries.
Natural Language UIs	• Natural Language AI Agents capable of automating data enquiry; problem investigation, and simplification of access to complex system capabilities. Objective is to make simultaneous jumps in both productivity and efficacy by automating mundane tasks. Often front ends a combination of the system (e.g.: NSP); system data; and related data such as network equipment documentation
Agentic Fluent APIs	• Enable – often using MCP – the exchange of context rather than specific queries, and both opens (and gates for security) the system and its capabilities to external innovation, whether by the CSP or a 3rd party.
Agentic AI "studio"	• Enables definition of new tasks; running of multi-agentic workflows (often self-created) • Enables definition of guard rails and security/access rules • Enables definition of data sources and I/O rules • May be used by the system supplier (initially: likely); by the CSP; and eventually by third parties
App Sharing Environments ("App stores")	• While rare today (app stores need libraries of apps after all), many players acknowledged the value app exchanges, leveraging Agentic-native (e.g.: MCP) Apis on innovation and the sharing of good ideas. We expect to see these emerge over time, with maturity and once risks are better understood.

Source: Appledore Research primary market research

A relevant and related view of the layers of data and capabilities supporting agentic operations is presented in this research note. Interestingly, this was published and distributed just ahead of the 2025 DTW show in Copenhagen to document Appledore's position and engage suppliers and CSPs. The fundamental point was that worrying about which LLM, for instance, was being used, or turning AI loose on existing general data lakes, was missing the point. Rather, **success will depend on relevant data, structured to support and simplify AI reasoning.** Structuring occurs at several levels from well-established AI database characteristics through network-specific ontologies. The data, ontologies, etcetera presented above, along with data structuring tools (typically provided by 3rd parties) such as vector databases, are all evidence of this illustrative layering's fundamental truth.

NSP and Digital Twins

It's worth taking a moment to look at the journey Nokia has been on for its NSP "digital twin". For many years, Nokia has had a structured copy of the network including equipment, configurations, state, overlaid services, and other relevant data. Such data is necessary for automated management of the IP network. Several years ago, this was upgraded to become a "what if?" environment; specifically to understand the implications of service and load changes across the network as configured at that point in time. Last year, Nokia launched the first version that it refers to as a digital twin. This version supports emulation and allows changes to the network—with an emphasis on configuration changes—to be tested. From this we draw 2 takeaways: first, that the NSP team has been at this game for many years; and second, that they treat the word digital twin quite seriously and reserve it only for their most advanced iteration.

In summary, over the past 18 months the industry has gone from viewing AI as generic magic, to a **relatively nuanced view of the infrastructure necessary to support far more effective agentic AI**. Similarly, it has begun to distinguish between building AI models and databases, and focused on our industry's core competence: understanding the data and relationships that populate those layers of data. The latter demands very specific network expertise – and often long experience. Finally, the relationship between agentic AI an intent based control loop has been widely agreed upon at least by those at the leading edge of such architectures. All of these mean there is a relatively well understood and safe best path forward.

One Agent, Many Problems at Once

A large broadband operator's NOC was drowning in alarms — thousands a day — yet the questions that mattered went unanswered: *what's the real root cause? Which services and subscribers are affected? Is redundancy still holding? Did last night's change cause this?* Worse, traditional alerts leave some thorny problems undetected: threshold monitoring missed "quiet failures" entirely; a lost redundancy path could stay invisible until a second fault took customers offline. Engineers stitched the picture together by hand across tools, often dispatching several teams before the true cause emerged — and every incident still had to be written up in a fixed format.

The operations team – capacity limited – performed traditional operations actions as separate problems: noise here, root cause there, impact analysis, reporting, next steps. A leap forward, however, occurred when an AI troubleshooting agent was placed on top of the controller data and looked more holistically — detecting the incident, correlating events, enriching them with topology and operational knowledge, determining service impact and root cause, then returning a formatted report with recommended actions, query-able through dashboards or natural-language MCP interfaces.

AI, with its tremendous ability to sift through both large volumes of data and to correlate data from multiple topics, gave the speed and scale to make this jump. Yet AI alone was not a panacea- rather it, in turn, depended on and benefitted from learnings and structured data built into NSP by years – maybe decades -- of engineering and operational experience.

That was the a-ha insight. These were never separate problems; they were one missing layer — the reasoning between a raw alarm and a decision. Close that gap once, and faster MTTR, cleaner tickets, and earlier catches on silent, cascading failures all flow from the same place.

Choosing a supplier for Domain Specific Agentic AI

From where should you get your agentic AI for a network domain?

It is becoming clear that for AI and Control Loops *that diagnose and touch specific network technologies and topologies*—**domain expertise is paramount**. The underlying technologies, such as vector databases and generic LLMs can be and will be frequently purchased (or in the case of open source, embedded).

The closer you get to the network, the more domain expertise counts for. Reinforcing this, in control theory a fundamental maxim is to act as close to the source as possible. Conversely, as you rise up the layers to cross domain network services, end user services, and ultimately to BSS, the greater the level of abstraction will be and the more important it will be to understand more generic business and service-related objectives and operations. Consequently, the ideal supplier for your service layer agentic AI infrastructure and your IP domain agentic AI infrastructure are likely to be very different (and yet both experts in their own areas).

Network Domain Agentic AI supplier categories and strengths / weaknesses

AI Supplier Category	Germain Characteristics
Domain Experts (e.g.: IP equipment providers, RAN providers, IP operations specialists)	• Large players have the scale to invest across many CSPs and support a product for the long haul • Have already spent years wrestling with the details of AN control loops and intent (both of which spill over into Agentic AI) • Focus is entirely on one domain (even large players have multiple products and teams, each with focus) • Know the operating characteristics, Metrics, useful Ontological structures, and how to model topologies - but only for their specialty • Typically, not database suppliers, LLM suppliers • Standards (BBF, TIP, others) are making multi-vendor more and more practical.
NAS/OSS ISVs (e.g.: Orchestration firms, full-suite providers)	• Typically operate across domains • May have scale at cross-domain layer. • Typically, insufficient scale and expertise at domain layer (but notable exceptions exist, see our domain profiles) • . . . And are consequently skilled and natural fits for cross-domain Agentic AI in which domain specifics are, by the nature of ANs, abstracted • Some have developed expertise in one or more domains, emphasizing vendor neutrality.

AI Supplier Category	Germain Characteristics
Hyperscalers (often AI Giants) ***e.g.: Google, Anthropic, AWS, Microsoft, Google, Oracle (as database and AI firm, not OSS organization)***	• Often leaders with LLM based products that can be later customized for specific industries (e.g.: a network domain) • Often have useful tools on which others build – databases, vectorized databases for accelerating AI, "fabrics" for structured data – all of which may be used by others • Runtime environments (cloud) • Scale in operations, training • Leaders in true plumbing infrastructure selling to all buyers, specializing in none
DIY (CSP IT shop)	• Historically the source of some of the longest-lived legacy headaches in the industry • Even the largest CSPs have questionable scale (divide/1) • Cannot benefit from multiple CSP deployments learnings and operating models • CSPs get precisely what they want (or think they want) • Caution: history says long term support will kill such products. • All other industries have shifted to a modest number of large and well-funded players – including the first autonomous domain- cloud native datacenters.

All mature industries consolidate to scale. Auto manufacturers no longer make their own tires or steel; only a handful of computer firms make chips; "phone" companies no longer make their own equipment. Rather they rely on specialists with the ability to focus and excel within a narrow confine. Scale wins because scale can afford ongoing investment and expertise. And done right scale wins because of commonality and consequently far easier integration.

"History doesn't repeat itself, but often it rhymes" – Mark Twain

We (Appledore) are increasingly convinced that the combination of autonomy and agentic AI demands a level of system engineering and ongoing investment that can only be shouldered by a relatively small number of players. We have published extensively over the past few years how concentration is occurring in the network automation software component of our industry. We first saw the practical evidence of this in the market for cloud native software automation, where individual suppliers' versions of "NFV", and even their attempts at cloud native, ultimately could not compete with the larger players (often based on open-source projects) with huge resources and a large client base. We are convinced that a similar concentration will occur in the markets for autonomous network domain software / network operations focused agentic AI, and that it is the right decision.

There remain, in this scenario, two important issues that concern buyers. First is the historic lack of success at true multi-vendor parity when any equipment vendor was involved. Yet we are seeing far stronger standardization in everything from IETF to broadband forum, TIP and other organizations. Second is the concern of closed systems and lock in. Yet modern open methods such as the APIs that are being widely adopted along with open-source foundations and more modular—if not necessarily microservices architectures—mitigate this concern. Nokia has been a leader in open API access to essential data and functions, mitigating this concern as well.

Summary & Recommendations

Agentic AI promises to be a tremendous boon to network autonomy, and also to basic operations productivity even where autonomy is not yet a reality. In this paper we have looked at the rapid evolution of thinking in terms of how best to achieve effective agentic AI—with a specific focus on the unique environments and challenges of network operations.

Based on the data above and our read of both the marketplace and what will work long term, we want to leave the reader with a few summary recommendations. We have written these with the presumption that the reader is in a leadership position at a CSP tasked with effectively harnessing agentic AI to spur automation and long-term success in one or more specific network technology domains.

First, recognize that both the software and the operations processes for long term network **autonomy and those for agentic AI are inextricably linked**. In fact, a properly architected autonomous network controller is a foundational enabler of agentic AI. *Therefore, focus on supply choices with BOTH autonomous fundamentals AND a strong Agentic infrastructure story*

Second, recognize the importance of domain expertise, domain data, network and service metrics, topologies, configurations, and operational actions are – and how inter-related they are. The killer expertise in crafting agentic AI for a routed IP network is understanding the routed IP network. The AI can be bought. *Therefore, in network domains prioritize network domain experts.*

Third, focus on the enabling platform before the fully baked use cases. Once that platform (with a focus on data and embedded ontological knowledge) is in place it is faster cheaper and simpler to build use cases—and they should operate far more effectively, using less power, and at less cost. *Therefore, build a platform – a system – on which you can innovate for decades.*

Fourth, remember that domain level autonomy is compartmentalized and **operates as a modular component of a much larger end to end autonomous network**. A foundational principle of control theory is to act as locally as possible; then abstract and remit larger problems to the next layer up. The bottom line is not to try and solve problems beyond the domain and to make sure that everything can be packaged and abstracted to another specialist system which looks across those domains. *Therefore, build to an end-to-end architecture and work to push decision making down to the lowest level possible.*

Fifth, Choose your supplier carefully. We recommend choosing a network savvy specialist; a domain expert for a specific domain; a cross-domain specialist to link them together. Built to a corporate end-to-end autonomy blueprint.

Sixth, choose the supplier with the scale and the breadth of experience to be a valuable partner and help you innovate for the life of the solution—because that may be one or two decades.

In the end the transition to autonomy and the introduction of agentic AI will be marathons not sprints. Yes, we need some quick hits, but the greatest value from those quick hits is proving what works well and what does not.

For those who are curious to learn more or want to discuss or challenge any of these points we encourage you to reach out to Appledore Research and of course to our sponsor for this document,

Nokia. We have also linked quite a few original independent research reports within this document and present a bibliography below.

Bibliography and Further Reading

1. Agentic AI In Autonomous Networks – Supplier Review
2. Agentic AI – Advancing Autonomous Network Operations
3. CSP Survey
4. The AI Model Illusion
5. AI in the autonomous control loop
6. Best Practices in Network Autonomy (Parts 1,2 and 3)

Further information, specifically on capabilities of Nokia NSP see:
https://onestore.nokia.com/asset/215425

About the Author

Grant Lenahan provides a unique combination of management and technical acumen, combined with 30 years of successful innovation in both technology and business models. Prior to co-founding Appledore, he served in the office of CTO for Ericsson, after many years with Telcordia Technologies. Through his career, Grant has specialized in transforming telecom software and service businesses in the face of dramatic market and technology shifts, positioning the businesses for survival and growth in new environments.

Grant has deep experience in understanding market and technology shifts, and the consequent opportunities and threats that these shifts create. He has consistently guided the Telcordia and Ericsson software product portfolios to thrive on these changes. Grant holds a Bachelor of sciences from Drew University and an MBA (SM-Management, SM-Engineering) from Massachusetts Institute of Technology.

Email: grant.lenahan@appledorerg.com

Insight and analysis for telecom transformation.

@AppledoreVision

Appledore Research

www.appledoreresearch.com

info@appledorerg.com

+1 603 969 2125

44 Summer Street Dover, NH. 03820, USA

Appledore RESEARCH