

Building success in agentic AI operations

Agents aren't the secret: broadband expertise is.

Author: Grant Lenahan, Partner & Principal Analyst



In partnership with

NOKIA



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

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

© Appledore Research LLC 2026. All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means – electronic, mechanical, photocopying, recording or otherwise – without the prior written permission of the publisher.

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.

Appledore Research LLC recognizes that many terms appearing in this report are proprietary; all such trademarks are acknowledged, and every effort has been made to indicate them by the normal USA publishing standards. However, the presence of a term, in whatever form, does not affect its legal status as a trademark.

Appledore Research LLC maintains that all reasonable care and skill have been used in the compilation of this publication. However, Appledore Research LLC shall not be under any liability for loss or damage (including consequential loss) whatsoever or howsoever arising because of the use of this publication by the customer, his servants, agents or any third party.

Publish date: 14 September 2026

Cover image: Perplexity

TABLE OF CONTENTS

FOREWORD & EXECUTIVE SUMMARY 2

DEFINING SUCCESS IN NETWORK AUTONOMY 3

SUCCESS FACTORS FOR APPLYING AI TO AUTONOMOUS NETWORKING..... 4

 Theory & Market Findings..... 4

 The Foundation of Autonomy: Control Theory..... 4

RECENT RESEARCH REVIEW: WHAT MAKES FOR SUCCESSFUL AGENTIC AI IN NETWORK OPERATIONS?..... 5

 Research Findings and An Emerging Consensus 6

 Success factors for Agentic-augmented domain autonomy 7

INTEROPERABLE AND AFFORDABLE AI 10

 The practical issues in enabling innovation and efficient scale 10

 Affordable (and potentially more dependable) AI.....10

 Interoperable AI 11

CHOOSING TECHNOLOGY SUPPLIERS 12

 Measuring Altiplano against these objectives.....14

SUMMARY & RECOMMENDATIONS..... 17

BIBLIOGRAPHY 18

Foreword & Executive Summary

This paper examines success factors in the move to Agentic AI-augmented network autonomy, and look at the specific factors in the broadband access market. We conflate these four concepts (Autonomy, Agentic AI, Network, Broadband access) intentionally – because success demands a system implementing all four. We summarize the foundation (and theory) of network autonomy; include the latest research and industry experience with Agentic AI applied to autonomy; and focus specifically on the branch of Agentic AI applied to **network operations**. This last point is also critical, since the application of AI and Agentic AI to customer and billing operations is significantly different – both in nature and maturity.

To support this paper, Appledore draws on our extensive primary market research corpus; performed over the past year, and including dozens of deep briefings with CSPs, network suppliers, ISPs and hyperscalers. From this we have identified priorities, as well as a broad consensus about what makes Agentic driven network autonomy successful. The original research is listed in the Bibliography.

In the end we hope to accomplish the following, along with supporting facts:

- Establish **success factors** for Agentic-driven network autonomy
- Discuss the **foundational components** of a state-of-the-art solution
- Analyze the **strengths and weaknesses of various AN and Agentic AI supplier segments**;
- **Summarize** the state and direction of **Nokia's Altiplano** product against these criteria.

Autonomous network automation will typically be delivered by a federation of autonomous domains and technology modules, with each domain (and cross-domain function) operating autonomously, yet in concert with others. In each domain and at each layer, there may be different suppliers with distinct specialties. This is to be expected. In these complex and evolving systems, CSPs (buyers) must negotiate the trade-offs among DIY, single-vendor solutions, and best-of-breed architectures.

Regardless of choices, **it is becoming increasingly evident that the most important expertise in both Autonomy and in created AI solutions is the domain under management**. There are huge differences between broadband access; 5G RANs; cross-domain management; and – to extend the point – medical automation. And in each and every case, as autonomy is applied and AI simplifies “programming”, the essential skill behind success is domain knowledge.


In the end, we hope to explain why Nokia, in Altiplano, has made the design and investment choices that it has; how those map to our independently generated best practices, and why Appledore regards Altiplano as a strong choice for CSPs in implementing Agentic AI in the broadband access segment of their Network Autonomy journey.

Defining success in network autonomy

While network technology gets all the attention, operating expenses dominate the financials of all network operators. Autonomous networking is merely the latest in a series of technologies intended to automate and improve the efficiency of networks and the productivity of the workforce that operates them. Similarly, AI, and in particular agentic AI, is simply the latest generation of algorithms to help power network autonomy— sifting through data; finding correlations; understanding root causes; making predictions; suggesting resolution; and sometimes taking action.

For over a decade Appledore has periodically amalgamated CSP financial operating data to ask: “where does the money go?” and “what really matters?”. As the chart shows, the lion’s share of total operating expenses are network operations; by contrast, capital (mostly depreciation) is a small fraction. And the percentages shown have changed little over the past decade.

The implications are clear: first, operations must become more efficient; second, surging traffic and network complexity are exerting upwards pressure on operations costs. We can expect that in the future, Autonomy and Agentic AI will be essential for telco’s operations experts to accomplish more with the same resources.



Where money goes in network operators today
Source: Appledore Research

“Money is in the Poles and the Holes”

Broadband access networks may be the most distributed and labor-intensive operations of any segment of the network. There’s an old adage that “money is in the poles and the holes.” Mobile RANs may come in a close second, but whereas there may be 10s of thousands of cell sites in a mobile network there are likely to be millions of fixed broadband endpoints (and some of them will be cell backhaul!). This huge scale and geographic distribution point directly to complex NOC operations and complex expensive truck rolls, both of which need to be simplified, improved in precision, and cost reduced.

“The beauty of properly implemented autonomous operations is that it is not only possible, but almost necessary, to have at the same time improved accuracy; lower cost; faster time to repair; and as a consequence, better customer satisfaction” – Appledore Research

Properly implemented autonomy, (including a properly implemented closed loop, separation of concerns, abstraction, and intent) is naturally far less maintenance and integration intensive. Similarly, if the platform for AI is properly designed as a system, rather than having AI “bolted on”, it too becomes maintainable, upgradable, extensible, and an open platform for innovation.

Agentic AI may be the missing piece in the autonomy puzzle. For years suppliers and CSPs have struggled with algorithms to implement intent; algorithms to understand root causes and implications; and myriad other tasks in a control loop. Agentic AI, **built on a foundation of domain expertise**, may be simultaneously simpler and more powerful.

Success factors for Applying AI to Autonomous Networking

Theory & Market Findings

In this section we look at what makes AI effective in network operations. We show that the most important choices that operators make are related to operations domain specific, rather than to AI choices themselves.


When Generative AI burst on the scene it appeared almost magical; a panacea that appeared to have extensive and detailed knowledge of nearly everything. As time progressed it quickly became apparent that GenAI also had drawbacks, making errors, and repeating false information – sometimes referred to as “hallucinating.” Both of these are a direct result of its training method and data; and therefore, preventable with high quality, relevant and well-structured data (as opposed to the general internet corpus referred to as “the pile”). ***No surprise that our research showed that possibly the most important foundation is valid, useful, structured data on which network domain AI can reason.***

For the better part of a decade the industry has been struggling with how to achieve higher levels of productivity and higher levels of automation, cycling through ever more sophisticated concepts including cloud native, closed loops, intent intent-based networking, and most recently, autonomous networks (which depend on both closed loop and intent). ***Nothing in Agentic AI changes the fundamentals, nor the importance, of autonomous network architectures.*** Rather, Agentic AI promises a potentially world-changing jump as they supplant some existing autonomous networking algorithms establishing the next generation of complex reasoning. However, success depends on a good system design – and that design on a foundation rooted in the task at hand.

The Foundation of Autonomy: Control Theory

The following diagram illustrates the relationship between Agentic AI, Autonomous Networks, and Control Theory, using a standard “textbook” view of a control loop. The three are tightly interlinked, and Agents rely 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 network. Performance data is collected by the “sensor”; irregularities identified and remediations calculated by the “controller”; and Agents performing myriad tasks along the way. ***Nokia’s Altiplano is an example of a controller (with some sensor capabilities as well) – specific to Broadband Access networks.***

Figure 1: A control loop, with Agentic AI examples overlaid



Source: Appledore Research

“If control loops and intent didn’t already exist; Agentic AI would have to invent similar systems. Fortunately, many domain specific controllers do exist . . .”

Control loops are established engineering doctrine and underpin myriad systems from rocketry to robotics. With the addition of Agentic AI, the goal is to apply similar systems to networks, with network experts acting (in the case of Altiplano) as the broadband access domain’s “rocket scientists”. Going one step deeper, intent defines the criteria for success, providing a target for algorithms to achieve. If control loops and intent didn’t already exist; Agentic AI would have to invent similar systems. Fortunately, many do.

Recent research review:

What makes for successful Agentic AI in Network Operations?

If we look back between 12 and 24 months, there was a general recognition that AI and later agentic AI had huge potential to telecom operations; however, the notion of *how* was vague. The first successful applications of agentic AI at scale in telecom were in customer service and billing—where the lessons learned from public AI were highly relevant; the ability to chat or interact with humans, combined with the ability to train on huge quantities of data in order to provide answers to a very wide range of unstructured questions. Since, in that environment, a customer service agent must be a generalist – existing LLM based solutions of the time were directly applicable. Hence, initial scale deployments of LLMs/AI occurred in billing and customer service.

Beginning about 12 months ago, we began to see more serious discussion of applying AI to *network operations*. Here the problems are very different: the problem space is narrower; and the level of complexity within it is significantly greater; and the stakes of misdiagnosing or mis-configuring a network are greater. Precision and deterministic operation are essential.

Research Findings and An Emerging Consensus

Appledore was convinced that Agentic AI could dramatically improve efficacy and productivity in complex network operations. We wanted to understand the real state of the market: what were the use cases? What were the productivity improvement metrics? Where was AI being used at the greatest level of scale? What were the success factors?

We spent the past year-and-a-half talking to leaders in the industry on both the supply and demand sides, as well as academics and industry organization leadership (TMForum, ETSI, MIT, Mplify Alliance).

Our first coherent view of industry status and direction began crystallizing at 2025's TM Forum "DTW" in Copenhagen; where, behind closed doors, we had highly consistent discussions with myriad vendors and operators. **A common recipe was emerging:** success demanded a focused set of highly relevant, authoritative, and structured data, embedding knowledge of the network and operations, that could be used by AI agents to quickly come to relevant and correct conclusions. In fact, after the show we postulated a three-legged stool composed of 1) agentic AI, 2) well-structured data, and 3) the evolution of inventory and topology which, together formed the foundation for digital twins.

This 3-legged stool in turn depended on **two critical foundations:**

1. an **intent based autonomous control system** built into the next generation network management software
2. **deep knowledge of network operations themselves;** most often exploiting expertise in a particular technology domain. (for example, Broadband Access or 5G RAN)

Our first broad-scale, primary research project in this field, aimed to understand the state of agentic AI deployments to support *network operations*, 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. 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.** Moreover, similar to our experience at the TM Forum event, suppliers and CSPs were coming to very similar conclusions as to "what makes for effective and efficient Agentic AI in network operations." We summarize them in the next section and link them in the bibliography.

Microsoft Research: Data Quality Beats Model Size

The following comes from a **Microsoft Research** paper titled "textbooks are all you need", and **underscores the importance of concise, domain-expert, curated data.**

When models options are all credible, the decisive factor in an LLM's performance is the data, not the model. The quality, relevance and curation of the training data do more to determine real-world results than which model is chosen — as Microsoft Research showed with phi-1, a model trained on curated "textbook-quality" data that outperformed code models many times its size.

"We introduce phi-1, a new large language model for code, with significantly smaller size than competing models... Despite this small scale, phi-1 attains pass@1 accuracy 50.6% on HumanEval."

— S. Gunasekar et al., Textbooks Are All You Need, Microsoft Research, 2023 (arXiv:2306.11644).

Success factors for Agentic-augmented domain autonomy

Building on our research, 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.

“The majority of today’s investments are foundational (“platform”) investments, often built into domain controllers or cross-domain systems, that make AI more accurate, more efficient, more effective, and easier to program.”

The table below expands 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 fully implemented all.

With these platform capabilities in place, successful Agentic use cases become far easier to develop and more accurate/effective.

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.
Role of Mass market AI firms and Hyperscalers	Another finding of our research is the general recognition 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 automotive OEMs no longer make steel or tires, we are seeing a maturation of the layered AI supply chain.
Domain Expertise	As in other areas of AI - the critical needs is domain knowledge: what are the problems; the topology; the secrets; operational requirements. The AI part them becomes easy. The next three build on domain expertise.
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” are not the answer, rather, success depends on more focused, well structured, and valid data. 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) build domain Agentic foundations.

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 a Digital Twin. Note real world digital twins are task specific – intended to simulate configurations, capacity and performance or other specific and limited queries. (Again, driven by domain expertise and operational needs)
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. Natural Language Processing often augments 1) the system (e.g.: Altiplano); 2) system data; and 3) 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. • A wide range of APIs are critical to open-ness; and open-ness is critical to innovation. It unlocks the ability for anyone (supplier, CSP/operator, third parties) to innovate, and potentially to then share that innovation (see App/Agent sharing environments below)
Agentic AI framework supports app/user case innovation	• Enables definition of new tasks; running of multi-agentic workflows (often self-created) • Enables selection of guard rails and security/access rules • Enables definition of data sources and I/O rules • Safe access via MCP • Allows app/use case development by Nokia, CSP and 3rd parties, spurring innovation
App Sharing Environments (“App stores”)	• Application and Agent sharing environments are all about open-ness and facilitating innovation. • 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 and believe that they can unlock huge potential (the “long tail”, among others).

Foundational Capability	Commentary / Why? Important Points
AI is not always the answer <i>Key point – its not either/or – its “both”</i>	• We were surprised and pleased to hear an (unpopular) recommendation that we have made for years going mainstream: AI is not always needed, nor desirable for all tasks. Often a combination of LLMs/Agents and other methods, each used where it is well suited, is ideal. • Non-LLM methods and deterministic automation, including sophisticated cascading rule sets and even simple scripts are often cheaper, less power hungry, faster, and critically These non-LLM/Agentic solutions can be more deterministic and precise than LLM-driven solutions. • The combination of both AI and deterministic automation enables trustworthy, accurate and token efficient solutions. • Example: LLM based agent finds likely root causes. Once identified a script or rule set enacts a known, correct configuration change.

Source: Appledore Research primary market research


In [this research note](#), we present Appledore’s view of the layers of data and capabilities that support agentic operations. Part of the focus was to re-focus attention on data and domain expertise. Interestingly, this was published and distributed just *before* the 2025 DTW show in Copenhagen and was intended to 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.**

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 AI- friendly 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 and 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.

Real world application: Agentic AI for network troubleshooting and ticket handling

The following actual use case implementation illustrates modular (re-usable) Agents, and how they interact with the underlying domain expertise driven data tools embedded in Altiplano.

Nokia has worked with a Tier 1 operator in APAC to implement an agentic AI framework for troubleshooting network issues, both reactively and proactively. It covers multi-agent flows to solve end-to-end network issues leveraging Nokia Troubleshooting Agents from multiple network domains. The benefit for the operator are higher first-line resolution rates, incident qualification within minutes, reduced ticket handling costs, and improved customer satisfaction.



The Altiplano Troubleshooting Agent helps to cut through the alarm noise, isolating/localizing faults, pinpointing root causes and recommending corrective actions in fiber-to-the-home (FTTH) networks. The agent builds on Altiplano's digital twin data and fiber health analysis tools, and uses LLMs for reasoning and building human-like responses for customer care and support teams. The agent's capabilities are exposed via MCP (Model Context Protocol) to interact with the operator's agents.

Interoperable and Affordable AI

The practical issues in enabling innovation and efficient scale

The terms "interoperable" and "affordable" AI are terms suggested by Nokia. Yet we include them because they are concepts that we have emphasized for many years using different terminology. Quite possibly Nokia's terminology is better.

Affordable (and potentially more dependable) AI

Affordable AI ensure good value for money, whether in terms of power, heat, tokens or any other measure.

Conveniently, "affordable AI" often also brings along low error rates and near-perfect deterministic performance. There is no question that AI delivers a significant step forward for many complex problems. However AI, and in particular LLMs, exhibit some undesirable characteristics, which make them

undesirable – in both dependability and cost – for certain tasks within a process workflow. Specifically, LLM based Agents are nondeterministic; providing an output that may often be excellent and innovative but is neither consistent nor dependable. Second, LLMs, as they implement “deep search”, consume an enormous amount of compute power and resulting cost. Beyond direct costs, compute power means money in capital and money in electricity; both of which drive incremental costs that make scaling them hugely expensive. Both within network operations, and across other industries – the early management guidance to “token max” – and apply AI to everything, is being reversed as the true costs become apparent.

Appledore has made a similar argument for many years; specifically, that AI algorithms should not be used universally – but rather where they provide a meaningful improvement is speed, and automation, along with acceptable precision. We have argued and demonstrated that highly effective closed loop autonomy can often be achieved with very simple algorithms such as rules or much more advanced cascading rule sets. Both operate quickly and cheaply, and give exactly the same result every time. A great example where cascading rule sets outperform LLM based agents is configuration of a network port or port assignment – where the problem is well understood and there should be no variation. Rules do this. Agents might get, um, creative.

None of this is intended to say that AI is “bad”; but rather to suggest that it is appropriate for some tasks, but undesirable and unnecessary for others, which may be solved more cheaply and precisely. This is Nokia's concept of affordable AI and is a key part of how they choose to implement AI in their system—use deterministic tools in combination with AI models with proper guard rails to help reduce hallucination and get more reliable responses.

Interoperable AI

Interoperable AI effectively means AI that is open and easy to communicate with whether you are an onboard system and offboard system a third party or a telco domain expert. Within Altiplano, several design choices facilitate interoperability. First, all functions are exposed as APIs (often these are MCP interfaces). The same APIs that Nokia uses to build capabilities are available to your staff or your partners. Via these interfaces CSPs (or others) can connect chatbots or apps, via safe MCP APIs to system tools and data, facilitating open innovation, yet ensuring guard railed and consistent actions. Last but not least, Altiplano -- almost unique among commercial domain controllers—offers an application store; encouraging and simplifying the rapid adoption in distribution of good ideas. Appledore has argued that such “App Stores” ought to have been part of the industry’s move to cloud native and would likely have spread the adoption of leading practices. It's good to see one of the leading players in the segment taking both openness in collaboration to heart and putting their investment money on the table.

Choosing Technology Suppliers

Where should you get your agentic AI for a network domain?

It is becoming clear that **domain expertise is paramount** for AI and Control Loops *that diagnose and touch specific network technologies and topologies*. On the other hand, related platform components, such as vector databases and generic LLMs, can be and will be frequently purchased (or in the case of open source, embedded) from leading AI and database suppliers (Oracle, Microsoft, other).

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 agentic AI supplier for end-to-end service layer autonomy, as opposed to the ideal choice for your broadband access domain are likely to be very different (and yet both experts in their own areas)¹.

Below we list various segments of the supply side, along with their strengths, weaknesses, and germane characteristics. Note we include CSP’s “doing it yourself” (DIY).

Table 1: Summarizing network Domain Agentic AI supplier categories and strengths / weaknesses

AI Supplier Category	Germane Characteristics
Domain Experts (e.g.: Broadband Access suppliers; RAN providers,...)	• 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, TAP, others) are making multi-vendor increasingly 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.
Hyperscalers (often AI Giants)	• Often leaders with LLM based products that can be later customized for specific industries (e.g.: a network domain)

¹ This deserves a quick footnote: an independent line-of-business within Nokia designs, builds, and sells AI-driven closed loop autonomy at the cross-domain layer (Nokia “Digital Operations”).

AI Supplier Category	Germane Characteristics
<i>e.g.: Google, Anthropic, AWS, Microsoft, Google, Oracle (as database and AI firm, not OSS organization)</i>	• 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. Car manufacturers no longer make their own tires or steel; only a handful of computer firms make chips; “phone” companies no longer manufacture 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, where it can be spread over a large number of buyers/users. We have published extensively over the past few years how concentration is occurring in the network automation software (“NAS”) 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. A tangible example is Nokia’s own partnership with RedHat, whereas they previously built their own cloud native platform. 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.

In the technology domains, we believe that the domain technology firms, maybe joined by a tiny number of specialist ISVs – must and will win. 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 the primary software developer/supplier. Yet we are seeing far stronger standardization in everything from IETF to broadband forum, TAP 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.

Measuring Altiplano against these objectives

We believe Altiplano is on a credible roadmap path to implement all of our best practices for Autonomy and AI. To be fair, no systems today are fully mature and achieve a mature level of autonomy, let alone a mature agentic AI environment. Yet many are progressing rapidly down this path, making significant breaks from the legacy past, and have a very credible plan to deliver a proper Autonomous Networking platform combined with a proper Agentic AI platform to serve that autonomous loop.

Our assessment draws on extensive investigation; not only normal documentation but detailed 1:1 briefings for this project and for our ongoing series of market outlooks in the network autonomy research program.

We will now step through the major success characteristics one by one. We also refer readers to our existing Altiplano primary market research survey and analysis ([here](#)) and our upcoming profile (expected 3Q 2026).

Table 3: A Consensus Emerges – The Components of an Effective Agentic Platform for network Domain Operations


Foundational Capability	Commentary on Altiplano
Intent Based Control Loop	• Altiplano is built around a proper intent-based control loop. • Altiplano has been intent based for years and Appledore has talked directly with operators utilizing it and exploiting it to improve and simplify their operations.
Domain-specific, structured data	• See also “digital twin” and “ontologies”, below (all are inter-related) • Altiplano maintains a curated base of valid, relevant, information—based on increasingly rich and real time network data collection.
Ontologies	• Altiplano structures its data ontologically, using its deep experience and expertise in broadband access network operations. • Their ontology includes dependency information, and other associations that simplify and speed reasoning toward valid solutions. • This is specifically the kind of knowledge that can come only from domain expertise; codifying what is known by experts—and making it usable at scale.
Digital Twin(s)	• Altiplano contains digital twins (plural), since at today's level of technological sophistication, a twin must be designed with certain queries and outcomes in mind. • It is interesting to note that Nokia has been an advocate of digital twins across its domain product lines long before it was trendy. • A digital twin often benefits from the ontological data and domain information discussed above. • One of the foundations of the Altiplano digital twin is knowledge about network topology, service topology, customer topology, and related.

Foundational Capability	Commentary on Altiplano
Natural Language UIs	- Altiplano has built a front end to its system, providing natural language access to Altiplano functionality and data. - This allows users to quickly achieve expert user status; and dramatically increases the productivity of network operations staff.
Agentic-Fluent APIs	- All Altiplano data and operations are accessible via APIs and in fact Nokia do their own development utilizing these APIs. - Altiplano supports communication with external agents using MCP
Agentic AI framework (supports operations task development)	- Altiplano offers an Agentic AI framework allowing agents and applications to be efficiently built exploiting the underlying infrastructure -- data, twins, and other capabilities. - Framework is enabled, yet also kept safe, via MCP APIs - Allows 2nd and 3rd parties to develop apps/use cases - One concept is that primitive tasks (data sources, guardrails, re-usable specialized agents) can be built for extensive re-use. - According to Altiplano, one of the most important capabilities are the guiderails and guardrails developed, often specifically for use case types to ensure as desired results.
App Sharing Environments (“App stores”)	- For nearly a decade Appledore has been advocating greater cooperation and sharing amongst CSPs and suppliers-- so that basic building blocks are not reinvented time after time and best practices are spread and adopted most quickly. - Altiplano is relatively unique in the fact that it has implemented an application and agent sharing environment—an “App Store” - They envision multiple types of developers innovating and making the results of that innovation available - This App Store, along with the rest of Altiplano’s open (“interoperable”) environment democratizes innovation – which has proven to be the killer app of many platforms (think: iPhone).

Below is a simplified view of Altiplano’s Agentic AI environment. Some key takeaways are:

1. The Agentic framework treats Altiplano, its data, and internal apps as resources.
2. All resources and framework elements are exposed.
3. LLMs are external and largely left as a choice.
4. A significant portion of the framework is dedicated to security, guardrails, supervision and tasks that ensure safe, secure, operation.


Figure 2: Altiplano Agentic Infrastructure (Illustrative)



Courtesy: Nokia

The next diagram, also courtesy Nokia, provides a graphic summary of AI and ML capabilities provided in Altiplano. It is a combination of use case foci and foundational platform capabilities. Note the focus on highly labor-intensive operations where improvements in productivity and time-to-results can be significant. Examples are field installation; truck rolls, and (from the diagram above) “Ask Altiplano.” Willie Sutton robbed banks because “that’s where the money is”; Altiplano applies AI to labor-intense tasks because.... “that’s where the money is.”

Figure 3: AI/ML in Altiplano



Courtesy: Nokia

Summary & Recommendations

Nearly every supplier of autonomous networking software and agentic AI for network operations is on a long-term march, working to crack the code. None are yet at their destination. That said, Appledore sees mature thinking emerging across the spectrum. Some of the most mature thinking is in technology domains, led by domain experts. This ought not be a surprise given the fact that domains present a more bounded environment, where existing expertise and knowledge may be applied, and the scope of models and algorithms defined.

Significant work is also going on a layer above in the realm of cross domain orchestration and service orchestration. Here the problem is different: a higher level of abstraction simplifies; while a broad an open-ended problem space makes problem definition far more complex. This is where more traditional approaches to LLMs and larger data sets may thrive; along with the recognition that abstraction means that many details may simply be handed back to individual domains once the bigger picture is sorted out.

At the end of the day autonomy, like a physical construction project, begins with a solid foundation and builds upward. Part of that foundation – within an Autonomous Networking software controller -- is formed from domain expertise, data, ontology, and topological relationships. Stepping back and looking at the bigger picture, individual domains and even micro domains form the foundation across which one or more layers of multi domain orchestration may also operate. Without autonomous domains; end-to-end autonomy is at best difficult. Similarly, without a proper intent-based foundation for effective agentic AI, in-domain autonomy is similarly difficult.

Pivoting to a look at Altiplano, who commissioned to work (but left us with broad editorial control), we see a product that is on a clear and documented path toward both proper autonomous networking

control and in building a solid base for effective, safe, and efficient agentic AI to help power that autonomy. As we presented in the table above, Altiplano either embodies or is on a path to implement each and every one of the best practices that we established in our independent Agentic AI research, as applied to network operations.

We are particularly interested in three attributes of this product that often don't get enough attention as topics in general.

- First is Nokia's focus on **openness**. This primarily comes in the form of a rich set of APIs available to all--and therefore a near elimination of lock in risk.
- Second, and closely related, is Nokia's application & agent **sharing environment – or “App Store”** - intended to encourage broad based innovation and the sharing of best ideas.
- Finally, the Altiplano team raised a topic (they call it “**affordable AI**”) that Appledore have been raising for years: AI may be sexy, but it is not always necessary nor even desirable. Some tasks can be performed with greater certainty and lower cost by other means – including flexible rule sets². For well-defined problems; scripts, rules, rule sets and similar methods can be far more deterministic, faster, cheaper to operate, and just plain better. Such scripts can be under the control of broader logic implemented by agents however. Using deterministic tools in combination with AI models that are properly guard railed helps reduce hallucination, speed performance, lower cost, and get more reliable responses.

Appledore has long advised that scale and specialization will win, and that “DIY” – cannot support the level of expertise, research, and investment that excellence demands. For evidence, consider the scale of AI data centers and hyperscale cloud providers. Simultaneously, we observe that the best suppliers are eliminating the weaknesses and concerns that always made CSP's wary of dependency. Altiplano demonstrates a commitment to open-ness, multi-vendor support, and industry standards – as well as participation in the industry efforts that will mature those underlying standards. Its open APIs, Agentic AI framework, App Store, and multi-vendor abstraction are all examples of this commitment.

Further reading is linked in the bibliography below. If you would like to discuss this, please reach out to Appledore. We always enjoy operational discussions with our telco clients and friends.

Bibliography

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 \(Parts 1, 2 and 3\)](#)

² note we are not talking static, brittle individual rules or logic trees, but flexible sets that perform a variation on linear algebra)

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 LLC 2026

