AI and automation drive the evolution to open networks

MAVENIR FUJITSU VIAVI intel NETSCOUT

Go back to the top

Roadmap

Report

Introduction

Setting the stage: AI, ML, GenAI

Openness, automation and AI: How it all fits together

The journey from automation to autonomous networks

The AI toolbox

Data as the foundation of success

A new approach to network operations

Humans are here to stay

A telco-specific, domain-specific approach

Choosing the best use cases

The business case for AI and ML in automation use cases

The good and the bad of AI, ML and GenAI

Q&A

Does automation require AI?

How close to L5's autonomous networks are we?

Are we ready for AI-native networks?

Is AI-driven automation sustainable?

Do we need a CAIO?

Should operators build their own LLMs?

Who do you trust to develop, train, run and host your models?

Which use cases need AI and ML?

What is a use case?

For the impatient: Takeaways

For more detail: References

Conversations

- **Mavenir:** The future network is open and autonomous, with Brandon Larson
- **Fujitsu:** How AI, ML and GenAI can increase network efficiency today, with Blake Hlavaty
- **VIAVI:** The decomposable AI landscape, with Per Kangru and Takai Eddine Kennouche
- **Intel:** AI everywhere, from the edge to the cloud, with Cristina Rodriguez
- **NETSCOUT:** Rethinking AIOPs – It's all about the data, with Rick Fulwiler
- Will AI change how humans operate in telecoms? With Azita Arvani, formerly at Rakuten Symphony, Paul Patras at Net AI and the University of Edinburgh

Report sponsors

FUJITSU

MAVENIR

intel

NETSCOUT

VIAVI

Watch the videos of the conversations

Go back to the top

Analyst report

Watch the videos of the conversations

Go back to the top

Phew! We are finally out of the initial 5G frenzy and not yet into the 6G one, so we can focus on the broader transformation towards open and autonomous networks—networks that are also virtualized, multi-vendor, digitalized and software-based. It is a long-term process, but also one that may have a deeper impact on how we run networks than a new radio access technology and one that requires more operational changes than technological ones.

Wow! The network evolution towards openness carries an inevitable increase in complexity. This is where AI and automation come in; we need them to manage complexity and benefit from the flexibility and choice that complexity offers. AI is finally coming of age and becoming an essential tool to enable automation and the eventual move to autonomous networks, and, in turn, the transition to open networks. We have all the fundamental ingredients for the transformation, but there is still much work ahead.

Argh! There is undeniably hype about AI and, even more, about GenAI. Much of it is driven by the curiosity for technologies that are new and not well understood but can elicit equally unrealistic expectations and fears. Hype is not entirely bad: it forces us to pay attention to technologies that, for some time, have been left in the background. But we need to inject some more grounded understanding of what AI and automation can do in telecoms, what benefits it brings, and how to protect us from the risks it carries.

Hmm... There is still much that we need to learn about the role of AI—including ML and GenAI as part of AI—and to dispel the too rosy or too gloomy projections that still dominate the debate. A growing body of R&D, trials and deployments are taking us down from hype to reality. In this report, we present the work and perspectives of those who have been working on AI, automation, operations and open networks, and discuss the contribution of AI and automation to open networks based on what we are learning in this journey.

AI, ML and GenAI are overhyped and yet (or because of that) attract most investment in telecoms

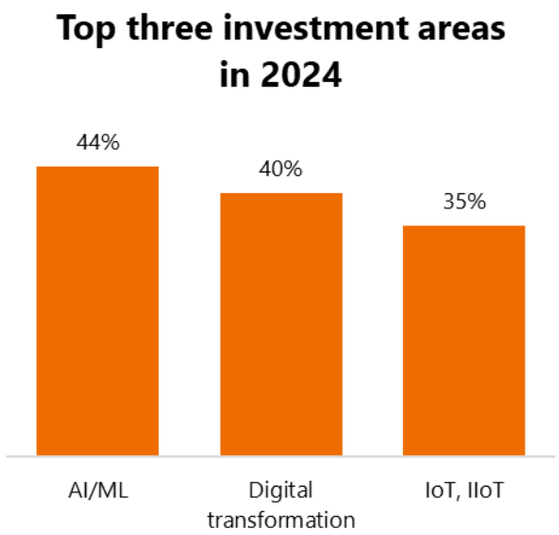

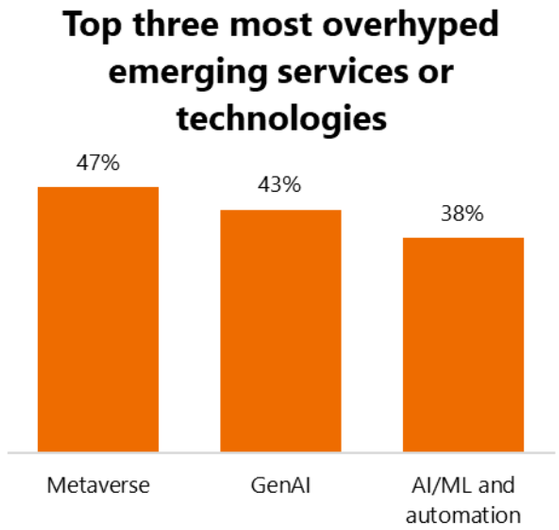

Percentage of respondents
Source: Annual industry survey, Telecoms.com, 2023

Go back to the top

Setting the stage: AI, ML and GenAI

The sudden popularity of ChatGPT and GenAI has led many to equate GenAI with AI—or even worse, with the latest generation of AI—and this is creating much confusion about the different roles that AI, ML and GenAI can play in telecoms.

AI was initially defined by John McCarthy as "the science and engineering of making intelligent machines" in 1956 and has the widest scope. In telecoms, we use AI models that are capable of learning, generalizing, and making predictions and inferences from data through learning. Because they typically need large data sets for the training, training requires time and computational resources, but they can produce recommendations or other outputs that humans would not be able to come up with, or that would take a long time.

ML is a subset of AI with a narrower focus that makes it especially relevant for telecoms because it is less committed to AI's "intelligent" aspects and more to the data crunching. ML algorithms can analyze massive data sets, identify patterns, generate predictions, make recommendations that can be used to automate operations, optimize network performance and support new services. It is not an exaggeration to say that AI and ML will be everywhere: they are tools that can help in any task that is repetitive, needs a scale humans cannot provide, or can be optimized.

AI and ML are often referred to as the "**predictive AI**," as opposed to "**generative AI**" (GenAI). GenAI is also capable of processing impressive amounts of data, but it does so to generate content based on the data fed to the model. Because GenAI creates new content derived from the input data, the output's quality depends on the input's quality—the well-known "garbage-in-garbage-out" model. This explains factually wrong outputs, often referred to as hallucinations, but also biases in the responses. Inaccuracies and biases in the input data are bound to appear in the output unless the data is curated or filtered.

Because of this, the scope for **GenAI** in telecoms is more limited than AI and ML and presents specific challenges and risks. GenAI is mostly suitable for tasks such as customer support services or making available data among staff that require linguistic or visual output. AI and ML are better at recommending actions to reduce power consumption; GenAI is better for chatbots. Large language models (LLMs), the main workhorse of GenAI, use an extremely high number of parameters to process and generate language data, but they are not well suited to work with network KPIs.

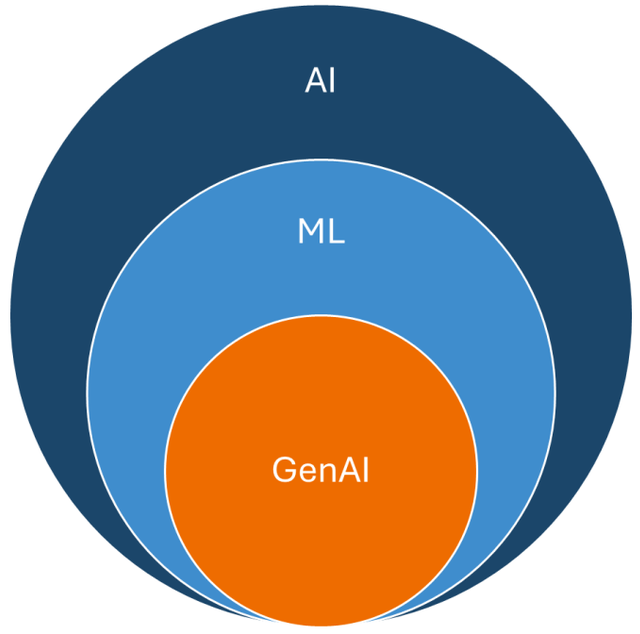

> AI and ML are pushing the operational envelope further, beyond what we can achieve with automation alone. When we add AI and ML, operations people can see things they could not see before. They may get a recommendation on how to tune a parameter in a way they have never had the insight to do before. We can get to a level of performance that was not achievable before.
>
> AI and ML are about efficiency, achieving a higher level of operational capability as network complexity increases. They are not just about removing the human element; they are about figuring things out to make decisions in a more complex environment with a heterogeneous service demand.
>
> **Brandon Larson, SVP, GM, Cloud, AI & IMS at Mavenir**

Go back to the top

Openness, automation and AI: How it all fits together

Openness, automation and AI each bring an independent and crucial contribution to network evolution, but it is their **combination** that provides the essential enablement and adoption drivers.

The transition to **open networks** is more than about open RAN and open interfaces. It extends to all network elements and functions, from the chipsets to system integration, from non-proprietary hardware to open-source software, from virtualization to disaggregation, and from edge to centralized architecture. Openness gives operators unprecedented flexibility to optimize, innovate, differentiate and personalize their services. But this comes with a hefty price: complexity.

Automation allows operators to tame complexity and benefit from the flexibility and openness that technology, standards and a more diverse ecosystem enable. Automation is no longer a nice-to-have, optional capability; it is a necessity to run networks with the efficiency and reliability operators need for critical infrastructure. Hiring more people to run networks is not only prohibitively expensive; it is increasingly an insufficient and inadequate solution. Human computational capabilities are tightly bound—our short-term memory is limited to less than ten items—and make us inefficient at continuously tuning the 1,000+ parameters that may be available to optimize a base station in real-time.

AI and ML are an excellent complement to humans. They learn from our expertise through training and historical network data and use their processing capabilities to apply their learning to automate and optimize networks. As our AI/ML models become more reliable and we learn to trust them, they can be more broadly used for more tasks and with decreasing levels of human supervision in the training, testing and verification phases.

Openness, automation and AI create a **mutually reinforcing loop**. As network openness and complexity grow, so do the automation needs. As automation capabilities expand, so does the ability of operators to move forward in their transition to open networks. At the same time, higher automation levels require a more extensive use of AI/ML. The increased availability and understanding of AI and ML establishes the foundation for the transition to autonomous and zero-touch networks.

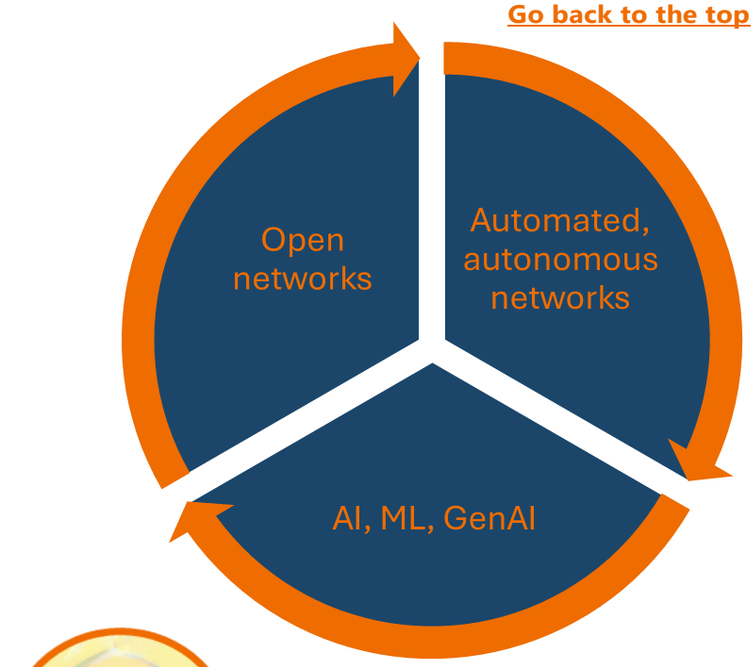

Q&A

Does automation require AI?

There are different levels of automation. Initial automation scripts only needed Excel or simple software tools. Most automation algorithms still do not use AI and may not even need it. If we know how to automate a task, we can have a deterministic if-then algorithm that is simpler and more efficient than an AI-based one. AI adds value and eventually becomes necessary with more sophisticated automation tools– for instance, when there is scope for optimization, but we do not know what the best way to go about it ahead of time, or if the task is so complex that it is difficult to create an algorithm that explicitly covers all the possible scenarios. When we need a model to learn from our input data to generate an outcome or to figure out relevant relationships among variables, AI provides multiple learning models that require good input data and training to be effective.

As network complexity grows, simple, deterministic algorithms are no longer sufficient to optimize automation and to take us to autonomous, zero-touch networks.

Go back to the top

The journey from automation to autonomous networks

In the short term, the goal is to automate as many tasks as possible across all network domains, and to do so by increasingly adding AI and ML capabilities to make automation more efficient and valuable. The long-term goal, however, is to move to **fully autonomous networks**, where all functions and services are automated, and humans are no longer needed to directly operate them.

The TM Forum has developed a 6-stage framework to describe the evolution from L0 with networks operated by direct human intervention to **L5**'s full autonomous networks. Some question whether we will ever be able to get to L5 or whether it is desirable to have fully autonomous networks, even if we have the technological maturity. Similar worries were initially raised about virtualization, cloudification and open RAN, too, but as the technology and ecosystem mature, acceptance grows (even though adoption may be slow and initially limited). Once we get closer to L5, we will be better positioned to see if there may be parts of the network where automation does not bring value. Overautomation could become an issue, but not one we should yet worry about.

The journey to L5 progresses through increasing levels of automation in a very specific way. The main driver of change is not the depth of automation for a single use case or domain but how automation progressively gets **integrated** across use cases, domains and services. Initially, operators pick use cases and deploy automation within their narrow scope. Then, as they adopt automation across multiple use cases and domains, they will have to integrate their work to gain all the benefits of automation, to avoid duplication and conflicts, and, eventually, to get to a holistic approach to automation.

In L5, automation becomes something fundamentally different from what we have today, and the change will require much work, most of which is still uncharted, as we are still learning. Some noted that the steps become increasingly steep as we go from L0 to L5, and they may take longer to move to the final stage. While there is considerable variability in the pace among operators, it will be a **long process**, with the costs more front-loaded than the value. In assessing the business case for automation and autonomous networks, adopting a long-term strategic and financial view is imperative.

Autonomous Levels	L0: Manual operation & maintenance	L1: Assisted operation & maintenance	L2: Partial Autonomous Networks	L3: Conditional Autonomous Networks	L4: High Autonomous Networks	L5: Full Autonomous Networks
AN services (Zero X)	N/A	Individual AN case	Individual AN case	Select AN cases	Select AN services	Any AN services
Execution	P	P/S	S	S	S	S
Awareness	P	P	P/S	S	S	S
Analysis/ Decision	P	P	P	P/S	S	S
Intent/ Experience	P	P	P	P	P/S	S

Personnel (manual) Systems (autonomous)

Six levels of autonomous networks ***Source: TM Forum***

How close to L5's autonomous networks are we?

The TM Forum has developed a powerful framework to assess progress towards L5 that will help the industry keep track of the evolution. However, it is clear that we are far from L5; nobody claims to be close to it.

There are major differences among operators. Some are still at L1, and a few announced plans for L4. China Mobile has been very aggressive in automating its network and expects to reach L4 in 2025. Most large operators are in the L2 stage, still working on selected use cases and often on trials rather than full-fledged commercial deployments.

The progress among vendors is also uneven. Some have a stronger commitment and capabilities than others, with automation more deeply integrated within their solutions.

How long will it take to get to L5? It will take years to get even the most ambitious operators there—and a decade or more to include most major operators. As with all major transitions, it pays to set realistic expectations and avoid rushing change.

Go back to the top

The AI toolbox

AI is **no magic wand**. It does not come as an off-the-shelf solution that operators can add to the network and suddenly create increased efficiency and revenues. It is a toolbox that requires investment, work and operational and organizational change to deliver on its promise—and to avoid the risks it poses.

The AI toolbox can be used to empower advanced automation, but also for other tasks. Some models can be used for multiple automation and other use cases, and leverage the same data. In other cases, a model can be tweaked and reused in a different context, possibly with different data inputs. You can think of it as a **composable AI** approach, that makes you build your AI arsenal gradually as you include new use cases. It requires a holistic approach to the integration of AI in the network and in the organization and an appreciation of when models and data can be reused for a new use case, and when they are not the best choice.

We are not there yet. Instead, as operators go through trials and limited deployments, they often follow a **fragmented approach** with multiple models coexisting. Initially, this is a natural path to follow as operators explore different solutions, but in the long term, this can lead to duplication and conflicts and, in turn, to higher costs, increased operational complexity and reduced benefits.

AI models **vary** greatly across multiple dimensions, in addition to the specific architecture and algorithms they use. There are AI models with a narrow focus that may be not only telco-specific but also network-specific or use-case-specific, and, at the other end, broad foundation models, such as GenAI LLMs. In the middle, there is a growing trend toward customizing large foundation models (e.g., for small language models, or SLMs) to fit a circumscribed scope. The customized model retains the capabilities of the foundational model, but the smaller size makes it faster, more efficient and, if the right input data is used, more effective.

To learn, AI models need **training, fine-tuning and testing**. The use case, model type, trust in the model and reliability requirements determine the best way to do this. Initially, as we learn how to develop, use and trust AI tools, humans still need to be in the loop, taking an active role in ensuring the model performs as intended and does not cause harm. With time, models will require less human supervision and move towards closed-loop learning, in parallel to the transition to autonomous, zero-touch networks.

AI's powerful learning capabilities can also be used in **digital twins**, where virtual

> The process starts with data collection and figuring out what AI priorities you have for your company. That needs to be based on a holistic view from the top to direct the data collection.
>
> Humans play a large role in collecting, labeling, curating and possibly augmenting the data.
>
> The model selection is super important. Picking the right model and doing a lot of trials and experimentation with that is crucial to getting the cost versus performance tradeoffs right. We all want great performance, but it comes at a cost.
>
> Training, fine-tuning and alignment give us the right optimization for the objective we want in the budget that we have.
>
> **Azita Arvani, formerly CEO at Rakuten Symphony NA**

Q&A

Are we ready for AI-native networks?

A simplified way to think about AI-native networks is AI running itself, and that may sound scary or dystopian. Certainly, we are still far off from this scenario.

Moving towards an AI-native approach is compelling and should become a priority today if we take it as the path to integrating AI throughout the network, end-to-end, top-to-bottom.

Deeply integrated AI models require high levels of trust. In turn, we must have the right guardrails, governance, regulations and social policies in place to build a solid foundation for the trust we need. Some will be telecom-specific, but others will have a broader reach.

We also need more transparency and visibility into AI models. Because they can learn new things from their input data that we don't know or expect, it is difficult for the human out-of-the-box observer to understand how models generate their outputs.

We need a native-AI approach for the success of AI and, in the long term, the transition to autonomous networks, but it will require hard work along multiple dimensions and take time.

Go back to the top

representations of existing networks can be used to explore scenarios that simulate the adoption of various AI models. With digital twinning, operators can assess the impact of AI, fine-tune and compare different AI models, and decide whether to deploy them in commercial networks. Especially in the initial stages of AI adoption, digital twins can contribute to building trust in AI-driven automation.

Data as the foundation of success

The AI models used to automate and optimize networks are expensive to develop, train, fine-tune, test and run. Operators can choose among different models and need to compare how they perform and the value they bring. But possibly the factor that plays the largest role in the successful deployment of AI models comes before them: **feeding the right data**—in terms of relevance, quality, granularity and volume.

AI models can crunch any amount of data, but that does not come for free. ChatGPT models illustrate this well: Qualcomm estimates that a ChatGPT query costs 10 times as much as a traditional search. High **per-request costs** reflect a greater processing requirement, and have an impact both on the ROI and on scalability. The resources large models need may not be available where they are needed, for instance, at the edge.

The larger the data set, the longer it takes to train the model and to generate recommendations or other inputs. As we move to real-time network management, **timing** becomes a crucial metric. If the AI model slows down network optimization, the business case for it may vanish.

Throwing data at an AI model and letting it sort out a solution is not an effective strategy for managing and automating networks. ChatGPT and other GenAI models use an approach of this type to capture vast amounts of data without vetting it; it is successful because that's exactly what they are designed to do: create new content based on what's available.

In telecoms, the goal is not to reproduce network behavior but to understand it and improve it. We want AI models to tell us something we do not know and give us **valuable and actionable recommendations** on scale, reliably and quickly. This requires not only containing the amount of data but also reducing the amount of noise in the data that comes from inaccurate or irrelevant data.

Operators have **massive amounts of data** they can feed to AI models. Doing so may

Is AI-driven automation sustainable?

Not necessarily. It can require more resources to develop and run than it saves if you do not choose the right model or the right data.

For instance, using an LLM and feeding it all the network data available to reduce network energy consumption is not likely to be the best approach. But a telco-specific model, fine-tuned to a specific network and operator requirements, will require fewer resources, and be faster and more effective.

Resource needs and costs of running AI models should be factored in the assessment of available options. However, generic claims, often based on GenAI, that all AI is too expensive or wasteful miss the point that there are many ways to deploy AI for automation and optimization. Predictive models, which are more appropriate in this context, are smaller and have a better ROI. Feeding smaller data sets to models also reduces model size and resources needed.

The question is not whether AI models are sustainable but how to make them sustainable and effective.

To meet the customer's need for automation, scalability, intelligence and optimization, we need to operate in an AI/ML ecosystem with third-party providers with pre-trained capabilities, using data sets from different sources, and applying different ML modeling and training techniques.

They will all need to interoperate and fit together nicely to provide value to the customer. Interoperability will have an impact here. If the business value depends on four orchestrated capabilities, how do you vet and test them—not just when deployed, but continuously to ensure operations run smoothly? You need to think about the output quality of a single model on one hand, and the quality of the output of a pipeline of different models in delivering the final business outcome in an operational environment on the other hand.

Takai Eddine Kennouche, AI/Data Modeling Architect, CTO Office at VIAVI Solutions

Go back to the top

reduce their upfront efforts, but also the effectiveness of their models. Selecting the right data, curating it, labeling and contextualizing it are emerging as one of the top lessons learned through the early AI-driven automation trials, and one of the key success factors.

It is also becoming increasingly clear that this is a complex task that requires human expertise and judgment. We don't only have to eliminate inaccurate data; we also have to find out what data and what data granularity are **relevant** for a specific use case. For instance, two use cases may need similar data sets but at different temporal resolutions. A fine temporal resolution may make it less straightforward for a model to identify a pattern that can be identified over long time intervals. A coarse temporal resolution may make detecting a more time-sensitive pattern impossible.

In some use cases, we want AI models to predict responses to scenarios that may happen but that we have not seen before and for which we have no historical data. This is where **adversarial learning or synthetic data creation** may help promote AI models to learn how to deal with novelty by expanding the learning ground. Security is an area where being able to respond appropriately to an anomaly that was not observed before can be highly valuable.

> AI and ML are complex technologies. AI will replace the way we do things today. In networking and in the RAN, there are algorithms that have been used for a long time and operational practices that have been around forever. AI will replace them, but it will take time. You have to test what you have created, you have to get the data, you have to create and train the models, and, finally, apply and execute these models.
>
> You need to develop trust in these models. This will require testing and fine-tuning before you switch them on.
>
> **Cristina Rodriguez, VP Network & Edge Group and GM Wireless Access Network Division at Intel**

A new approach to network operations

AI and ML will play an increasing role throughout our networks—and our society and personal lives—but within network automation, the biggest impact is arguably on **operations**. Many boring and repetitive tasks that are now performed by humans will be executed faster, more accurately and more reliably by software. But this is not the most important aspect of the change in operations. As we introduce AI and ML models, learn how to efficiently integrate them and trust them in our networks, and expand their capabilities, we will be able to optimize and automate network performance in a more effective way that is more fine-tuned to the specific operator's needs, more customer and service based, and with higher temporal and location granularity. This will require both an expanded role of end-to-end orchestration and management, and capillary interventions in real-time and throughout all network elements. In turn, this means more interactions among teams and a breakdown of organizational silos.

This is where the new flavors of operations—**AIOps, MLOps, MLLOps, DataOps, RAGOps,** and more—come in. They address the need to change operations to adapt to the new approach that AI/ML and automation introduce in our networks. AI/ML models

> We call data the hidden gold within the network. You need to mine the gold to see it. We are enamored with the cool AI parts of AIOps, but we must avoid losing sight of the critical data that drives the AIOps engines.
>
> In 5G standalone networks, the hidden gold needs not only to be mined but also refined to become consumable by the AIOps pipeline. That is critical.
>
> The data refinement process helps us avoid this by providing context for the data. In turn, it requires a level of network quality to collect data about who, what, when and where is connected, what the network conditions are, and how they affect the quality of service.
>
> **Rick Fulwiler, Chief Solutions Architect at NETSCOUT**

Go back to the top

alone will not deliver the benefits they are capable of unless they are supported by an accompanying change in how we run the networks.

How should these new styles of ops **integrate** with the current ops? Is there a risk that they will create new silos, as we are trying to get rid of the existing ones? While we need to call out the need for specific types of ops capabilities, especially in the initial stage of AI/ML and automation deployments, we also need to ensure that they build on top of existing operations, enhancing rather than replacing them. We may think of them as ingredients in preparing a dish, and each may be required for one dish but not another. Operators and vendors may require hiring new people to introduce AIOps, for instance, but the new hires should become part of the existing ops rather than joining stand-alone units that cut across functions. At the same time, those in ops who work with AI/ML models have to embed AIOps as part of the scope of their jobs.

Humans are here to stay

The need to integrate AIOps and other emerging ops types into the existing ops organization brings up the wider issue of the **role of humans** as we move to higher levels of automation and eventually to autonomous networks. From a process point of view, if we see AI and ML as tools, we need to deeply integrate them into what we are already doing to succeed. From an organizational point of view, we need to carefully balance the role of humans and software to maximize the benefits of human skills and expertise, and the data-crushing computational power of machines.

The challenge is not to find the quickest way to get rid of humans but to find out how to best leverage their **unique and necessary contribution**. Humans are fallible and have limited processing powers, but they are essential throughout the journey to autonomous networks in multiple ways, which include:

- **Define** the targets of AI/ML models: what they should do, what output to expect, and what level of detail.
- **Decide** what model, model size, algorithm, and training is most appropriate, and revise this decision if results are not satisfactory.
- **Select** input data sources and level of detail; assess data accuracy; cleanse, label and contextualize data.
- **Train**, fine-tune, and test the models.

> AIOps is relevant for many reasons. In telecoms, we are going through a rough now. Revenue, growth and profitability are in decline. Operations costs have started to rise substantially, especially with the rollout of 5G standalone.
>
> How can CSPs operate networks more efficiently, drive down costs and look for more efficient revenue pathways?
>
> They are looking at automation technologies to cut costs and boost productivity. AIOps provide a new competitive edge that the industry needs to tackle some of the automation challenges and can drive untapped data monetization revenue streams.
>
> **Rick Fulwiler, Chief Solutions Architect at NETSCOUT**

> Our culture helps us get the value out of the cloud and AI. We bring cloud and data science people together and prevent them from talking past each other. This is always difficult, but we got the data scientists and the cloud scientists connected with their thinking process. And the same thing has to happen when you bring the operations guy into the room.
>
> We have telco, cloud and data science expertise – all working on the same wavelength. When we achieve that, magic starts to happen. We bring all these cultures together; this allows us to think about problems in a different way and to start innovating.
>
> **Brandon Larson, SVP, GM, Cloud, AI & IMS at Mavenir**

Go back to the top

- **Deploy** the model, integrate them within the network, and ensure they continue to work as intended.
- Over time, **move** to models that require progressively less supervision in training, continue learning on their own, and provide recommendations that can be executed with less or no human control.

These tasks are crucial to the deployment of AI and ML models, and yet they are often **left out** of the discussion, as most of the attention has focused on how to reduce the role of humans and what is the impact of doing so.

We should pay more attention to the role of humans because it is still central but changing in deep ways that affect the overall **organizational culture** and are related to the overall move to IT-inspired processes, digitalization and softwarization.

Operators and vendors are strengthening their data and AI/ML expertise with **new hires** but must ensure they can fit within the organization. They bring in new expertise, skills and ways to use them that may conflict with the established modus operandi at the organization. Long-term employees have essential knowledge, insight and experience but may feel threatened by the new talent. Top management has to encourage a close collaboration between the old and new, reskilling current employees, and getting the new ones to get telecom-specific knowledge. This is another reason why it is important that AIOps and the other new ops are well integrated with the existing ops, and are not homed in separate teams. Unfortunately, many organizations seem headed towards establishing special-purpose teams, perhaps to showcase their commitment to innovation or to push forward changes that may be difficult to introduce otherwise. We need a deep commitment throughout the organization to introduce the cultural change that successful technical innovation in the long term requires.

In the past, AI and ML and, even more, automation were often dreaded as a threat to **employees** and leading to potential massive layoffs. These fears have now largely subsided. We need automation to deal with increasing network complexity; without it, we would need to expand the workforce to financially unsustainable levels. But perhaps more importantly, automation can help operators and vendors deal with an increasingly aging workforce and the difficulties they face in attracting new talent more inclined to accept high-tech jobs outside of telecom.

Q&A

Do we need a CAIO?

To show their commitment to AI, it has become fashionable to have a Chief AI Officer, or CAIO, who takes responsibility for AI initiatives within the company.

While this may accelerate the adoption of AI and help establish a vision and commitment from the top, it is a risky proposition because it implicitly positions AI as a separate technology that intrudes into the existing technological, operational and cultural environment. As a tool, AI should be integrated into the organization as much as possible within the existing structure.

Organizations will have to hire AI experts, but the priority should be to get them to work as closely as possible with the staff and share—and eventually blend—the different skills and expertise they have rather than work in silos that may end up clashing and curtail the benefits that AI and automation can bring. Didn't we want to move beyond siloed organizations? If so, we do not need CAIOs.

"The telecom workforce is aging out, and we are unlikely to have the large staff we used to have, even though networks are more complex on the wireless and transport side. Traffic, the number of use cases, and costs are also growing.

Operators need automation to address all these changes. They will need more proactive staff to manage their networks, which, with automation, will spend less time reacting to issues.

Software updates to the network nodes are one use case we are looking at right now. This is a use case where automation saves time. Ultimately, automation is about reducing costs and incidents and improving reliability because we operate the network in a more repeatable way."

Blake Hlavaty, Director of Global Network Software Offers at Fujitsu

Go back to the top

A telco-specific, domain-specific approach

ChatGPT and other GenAI models that use massively large, unstructured data sets, are computationally demanding, slow and expensive, and notoriously prone to inaccuracy or hallucinations have attracted huge interest and have a bright future ahead. They will improve the quality and reliability of the results, become faster, and need fewer resources, increasing their value and use.

However, these general models that can create content in response to any query are not very helpful in telecoms, where we want models that only **know about networks and their users**, and generate results only within this context. The advantage is that such models are smaller, more nimble and more accurate. The (slight) disadvantage is that to make them smaller, we need to select data that is relevant. As discussed before, this requires effort and access to reliable data.

To get there, we need to move from general LLMs to telco-specific and domain-specific models, which can be further customized to become **network-specific models**. If an operator wants to develop an AI/ML model to improve energy efficiency, it may use a model that other operators use too, but using input data (e.g., historical network performance data) that is specific to its network. It would not have access to other operators' data, and even if it did, it would not be helpful. Models like ChatGPT are much more promiscuous and gladly accept data from all sources; they would not be of much help in reducing power consumption in a network.

While there is a general agreement over the need to move to a telco-specific approach, there is less agreement on what that entails with **GenAI**. (For predictive AI, models are almost inevitably telco-specific because they are developed specifically for telecom use cases.) Operators could use available public LLMs and feed them only network-specific or use-case-specific data, modify them to improve the fit, use open-source resources to develop their own models or build their own.

Building an LLM gives operators control over GenAI-driven services like customer support. However, operators may find it challenging to secure the deep expertise and huge financial resources to develop an LLM. More importantly, however, it is unclear whether building yet another LLM is necessary or even valuable—customizing the LLMs already available may lead to better results, a shorter TTM, and a lower cost.

If operators decide to **customize existing public models**, they face the choice of either doing so themselves or relying on vendors or system integrators. They must also decide

Should operators build their own LLMs?

SK, Deutsche Telekom, e&, Singtel and SoftBank Corp have recently created the Global Telco AI Alliance joint venture to develop their own multilingual telco LLM that operators worldwide can use to improve their customer support services through personalized AI-assisted apps and chatbots.

It is unclear how much of a departure from what is already happening. Operators have been using chatbots for customer support for some time already and are increasingly adding GenAI to improve the quality of customer interaction. It is a use case with a clear business case potential due to churn reduction, customer-support cost savings and increased efficiency, and where GenAI can provide much-needed performance and ROI improvement.

The project, however, seems to be more ambitious than this and eventually include a personal AI agent, as a ChatGPT-like interface that subscribers can use to submit queries on any topic. This would increase the scope (and cost) of the LLM, but it is not clear how it would improve and not simply duplicate the capabilities already available in public LLMs.

It is not a question of when this will happen. AI must succeed for operators to run their networks cost-efficiently and remain relevant. Hyperscalers have already started to move into the telco space. Incumbents may be threatened by more agile companies that are quicker to innovate. It is all about adoption, being open-minded, and being willing to try new paradigms to remain relevant.

Paul Patras, CEO at Net AI and Associate Professor at the University of Edinburgh

Go back to the top

which LLM or LLMs to use and whether to use proprietary or open-source ones. The extent of the customization is another crucial dimension, which will depend on the use cases.

A **small language model (SML)** may be more appropriate for many use cases, and vendors are working on such models. Again, if we think of AI as a toolbox, we should expect to see a GenAI component become an integral part of many new solutions or added to new versions of existing ones rather than being a separate offering. In this perspective, GenAI provides a valuable conversational interface that improves accessibility, ease of use and performance of the underlying solution, and that may be especially useful to newly hired, less experienced staff.

Finally, because of the demanding processing resources required, some operators may elect to host LLMs in the **public cloud**. Still, others may prefer to keep them within the network, especially if they have a tightly focused model. We don't know yet which options will become dominant, and multiple approaches will likely coexist, especially at the beginning.

For sure, **hyperscalers** are keen to expand their presence, and they now offer not only cloud services but also customization of their LLMs directly or through partnerships. Within the emerging AI ecosystem, hyperscalers are emerging as key players, working hard to showcase their expanding telecom experience and partnerships.

Choosing the best use cases

To get to L5 autonomous networks, we must start with **individual use cases** and then widen the scope to an entire domain and the end-to-end network. Every task, function or service can be a use case; there are innumerable lists of use cases, and they are invariably long. Because they include virtually every aspect of the network, it is pointless to reproduce such a list. However, the timeline for adoption will vary across use cases, as the graph from the TM Forum on GenAI use cases shows.

So, how to **choose** where to start with? What are the low-hanging fruits—the use cases that are straightforward to deploy, have a good ROI, and fit well within the operator strategy? The choice is going to be different depending on whether it is a GenAI or AI/ML use case.

GenAI's early use cases mostly involve customer-facing services or employee

Who do you trust to develop, train, run and host your models?

Some operators will build their own AI/ML models and automation platforms, and have the in-house and financial capabilities to do so. Most will prefer (or have to) work with vendors, system integrators or other parties to deploy AI and ML models in their networks.

Initially, operators may need more third-party assistance as they build up their internal capabilities and integrate AI and ML into their operations, but in the long term, outsourcing AI/ML-driven functions that have become an essential part of most of their operations will likely be inefficient and cumbersome.

At the same time, as AI and ML become fully integrated in commercial solutions, the need for AI and ML models to be developed separately and introduced as needed to automate or optimize specific functions will decrease.

So, this becomes a moot question: you will choose those you trust to provide a solution to your underlying requirements and have been able to integrate AI and ML within the solution. This could be a third party you are working with already or a new one—and many new vendors with a strong focus on AI and ML have emerged. But those most likely to succeed are those who combine AI and ML with deep telco-specific expertise, and have solutions tailored to the telecom market and, typically, to specific use cases.

When they try to expand and productize these trials, operators start looking at the costs of maintaining well-trained models with good hygiene and expandability. At this point, they realize that AI models have a structural cost.

As the industry is moving beyond the hype peak, we need rationalization. People have stopped building AI use cases just because they are AI use cases; they now build them because they are the right things to do. And at this point, we need to sober up and face the real costs. It is time to have this discussion today.

Per Kangru, Business Development & Strategy CSO at VIAVI Solutions

Go back to the top

productivity or efficiency. The first group includes chatbots for customer support, billing services, marketing campaigns, and B2B service interfaces. They aim to reduce churn and customer service costs, and to improve the user experience. The second group facilitates access to internal documentation and network data to support employee tasks, assists new, less experienced employees, and generates recommendations.

AI and ML support a broader set of use cases to support automation of any aspect of operations. Root cause analysis (RCA), anomaly detection, traffic prediction, predictive network maintenance, network slicing, detection of security threats, network optimization, resource scheduling, load balancing, capacity planning and scenario analysis are among the use cases most commonly cited.

But **energy efficiency** is widely considered the most compelling use case in the short term because AI and ML can correlate historical and real-time data more effectively than rule-based models, and more accurately predict traffic and optimize power consumption. It also ranks high among the strategic priorities of operators, and its value—improved energy efficiency—is relatively easy to demonstrate.

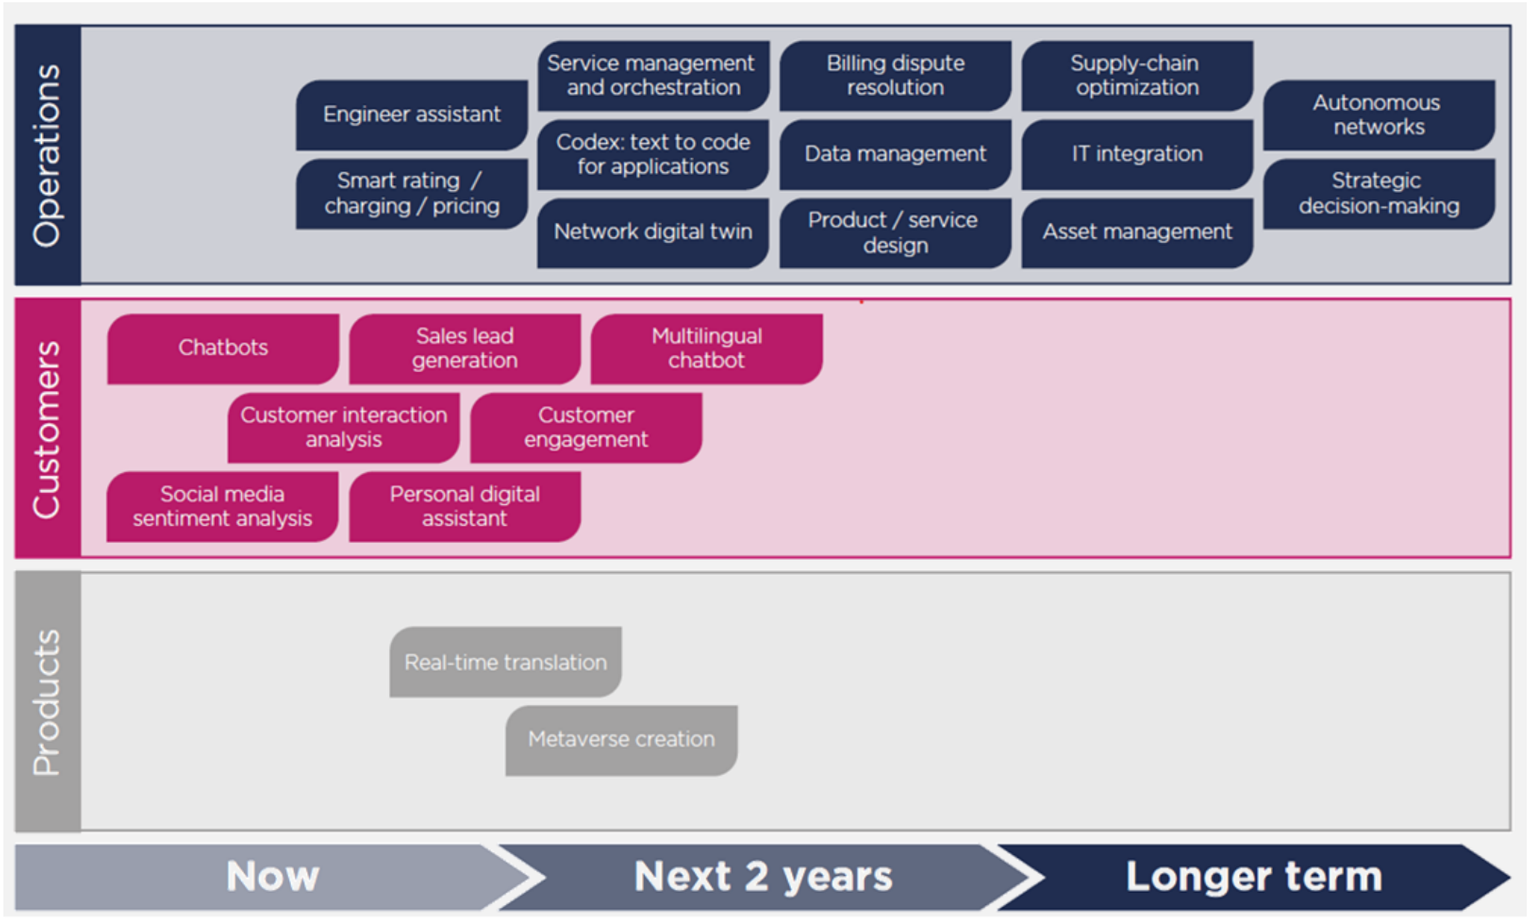

GenAI use cases deployment roadmap *Source: TM Forum*

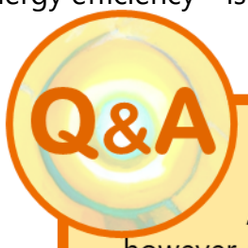

Which use cases need AI and ML?

As we move to autonomous networks, AI and ML will be pervasive, and finding use cases that do not require them will be difficult. Today, however, most use cases do not require AI and ML, even though they could benefit from it. Operators can improve energy efficiency without AI and ML, although they can do so more efficiently with their help.

The question then becomes at which point AI and ML bring value or become necessary to a specific use case.

AI and ML add cost, effort and complexity to a use case, so they should not be used unless they add value and have a positive ROI (but see more on how to determine that later).

More specifically, tasks that can be adequately addressed by more traditional, deterministic, rule-based algorithms because we understand them and know how to reliably complete them do not require AI and ML and would not benefit from them.

Of course, an AI/ML model could prove that there are better ways to approach the task and find new correlations and recommendations, but the improvement is likely to be smaller than for a more complex task that rule-based algorithms fail to address.

On the other end, some tasks may be too complex or too critical for today's telco AI capabilities, and they will require more time and experience before they can be cost-effectively deployed.

Go back to the top

Q&A

What is a use case?

have so far talked about use cases as if we know what they are. But there is confusing variability in what the term refers to. It could be a broad part of the business or the network, or a specific function.

We mentioned energy efficiency as the mostly talked about use case. It is a broad one, that may include all network domains from the RAN to the core, but it could also be confined to the RAN, responsible for most energy consumption. AI and ML models cannot improve energy efficiency directly by optimizing some metrics. They need to work on multiple fronts and optimize across different network elements and functions. AI and ML are helpful because they can correlate data from multiple sources and identify the appropriate tradeoffs. For instance, they can recommend sleeping modes in real time using traffic prediction, geolocation, spectrum and resource availability, and user experience requirements. They can use RCA and anomaly detection to adjust sleeping modes when there are network issues or unexpected traffic patterns. They can optimize power efficiency at a single cell level or across an area with multiple cell types.

In turn, all these components of the energy efficiency use case are, in turn, use cases on their own—some with a narrower scope, some with a wider scope. And they are all components of other use cases. For instance, traffic prediction can be a use case on its own but also contribute to the network slicing use case.

This is not a problem at all as long as we are aware of the implications. An effective energy efficiency model will require inputs from other models that, in turn, may also require AI and ML. The business case—and the attending costs and benefits—will depend on the combination of all these contributing models.

The business case for AI and ML in automation use cases

Ultimately, any use case—and more generally AI and ML in automation—have to provide **value** in terms of cost savings, revenue generation or both. We expect that increased efficiency will bring a financial benefit, but quantifying it accurately is not trivial—and operators are struggling to find realistic answers amid the inflated expectations that hype generates.

If we take the **energy efficiency use case** as an example again, a reduction in energy costs is undoubtedly a positive outcome. However, we need to factor in the costs of developing, training and deploying the models, of selecting and vetting the input data, and of integrating and managing the automation solution within the existing network. In the short term, during the development and initial adoption period, these may outweigh the power efficiency gains, especially since they may initially be contained as the models get fine-tuned. This is no different than business cases for other solutions, but in the

> With new technologies, it is important to consider how the customer will adopt them, not just in terms of sales, business or budget. You need to think about the person using the software day to day, or going through adoption.
>
> I have much empathy for my customers trying to do their jobs under the gun and facing increasing complexity. They need the right tools to get their work done and, specifically, to be able to adopt new technologies. We need to understand this as we develop new products. If we don't, we risk slowing down the adoption process or limiting the benefits of new technologies.
>
> **Blake Hlavaty, Director of Global Network Software Offers at Fujitsu**

Go back to the top

case of AI and ML models that require a welcome but nevertheless difficult operational change, the costs may be higher and come from a wider set of sources.

More importantly, however, as noted above, the energy efficiency use case may need AI and ML **models in adjacent areas**—e.g., for traffic prediction and anomaly detection.

Should the cost of adding AI and ML models to improve the efficiency and accuracy of traffic prediction and anomaly detection be added to the energy efficiency use case? Even if they were adopted exclusively to support the energy efficiency use case, they would be used for other use cases, and the cost should be spread across them. Initially, however, when only a few use cases are assessed or deployed, it is difficult to **split the costs** among present and future use cases. Similarly, some of the improved power consumption may come from increased efficiency in other areas—i.e., they would have occurred even if the energy efficiency use case had not been adopted.

As we move closer to L5 autonomous networks, AI and ML models will be adopted across the network, and, if properly integrated, they will reinforce each other and further **improve the ROI**.

Despite the difficulty in isolating costs and benefits from the broader adoption of AI and ML, we still need to assess the **business case for single use cases**. However, we need to keep in mind that they may overestimate the costs (and sometimes the benefits) as they do not distribute them across adjacent use cases that share AI and ML resources. In addition, the short-to-midterm ROI may underestimate the longer-term ROI, which may be more appropriate since adopting an isolated use case today will enable the transition to autonomous networks, which will deliver value in addition to that provided by individual use cases.

At the same time, it is worth trying to predict—or set an expectation of—the **overall impact** that the introduction of AI and ML will have in the end-to-end network in the long term to take a holistic view of the impact of the evolution to autonomous networks. The TM Forum predicts cost savings of 1.9% to 9%, mostly from increased efficiency in network operations in a bearish and bullish scenario. There will also be variability among operators, each following a different approach to automation and seeing different gains.

Similar considerations apply to the **revenue** side of the ROI for use cases that, unlike the energy efficiency, include that. While AI and ML are more likely to bring a more direct contribution to cost savings than revenue generation, they will enable new strategies and approaches to service creation and management, and increased efficiency will make

> With AI, operators can right-size the network based on the conditions they face and improve efficiency without compromising reliability and security.
>
> Operators want to improve the TCO after making a tremendous investment in 5G spectrum and infrastructure and get new revenues with network slicing or other new technologies. AI will play a crucial role in these efforts.
>
> With AI, they can automatically allocate network resources to meet the required SLA. They can do this in the background in a cost-efficient way that makes network slicing affordable.
>
> **Cristina Rodriguez, VP Network & Edge Group and GM Wireless Access Network Division at Intel**

Opex / capex item	Impact of AI	Bullish target % reduction	Impact ($ billions)	Bearish target % reduction	Impact ($ billions)
Network operations		20.00%	50.3	5.00%	12.6
Network infrastructure		0.00%	0.0	0.00%	0.0
Utilities		10.00%	5.4	2.00%	1.1
Content		2.00%	2.4	0.50%	0.6
R&D		0.00%	0.0	0.00%	0.0
Sales & marketing		10.00%	20.7	2.00%	4.1
General & administrative		5.00%	9.4	1.00%	1.9
Cost of Devices		0.00%	0.0	0.00%	0.0
Depreciation and amortization (D&A) expenses		0.00%	0.0	0.00%	0.0
Property, plant & equipment (PP&E)		5.00%	73.8	1.00%	14.8
Software licenses		0.00%	0.0	0.00%	0.0
Total opex/capex saving			**162.0**		**35.03**
Saving as % of total costs			**9.0%**		**1.9%**

No impact · Small impact · Moderate impact · Big impact · Transformational

Potential long-term impact of AI on capex and opex ***Source: TM Forum***

Go back to the top

more network resources available to support new subscribers and services. As on the cost side, it is difficult to assign the contribution of AI and ML to specific revenue streams, but it is useful to estimate the potential revenue increase that comes from end-to-end AI and ML adoption.

The good and the bad of AI, ML and GenAI

Support for AI, ML and GenAI across the industry is ubiquitous, although the commitment and adoption pace varies greatly among vendors and operators. Expectations of the **benefits** run high, but there are many lingering **concerns** that understandably are slowing down adoption but not reversing the trend. The list below summarizes the topics we discussed in the conversations that follow in the report.

Benefits	Concerns
Higher operational efficiency due to optimized use of network resources and staff	Trust and reliability, need for transparency and observability
Improved customer services, including support and marketing	Technological maturity
Better user experience	Lack of human resources, talent
Move to autonomous, zero-touch networks	Cost to develop and train models, operating costs to run models
Increased energy efficiency and network sustainability	Security, regulation, governance, and need to set guardrails
Better management of network complexity, composability, new ops approach	Energy and resource requirements
Value creation, mostly through opex reduction	Coexistence with legacy, duplication avoidance
Support for new services	Ecosystem maturity, risk of vendor lock-in
Address talent shortage, improved work quality, support to staff, new hires training and reskilling	Over-automation

Go back to the top

Takeaways

Automation driven by AI, ML and GenAI are essential in driving the transformation to open networks and to fully L5 autonomous networks. We must go beyond the hype about AI and, especially, GenAI and set realistic expectations.

AI, ML and GenAI help us manage increasing network complexity but require much work from everybody in the ecosystem. It is more than a technological change; the cultural and organizational transformation is just as challenging and demanding as the technological one.

Despite the huge interest in GenAI, its role in telecom will be more limited than AI and ML, which provide a wider set of tools better suited for the structured data sets and desired outcomes of network operations. We need both GenAI and AI/ML, but we must appreciate their different contributions.

AI, ML and GenAI are a powerful toolbox, not a stand-alone solution that adds a new layer. The new tools must be integrated into network elements, functions and services as part of the softwarization and digitalization process. Avoiding duplication and silos creation is also essential.

Humans are still essential, but their role is changing. With AIOps and other new ops, network operations require new skills and a new culture. The goal is not to disintermediate humans but to find a new way to use human expertise, insight and skills in new ways to drive innovation.

Telecom networks have specific requirements that demand telco-specific and domain-specific models. But they all will use existing AI, ML and GenAI algorithms, public models and open-source code as the foundation; we should resist the temptation to build proprietary large models unless needed.

AI and ML will be pervasive as we ramp up automation and move to open and autonomous networks. We need to choose wisely where to start and select use cases that set a solid basis for this journey, for which we have good data, a solid business case, and complex enough to make them challenging.

Go back to the top

References

Altman Solon (2023) The opportunity in generative AI for Telecom. Sponsored by AWS
Ashish Vaswani et al. (2017) Attention is all you need
Athanasios Karapantelakis (2023) Generative AI in mobile networks: a survey
Brandon Larson (2024) Integrating AI Into the Anatomy of Autonomous Networks
Deutsche Telekom (2024) Kick-off: Joint venture for telco-specific LLMs
European Union (2024) Artificial Intelligence Act
Fujitsu (2023) Advanced Management of Radio Access Networks using AI
Fujitsu (2023) Ignite 5G network analytics with machine learning and artificial intelligence
Global Telco AI Alliance (2024) Global Telco AI Alliance founding parties sign agreement to establish a Joint Venture focused on co-developing and launching a multilingual Telco LLM
Inclusively (2023) AI and ML: friend or foe of inclusion?
McKinsey (2023) The economic potential of generative AI: The next productivity frontier
Monica Paolini (2023) The open RAN
Monica Paolini (2024) Can we trust GenAI in telecoms?
Nvidia (2024) State of AI in Telecommunications: 2024 Trends Survey Report
O-RAN ALLIANCE (2024) Research Report on Cross-domain AI
President Biden (2023) Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence
Qualcomm (2023) Towards an AI-native communications system design
Qualcomm (2024) History of AI: How generative AI grew from early research
Senza Fili (2019) AI and machine learning: Why now? Network optimization in the age of 5G
Senza Fili (2022) The Open RAN ecosystem gathers steam. The future of the RAN is open
Senza Fili (2023) Automation at an inflection point: Are you ready to go?
Senza Fili (2023) Energy efficiency in wireless: RAN and beyond
Senza Fili (2024) Network autonomy or network automation?
Telecoms.com (2023) Annual industry survey
TM Forum (2021) Autonomous Networks: Empowering digital transformation
TM Forum (2021) Autonomous Networks: Exploring the evolution from Level 0 to Level 5
TM Forum (2023) Generative AI: Operators take their first steps
TM Forum (2023) Using AI to improve customer experience
TM Forum (2024) Building an AI strategy: telcos put the foundations in place
UKTIN (2024) Artificial Intelligence Future Capability Paper

Go back to the top

Conversations

Read the edited transcripts below, or watch the conversations here

Go back to the top

The future network is open and autonomous

A conversation with Brandon Larson, SVP, GM, Cloud, AI & IMS at Mavenir

It is not coincidental that networks are evolving towards openness and automation at the same time. Openness brings flexibility and choice but also increases complexity. Automation, aided by cloudification and AI, enables operators to manage complexity and benefit from the increased network choice and flexibility. Together, openness and automation are changing how we operate networks and deliver services in a profound and long-lasting way.

Brandon Larson, SVP, GM, Cloud, AI & IMS at Mavenir, talks about how openness and automation work together to innovate wireless networks and how AI, ML and GenAI play central but distinct roles in enabling this transformation.

Monica: Brandon, you cover a lot of ground at Mavenir, including cloud, AI and IMS. Tell us more about what that means in your daily work.

Brandon: I do have two roles at Mavenir. On one hand, I manage our IMS solutions, which cover voice and messaging. On the other hand, I run our cloud automation and AI solutions, and this is what we will be talking about today. I work with our ecosystem partners, like Red Hat, Microsoft, AWS and Google, to leverage state-of-the-art technologies for cloud, AI and automation to develop autonomous networks for our telco customers.

Monica: Let's start from the ground up. How do you define network automation and autonomous networks? How do they differ from AI?

Brandon: It is important to define the differences among these terms because there is a tendency in the industry to conflate them. They are complementary but distinct, and they solve different problems.

An autonomous network can operate with minimal human intervention; it is self-governing. It can think for itself and learn about network events, traffic patterns, and settings. Given a set of objectives or based on current conditions, it knows which actions to take. It needs AI to do this.

Automation, instead, does operational tasks that we could otherwise do manually, such as installing and upgrading software, updating a configuration, scaling out, scaling in, and terminating nodes. These operations are being automated today, with humans setting declarative objectives.

In an automation framework today, a human operator feeds the automation platform the desired network state, and the automation framework takes the network to that desired state. The operator doesn't have to tell the automation framework step by step what it wants it to do. Humans use declarative instructions to control the network's end state.

AI is different. It learns how the network works by ingesting network data and using this knowledge for problem-solving and decision-making. It understands the causation of network events and what leads to them. It is like understanding what causes people to have heart attacks: maybe because they smoke or because of family history. In AI, we call these attributes, and we assign a feature importance to them to define how they affect things. Sometimes, it is not a linear correlation.

These causation models give operators deeper insight into what is happening and why. Once they understand network behavior, they can forecast events and prevent problems. With reinforcement learning, recommendations

Watch the conversation

Go back to the top

become a big thing. What do we want the network to do? What actions does the network need to take to achieve that objective?

For example, AI can figure out how to increase spectral efficiency, or improve mobility or coverage by turning the right knobs to optimize the network in a way that a human could not even comprehend.

Monica: So we do not need AI in the initial stages of automation, but its role will progressively expand as we transition to autonomous networks.

Brandon: I think of automation as the arms and legs of the network. There, it all started with the move from OpenStack to Kubernetes, which brought in a more sophisticated orchestration layer. On top of that, we have GitOps. Git is a repository where we can store declarative information and do versioning and validation. Then, we have continuous deployment solutions such as Argo CD and Flux that help us create a declarative operational framework by synching the network to the declarative state. Finally, we bring in an orchestrator or workflow engine to manage the automation steps.

For infrastructure, automation is critical for bios and firmware upgrades, to reduce hardware failure and to effectively control lifecycle management and security. There are many infrastructure automation solutions on the market from partners we work with, such as Red Hat advanced cluster management, ACM. At Mavenir, we integrate infrastructure automation solutions with application automation to provide an end-to-end automation solution.

The last piece of the puzzle is the network functions. As a vendor, Mavenir develops software that is compatible with the automation framework. We must build our software as container images and need tools to create automation artifacts or blueprints – for instance, Helm charts, YAML files, or JSON configurations. We stick these blueprints into the automation framework, which provides the instruction set for automation.

To understand what an AI/ML framework is, we must understand what AI and ML are about. You often hear the term inferencing. It is reaching a conclusion based on your reasoning and evidence. This is what ML models do based on an algorithm – the reasoning aspect – and the network data – the evidence aspect.

Our machine learning development starts with the data. About 85% of the ML model's heavy lifting is on the data side–data collection, data imputation, filling in the blanks, normalization, scaling, and dimensionality reduction. We then feed the data to the ML model to make inferences —they could be a control action or a change in a steering traffic parameter.

The data collection must intersect with the cloud to do it efficiently. We then get into the ML part, where we process that data, make smart decisions, and feed them back to the automation framework for control actions. Whatever the ML model tells us to do, it intersects with automation.

I look at this automation framework as if we are building a robot, with the automation framework being the arms and legs that do everything and AI/ML being the brain that tells the arms and legs what to do.

Monica: We started introducing automation without the need for AI or ML, and there was much more caution in adopting it. What has changed over the last 20 years to make automation compelling and urgent?

Brandon: People always wanted automation since we started the cloud and virtualization journey in 2011. We went to NFV in 2016 with the promise to bring in more lifecycle management and automation.

These were the OpenStack days when we did not have the tools or ecosystem comparable to Kubernetes and had too much development to do. We had an ETSI MANO reference framework, but someone had to develop the components, such as an orchestrator and VNF Manager. Then, we had to run PlugFests to stitch all these things together.

Today, in the Kubernetes environment, there are off-the-shelf tools that work out of the box. Whether you are in telco or not, you do not need to code Kubernetes. Google did it as its third-generation orchestration system. We did not invent Argo or Flux either. They are just there and work.

The ecosystem is more sophisticated, too, even in infrastructure automation. We work with Red Hat and use their powerful ACM tool for infrastructure automation, which is not available in OpenStack. We also work with partners such as AWS, AMD and Intel, which straddle within and outside the telco ecosystem and enable us to bring innovation developed elsewhere to telcos. Our automation capabilities are much more powerful than they were even three years ago.

Go back to the top

Monica: We also gained more trust in what automation can deliver. We are no longer scared about what we used to call the black box. Building trust is a process that goes beyond technology. And introducing AI and ML requires an additional layer of trust.

Brandon: To remove the human element from the operational process, we need to know where we want the network to be, how we want to evolve, or how we want to roll out new software. Automation makes things happen faster and reduces operational expenses. It removes the human error component, too. Humans are fallible and make mistakes. With automation, we can avoid an unintended customer impact. This is the automation value proposition.

AI and ML are pushing the operational envelope further, beyond what we can achieve with automation alone. When we add AI and ML, operations people can see things they could not see before. They may get a recommendation on how to tune a parameter in a way they have never had the insight to do before. We can get to a level of performance that was not achievable before.

Automation and AI/ML are complementary because they work together to create an autonomous operational envelope by solving different problems.

AI and ML are about efficiency, achieving a higher level of operational capability as network complexity increases. They are not just about removing the human element; they are about figuring things out to make decisions in a more complex environment with a heterogeneous service demand.

Monica: Our networks are getting so complex that there are not enough people that you could hire, even if costs were not an issue, to operate the network efficiently.

Brandon: In a demo our data science people did, the head of operations gained some insight from the ML model's analysis of real data, which pointed towards something he was not aware of. He asked his operation guys if they understood what was happening. They did not. So, they started using the ML model to drill down deeper into the observed behavior. They got a deeper insight because the ML model figured out patterns in the network they didn't know.

To your point, this is not just about reducing headcount or hours of effort. It is about getting results that the human brain misses. There are only so many data points that a human can correlate. That is what AI does; it can look at everything that happens in the network, crunch all the data, and come up with insights that surpass what a human brain can do. That is the difference.

With AI and ML, we can reach an entirely different level of dealing with operational complexity. Automation can only take us so far without AI. If automation is like having a car, AI is like having a rocket ship.

Monica: Short-term memory in humans is 5-8 items. A massive MIMO has over 1,000 parameters that have to be set to optimize performance. We can't possibly manage that.

Brandon: This is getting worse as we move to 5G and eventually 6G. We want to deliver new services and use new technologies such as network slicing. Human brains are fallible, and we can crunch only so much data. To be efficient, we need to operate complex networks in real-time and to do so, we need AI and ML. This is what everybody outside of telco uses to solve complex problems and deliver personalization. We need to bring in and use these tools in telco to add value, as they do in other verticals.

Monica: You talked about AI and ML, but these days, there is a rapidly growing focus on GenAI. What is its role in telco? How useful is it? How different is it from AI and ML?

Brandon: Everybody talks about GenAI because they started using it for content creation out in the wild with OpenAI and ChatGPT. But AI and ML have been around for decades. In the 90s, ML was beating chess champions. ML is a subset of AI and data science, and GenAI is a form of ML. GenAI has a role in telco, but it is not the only ML capability that is available.

For instance, regression algorithms understand and learn the relationships among independent and dependent variables, inputs, or attributes. If I want to forecast a house price, I start with the current price, which is the dependent variable. The factors that influence its future price are things like square footage, number of bathrooms, and school system – these are features in an ML regression model. This is a very valuable model but is different from GenAI.

Our data scientists look for the best AI tool to solve an operational problem, and GenAI or LLMs are not always the right choice. LLMs have been trained on tons and tons of data and use a lot of horsepower. Using them in a telco network for some operations is like killing a mosquito with a bazooka. It's overkill. Instead, we pick the right

Go back to the top

algorithm for the right problem, and we do this in a very practical manner to avoid throwing a huge number of resources at everything.

We can turn to GenAI for very specific purposes. For instance, we can use it to clean our data before entering the ML framework by imputing missing values. We also use GenAI to deep fake the network and create digital twin models for offline validation. For this, we use generative adversarial networks, a type of GenAI that can learn deep network patterns and mimic the network by learning the relationship between a large number of attributes. Once the model is able to mimic the network, we can change any of those attributes, and the model will predict the impact of that change. We use this to validate the impact of a change AI may make before it is implemented in the network.

We have a reinforcement learning model that scares the death out of operations folks because it takes autonomous action. But before we let that ML model change the network, it will validate a proposed change using the digital twin to examine the anticipated network behavior. The digital twin model can do an offline validation before changes are deployed in production to lower the risk of running a reinforced autonomous closed-loop operation model in the network.

Monica: The confined role of GenAI in network automation may also diffuse the security and reliability worries specific to GenAI and more worrisome than those of AI and ML.

People worry about hallucinations in GenAI, for instance. How do we address these security and reliability issues?

Brandon: Many people think about Chat GPT, LLMs, and GenAI when they think of AI. These are not the only AI tools available to telcos; we use other methods depending on the use case.

For security, we use extreme gradient boosting, a type of regression decision tree for classification, to determine if a message is fraudulent. We use other algorithms, such as Random Forest, to determine causation of observed behaviors. Unlike a neural network, Random Forest allows you to investigate what the algorithm is doing. This gives you better insight into explaining why the algorithm is coming up with a given answer. For KPI forecasting, we use time series algorithms such as Recurrent Neural Networks.

We also use deep Q-learning for reinforcement learning applications. This is a reward-based carrot-and-stick system in which the agent tries to maximize a reward based on action. With this method, the system learns the optimum policy for determining what action to take in a given state to obtain the given objective.

These ML techniques can solve many telco problems for which LLMs are unsuitable.

LLMs are great for unstructured data but are only as good as the data you feed. If the data is not there or it is outdated, you may end up with hallucinations. LLMs can spit out a result with high confidence and yet get the answer completely wrong, and it is unable to cite the source of the answer.

To address these limitations, we use LLMs with RAG, a database that LLMs can tap into to get updated and relevant data, and provide source information. We use an LLM solution internally, which we named Copilot, to mine our operational database. For instance, you can input core file signatures to find applicable tickets in our Jira database. It has a user-friendly interface that helps staff figure out what's happening by extracting intelligence from a mountain of stored operational data.

We are using many data science tools to solve different and specific problems.

Monica: How do automation and openness relate to each other? In an open network, you have more independent elements you need to integrate, manage and control, often from different vendors. Automation is going to be more difficult to adopt but also more valuable.

Brandon: There are two types of openness. The first is about open interfaces. An operator can use an open interface to deploy a DU and CU from different vendors to increase supply chain diversity and speed up innovation.

What we are talking about here is a different type of openness. It is openness in the operational stack, with innovation coming from an ecosystem of partners, such as AMD, Intel, Dell, AWS and Google.

Autonomous networks require a one-time integration effort. This is true even for what we call a "black box," or a vertically integrated network.

Once automation is put together, it will continue to create value by managing network complexity and delivering complex services. Autonomous networks are mainly about managing network complexity.

Go back to the top

Monica: With openness, the partner ecosystem becomes crucial to success. How do you see Mavenir's role in this context?

Brandon: Our partners have good expertise in what they do. No one knows servers better than Dell. Red Hat is very good when we get into the cloud stacks, manage upstream in a Kubernetes environment, or build a security platform.

At Mavenir, we put all this together with the hardware in the telecom environment because we appreciate the deeper nuances of a telco. To solve a problem in a telco, you need to understand that telco in depth. We bring in our telco domain experience. We understand that problem. We understand how it needs to be solved. And we understand the constraints around solving that problem.

We are also educating our partners. What is the telco problem? What capabilities are out there? How do I apply a technology that is working for other enterprises in the telco environment?

Whether it is cloud technology or automation, we need to configure it, set it up, select the right components, and set the proper security policy to make it work for telco. Finding the solution requires a blending of minds and skill sets. And this is Mavenir's role.

What is special about Mavenir? It all starts at the top, with our management's commitment to openness. It is in our disruptive DNA to build software platforms using COTS hardware. We are not going to do everything ourselves; we are going to build a part of it and integrate it with other parts of the ecosystem. Part of the openness is the willingness to work with partners on a daily basis to solve problems and innovate in the telco space.

Our culture helps us get the value out of the cloud and AI. We bring cloud and data science people together and prevent them from talking past each other. This is always difficult, but we got the data scientists and the cloud scientists connected with their thinking process. And the same thing has to happen when you bring the operations guy into the room.

We have telco, cloud and data science expertise – all working on the same wavelength. When we achieve that, magic starts to happen. We bring all these cultures together; this allows us to think about problems in a different way and to start innovating.

Monica: Organizational and industry culture is a factor that we often overlook when adopting new technologies. AI and automation are areas where we can learn from other sectors. In telecom, we have specific requirements, but there is still scope for learning and avoiding reinventing the wheel. We need some humility, perhaps. Often, the culture is the bottleneck, not the technology or product availability.

How well do the operators you work with manage the cultural changes that AI and automation require?

Brandon: Operators have KPI thresholds to meet and P&L to deliver. Operators work in a heavily regulated environment with heavy performance and security requirements. We must respect that as we introduce new technologies. We have to build trust; we need to understand their problems. We cannot disrupt their operational models. We need to create value, not break things.

Monica: How do you create value?

Brandon: If you create a better user experience, reduce fraud, or improve security, you create value. You may protect your subscribers, prevent denial-of-service attacks, and reduce revenue leakage from fraud. On the operational side, better insight into the network's behavior can improve efficiency, coverage and capacity, which in turn improve subscribers' experience and sustainability,

Monica: How do you expect your AI and automation strategy to change over the next few years?

Brandon: Today, we are working towards AI everywhere. AI can help us tame network complexity, learn from historical data, and get more accurate causation, better forecasting and better operational recommendations. We will be able to personalize services in real-time, with solutions specifically tailored to the telco market's needs.

Go back to the top

About Mavenir

Mavenir is building the future of networks today with cloud-native, AI-enabled solutions that are green by design, empowering operators to realize the benefits of 5G and achieve intelligent, automated, programmable networks. As the pioneer of Open RAN and a proven industry disruptor, Mavenir's award-winning solutions deliver automation and monetization across mobile networks globally, accelerating software network transformation for 300+ Communications Service Providers in over 120 countries, serving more than 50% of the world's subscribers. For more information, please visit www.mavenir.com

About Brandon Larson

Brandon Larson is the Senior Vice President and General Manager of Mavenir's Cloud, AI & IMS Business Strategy. He joined Mavenir from Tango Networks, where he served as their Chief Strategy Officer. He is also a former employee of Mavenir, having previously run the Sales Engineering organization for America's region, where he was instrumental in many key Mavenir projects, including the world's first IMS VoWiFi, VoLTE and RCS 5.0 service launches.
Brandon has diverse experience, having held roles in sales, marketing, product line management, and program management at Nokia and Ciena. Brandon started his career as a Nuclear Engineering Officer in the U.S. Navy. He is a graduate of the Navy Nuclear Power School and received his nuclear engineer certification from the Department of Naval Reactors. He received a Bachelor of Science in Chemical Engineering with High Honors from Rutgers University.

Watch the conversation

Go back to the top

How AI, ML and GenAI can increase network efficiency today

A conversation with Blake Hlavaty, Director of Global Network Software Offers at Fujitsu

AI, ML and GenAI can help operators increase the value of their installed and new networks. With network-specific knowledge, they can empower operators with trusted solutions that increase network efficiency.

In this conversation with Blake Hlavaty, Director of Global Network Software Offers at Fujitsu, we discussed the practical adoption of automation powered by AI, ML and GenAI. Intelligent applications help operations teams navigate complex vendor ecosystems using AI/ML tools without being data scientists or knowing how to train data models. These applications embrace multivendor, multilayer network data, seamlessly integrating with existing infrastructure and eliminating expensive and disruptive infrastructure re-architectures to accommodate AI.

Monica: There is a lot of talk and hype these days on automation, AI and GenAI, so it is refreshing to talk to someone who is deeply involved in these areas and can tell us what the industry is working on and what is available to operators today. Blake, could you tell us more about yourself and what you do at Fujitsu?

Blake: I am the head of Global Network Software Offers at Fujitsu. My team develops and licenses control and analysis optimization software that we sell to network service providers. I have been at Fujitsu for 12 years and started in optical network design. I have been back and forth between sales and product management. I cut my teeth on SDN with the Open ROADM project. More recently, I have been working on open RAN architecture components such as service management and orchestration (SMO) and the RAN intelligent controller (RIC), as well as on AI and ML—not only in wireless but also in the transport network where Fujitsu has an additional legacy and heritage.

Monica: How does automation fit in at Fujitsu?

Blake: From an automation perspective, one of the big things Fujitsu has focused on is open networking. It is a pervasive concept across our transport and wireless technologies and a big part of getting to automation. It becomes even more important as new technologies such as AI and ML come in.

If you are tied to a certain vendor or don't have a strong ecosystem to drive technology innovation, you are limited in what you can do. Fujitsu understands that open networking will help operators overcome this limitation, and is a regular contributor to open networking groups like O-RAN ALLIANCE.

At Fujitsu, we provide the foundation for direct network access and management—or network abstraction. More specifically, today, we talk about network normalization. It requires access to all the network's data and the ability to manage and provision the devices. On top of this layer, we build the applications that do the analysis and use AI and ML policies to make changes to the network.

Monica: Before we delve into automation, could you tell us what the distinct roles of AI, ML, and GenAI are in network automation?

Blake: Machine learning is the brain of this process, the statistical analysis. In ML, we take the raw data and derive some knowledge or information from it using number crunching, neural networks, and classification models.

AI is the muscle: that's where the action happens. Based on the ML analysis and the information we have derived, we need to take action or make a

Watch the conversation

Go back to the top

decision. This is where AI comes in, in the algorithms and policies.

The next step is with GenAI, which is all about content creation: learning from large data sets and generating new ideas or content based on that. We see the applications of GenAI in our daily lives. It is all over the news, with ChatGPT and other players in this field, for text, image, and video generation.

We use AI as a broader concept, while GenAI is more specific. ML plays an important role that is sometimes overlooked.

Monica: What are the most promising use cases?

Blake: Fujitsu started to put a lot of effort into AI and ML six years ago when there was a lot of AI hype. As often happens in telecoms, the hype came down, and implementation came up. This is where we are now with traditional AI. Today, GenAI is adding interest to AI in general and to AI and ML use cases.

It is going to be an iterative, step-by-step adoption process, with network operators taking their time to develop confidence in the technology. In telecoms, we tend to be risk averse given the nature and importance of our function. Initially, we may start with just analysis, where you can do a lot without directly impacting the network. You are not relying on a machine to provision your network, for example. But you can use ML to derive insights from that data and use ML to gain those insights. This is where some of our early use cases are focused.

Monica: What areas are going to be the low-hanging fruits?

Blake: For GenAI, chat interfaces have started to improve customer experience and operations, for instance, in troubleshooting efforts.

You may have a network issue. You also have pages and pages of technical documents from multiple vendors where you may identify the issue. An LLM may help you index this data and bring you the solution or at least get you going in the right direction. This could accelerate some of the operations, save you time and increase network efficiency. That is the area where we will see low-hanging fruit.

Monica: How can we manage trust and reliability as we introduce automation and AI/ML in our networks?

Blake: Let's look at reliability first. If you are looking at improving efficiency or root-cause analysis, your model may not be 100% accurate out of the box. You need flexibility at the beginning because it will take time for the model to get trained and for accuracy to improve. In the meantime, you can still derive benefits – it may save you time and optimize your operations.

Building confidence is important for operators; it will take time and happen in steps. The first thing is just to analyze the data. Then, we need to take some time and build that confidence. Once I have confidence that the analysis and the information I am getting are accurate, that I have identified the cause of my outage, or that I found out why I do not obtain the efficiency I need, I can take the next step. That may be when GenAI comes in and suggests a resolution to the issue I see.

Maybe the system is not yet ready to implement the change on its own. It told me what the issue was and recommended a resolution. The human can now step in to review the recommendation and make the change.

Reliability and trust will come over time, and operators will be more willing to move forward with automation that uses AI systems.

Monica: What about the hallucinations that plague publicly available ChatGPT and other GenAI models? Understandably, many are worried that models that hallucinate are not suitable for wireless networks.

Blake: We are learning and getting feedback. The models will improve with every iteration, but in the meantime, we don't know what we don't know. Operators are watching cautiously and trying to figure out where GenAI can add value. But they are going to put an air gap right between this new technology and their network for some time.

Wherever we get an insight or recommendations, we still have a human who is going to process the output and filter out hallucinations.

Monica: AI, ML and GenAI are enablers for automation, but what are the drivers? Have they evolved over the past few years?

Blake: The drivers for network automation have remained the same, but we are approaching an inflection point that will accelerate some of the network operators' spending.

The demand for bandwidth continues to go up, but customers' willingness to pay stays flat. We have to find ways to cut costs to drive profits.

Go back to the top

The telecom workforce is aging out, and we are unlikely to have the large staff we used to have, even though networks are more complex on the wireless and transport side. Traffic, the number of use cases, and costs are also growing.

Operators need automation to address all these changes. They will need more proactive staff to manage their networks, which, with automation, will spend less time reacting to issues.

For instance, software updates to the network nodes are one use case we are looking at right now. It can be a very tedious process, especially at the beginning. With automation, we can have performance metrics that recognize when there is a problem and automate the backout.

This is a use case where automation saves time. Ultimately, automation is about reducing costs and incidents and improving reliability because we operate the network in a more repeatable way.

Monica: The emphasis you put on complexity is spot on. Our networks are so complex that automation is no longer a choice but a necessity. It is a way to manage and tame the unavoidable and growing network complexity. What is Fujitsu specifically doing in this area?

Blake: With new technologies, it is important to consider how the customer will adopt them, not just in terms of sales, business or budget. You need to think about the person using the software day to day, or going through adoption.

I have much empathy for my customers trying to do their jobs under the gun and facing increasing complexity. They need the right tools to get their work done and, specifically, to be able to adopt new technologies. We need to understand this as we develop new products. If we don't, we risk slowing down the adoption process or limiting the benefits of new technologies.

Monica: Automation is especially challenging because you are offering your customers something fundamentally new. It is not just a way to improve performance; it requires a new way of doing things, a change in the organizational structure, and new employee skills.

Network openness is also a factor. The more open the network is, the more complexity you have.

Adding more vendors, open interfaces, or services increases network complexity. How does automation fit in networks that are both more open and more complex?

Blake: In open networks, we are breaking everything apart and putting it back together with different people and products. This is something Fujitsu has been doing for a long time. Over the last several years, we have gained a lot of experience working with what we would have previously treated as direct competitors. We learned to do integration not only at the node level but also at the software level.

The interfaces have become very important. Open RAN standardization has required a lot of work, and more work has to be done in refining and defining those interfaces. We need agreement among the ecosystem players on how to build and implement the interfaces.

As we develop more partnerships in the ecosystem, we see different levels of adoption of open RAN standards and interfaces. It is challenging at the beginning; you need to be willing to work together to make the necessary changes in the interfaces.

Integrators play a huge part in creating an environment where different components can come together and in validating them.

Operators and service providers have to encourage the ecosystem players to move towards openness and meet the open RAN challenges. Automation through testing and integration will accelerate the deployment process for open RAN.

Once we have the interfaces set and the platforms in place, the applications on top will bring value and extract more benefits from automation, whether by improving the network's efficiency or delivering a new service to the end user.

Monica: How do two operators, one with open RAN and one with traditional RAN, benefit differently from automation?

Blake: The operator with traditional RAN has fewer validation points in the automation specs. However, the downside is that it is limited by the vendor's offers and capabilities. Being able to take advantage of the features and capabilities of other vendors is a reason to move to open RAN.

As the operator moves from a proprietary vendor implementation to a more open one, it will see differences in its ability to implement automation. For example, the operator may choose not to have different vendors at each level. It may choose to have a different vendor for the RUs, CUs, DUs and RICs, but then draw the line there initially. This is a good approach to get started. With time, the operator may open the network more to include more vendors.

Go back to the top

Where should operators start? We have seen different implementations from different vendors, some more open than others. We have not yet coalesced around what the best approach is, and maybe we never will, and instead we may see operators take different approaches.

Monica: It will be fascinating to see how that will work out in open RAN deployments with disaggregation. Do you think network automation will happen at the vendor or element level, system integration level, or end-to-end level?

Blake: We are still learning, and the key to doing so is for both vendors and operators to get out there and get engaged. As a vendor, we need to get engaged in customer labs and trials and collaborate with partners and other vendors because we are not going to learn all we need to know until we try it out.

In the past few years, it has been mostly theory, building what we didn't know. We are now at a point where we can implement what we learned. The learning phase with all the questionnaires and RFIs has been really exciting. Operators have been asking questions on automation in O-RAN, RIC, and SMO and wanted us to get into their labs to evaluate our software. We then bring in more vendors and introduce more interface points. This is how it starts, and then it gets moving quickly.

Monica: When you talk to operators, what do they tell you that they want to do with automation? What do they see as the main drivers?

Blake: We see automation coming increasingly into wireless networks. It is becoming very prominent in data centers as SDN takes off.

On the optical network side, where I used to work, we had automation in the commissioning, provisioning, and operation stages.

In the provisioning stage, automation may be used to get the initial configuration of a new node. We call this zero-touch provisioning – a term you will often hear. We worked on that and made many advancements at Fujitsu.

The next step is to commission the nodes. You need to validate that they are operating properly, monitor the performance metrics, and then put them online.

And finally, you can automate operations.

You can then extend automation beyond optical to other domains in the network, and you can add more automation levels.

We have seen different requests from operators. For instance, an operator may want to accelerate testing in a huge project because it is a tedious and redundant process that takes a long time. We can help them with requests like this.

Monica: Do you see a different level or type of interest across operators?

Blake: We see a different type of approach to automation depending on the operator's size.

A large operator is more likely to be interested in an AI solution. We will be talking in a room full of data scientists, with developers comparing different solutions.

A small operator may be more worried about how difficult it will be to adopt open RAN and use automation to manage it. How difficult will it be for the staff to implement and operate RIC?

Monica: Energy efficiency should be one of the best use cases in the short term. Manually switching off and on radios, for instance, is not going to deliver the efficiency that we need and that the technology can deliver.

Blake: We have seen a lot of interest in energy efficiency and sustainability. Many players are involved in developing applications, some specifically for open RAN. Government and legal requirements in Europe and the US, as well as cost savings, are big drivers for this work.

You can approach energy efficiency at multiple levels. The cell level may be the quickest path, with putting radios to sleep, for instance. You can also act at the network level, but that requires more data – and specifically drive test data – to understand the wireless environment you are working in. These are early implementations of automation, where you can start adding some ML technology to analyze and access the data.

We have seen interest in the impact of energy savings on the network and the user—i.e., the quality of service and the quality of experience. This requires more than setting a threshold to limit the impact on users. You need to consider tradeoffs. You may allow energy savings to have some impact on the quality of service and the quality of experience for some users but not for others. This approach allows service providers to offer additional savings and treat their customers differently. And this is the flexibility that some service providers are asking for.

Go back to the top

Monica: Does this mean that automation will create monetization opportunities by enabling a better segmentation of subscribers?

Blake: Some of the recent discussions have been around network APIs. We have been working in network control software, and for the last six years, we have focused on saving money for network providers. But they also need to increase revenues; this is where the network APIs will lead to new revenue streams for the operators.

Operators have so much data in their networks that they can make available to enterprises, such as retail stores, hotel chains, and financial institutions. These enterprises will pay for the right to access the APIs and use them to build new services. We still do not know how big this will be—will it be a side stream or something bigger?

Monica: Operators have always had massive data sets, but they can only now benefit from them.

Digital twin is also an area of intense work.

Blake: Digital twins can be used in many verticals, but they are a fantastic application for telecommunications. They create a digital mirror of your network. In a wireless network, you have to understand the environmental conditions—e.g., where buildings are placed and how that could impact performance and data rates. AI can bring a lot of value into this, but do you collect data, build and organize a digital twin, and where does the data come from? How do you get access to it? How do you keep the digital twin up to date? This may be the hardest task because you have to do it in real time. And this is where AI comes in to update and analyze the network model in the digital twin.

When this is done, you can monitor the network, update the model and predict performance for different scenarios. For instance, you may want to know the impact of adding a new service to a spectrum band versus another. Or you can use AI and ML to find the best way to optimize your network by exploring different possibilities.

Monica: A trusted and reliable way to update your networks is crucial to successfully implementing automation and using AI and ML. But you also need reliable data to start with.

If we look at the future, how do you see the work Fujitsu is doing evolving in this area?

Blake: We have done a lot of work in AI and ML, and over the next years, we will refine our work further in a step-by-step, iterative process on key use cases. We will identify new problems we can solve, develop confidence in our solutions, and then move on to the next one.

Over time, we will also move towards closed-loop solutions, which is where we will see the true value of AI. This will take a few years and will initially be limited to specific use cases that are easier to address.

I always listen to customers and am eager to learn and work on their challenges. My team is all about listening, learning and picking the next building blocks to deliver a successful solution.

Go back to the top

About Fujitsu Network Communications

Fujitsu Network Communications, Inc. is a leading provider of digital transformation solutions for network operators, service providers and content providers worldwide. We combine best-in-class hardware, software and services with multivendor expertise to enable cost savings, faster services delivery and improved network performance. Working closely with our customers and ecosystem partners, we design, build, operate and maintain better networks for the connected world. For more information, please see http://us.fujitsu.com/telecom or connect with us on LinkedIn at www.linkedin.com/company/Fujitsu-network-communications.

About Blake Hlavaty

Blake Hlavaty is the Director of Global Network Software Offers at Fujitsu Network Communications, where he leads the creation of cutting-edge software offers that help operators accelerate digital transformation. He has gained deep telecommunications industry expertise through various roles in network design, software engineering and consulting, and product management. Most recently, Blake has focused on network management software including open networking, SDN, cloud, AI/ML and open RAN technologies. The excitement and ever-changing nature of software products motivate Blake and his team to realize the vision of these technologies.

Go back to the top

The decomposable AI landscape

A conversation with Per Kangru, Business Development & Strategy CSO, Takai Eddine Kennouche, AI/Data Modeling Architect, CTO Office at VIAVI Solutions

AI will change the way we operate telecom networks. But what's the best way to get there? How should we introduce AI to get the benefits it promises but avoid duplication, conflicts among different AI models, and undesired outcomes? How do we know AI models will deliver what we expect?

Per Kangru and Takai Eddine Kennouche at VIAVI Solutions shared their AIOps practical experience and vision for integrating AI and ML in network operations in an efficient and streamlined way while preserving the reliability and quality of experience that are non-negotiable.

Much of the discussion was about how a decomposable AI approach allows operators to use a common framework where data and models can be used for multiple tasks and thus reduce costs, duplication and energy consumption, thus strengthening the business case for AIOps to deliver a higher operational efficiency.

Monica: AI touches many aspects of network deployments and operations, from R&D all the way to organizational culture, in addition to technology, services and, as always, ROI. In this conversation, Takai Eddine Kennouche and Per Kangru shared their experience, with Takai's more grounded in R&D and Per's deep involvement with customers. Per and Takai, please let us know more about your role in getting AI off the ground in telecom at VIAVI.

Per: I am based in Sweden and work globally to figure out where the telecom world is heading, how it is evolving, where this evolution is taking us, and how, at VIAVI, we should position ourselves in this context. VIAVI's strong foundations in AI guided us through the changes it introduced. I have been deeply involved in looking at what we can achieve—realistically, delivering substantial value to our customers and avoiding hype.

Takai: I am the other side of the coin, the lead AI specialist in the VIAVI CTO office, and I mostly do research and development. I also look at VIAVI's position in the market and our product portfolio legacy. I assess how we can assist our customers in their journey towards full automation with AI technologies as they try to reduce costs and increase the efficiency of their networks.

Monica: It is great to see that everybody talks about what you are working on—AI in your case—but sometimes you get a sense that it may be more than warranted and that the hype this generates is counterproductive. Do you see this happening?

Takai: Yes, this is the elephant in the room. Definitely, there is hype. We are probably at the peak of this hype, especially around GenAI, which has disruptive potential across the industry.

This may partly come from a misunderstanding of what these technologies are and what they are capable of. It may also come from a genuine desire for various industry players to progress in our journey towards automation. It seems this is just another one of the hype cycles that we often see in telecoms.

Our goal is to get networks to reach full autonomy, and to be self-organizing and self-optimizing. We are still far from achieving this goal, but with every hype cycle, we end up with new notions and techniques that augment our arsenal. Ultimately, we need a more fundamental and pragmatic approach to build solutions that gradually increase automation. There is hype, but there are also many moving parts to the adoption of AI.

Monica: Being pragmatic is going to be important if we want to avoid hype and focus on substance. What do you hear from your customers?

Per: AI helps us do things that we humans are not very good at. Humans get bored with repetitive tasks and are prone to make mistakes. It is boring

Watch the conversation

Go back to the top

and unproductive to have a human analyze 10,000 network clusters and come up with recommendations. It is a waste of human labor but a straightforward task for an AI model.

However, we still need a human to decide whether to accept the recommendations to improve these clusters.

AI will help us do things humans inherently are not good at volume. Humans are good at analyzing one, two or five clusters, but not 10,000.

Secondly, AI can help us see patterns that are too complicated for humans to identify.

In the late 90s, we learned about structuring the silent intellectual capital. It may take 20, 30 or 40 years of experience for someone to become an expert and be able to solve a problem in five minutes rather than spending hours or days. An AI model trained on large data sets can be reasonably good at instantiating this silent intellectual capital.

Monica: We should acknowledge that there are things that we humans are not good at and be glad that AI can do them and let us work on what we are good at and enjoy more. But how do we get this to happen in our networks?

Takai: When we talk about AI these days, we usually refer to data-driven systems that have proven to deliver value in various situations.

Because they are data-driven, data is the foundational element. You need data—good quality data. You need quality assurance when it comes to provisioning and acquiring it. You need systems that can deliver this data.

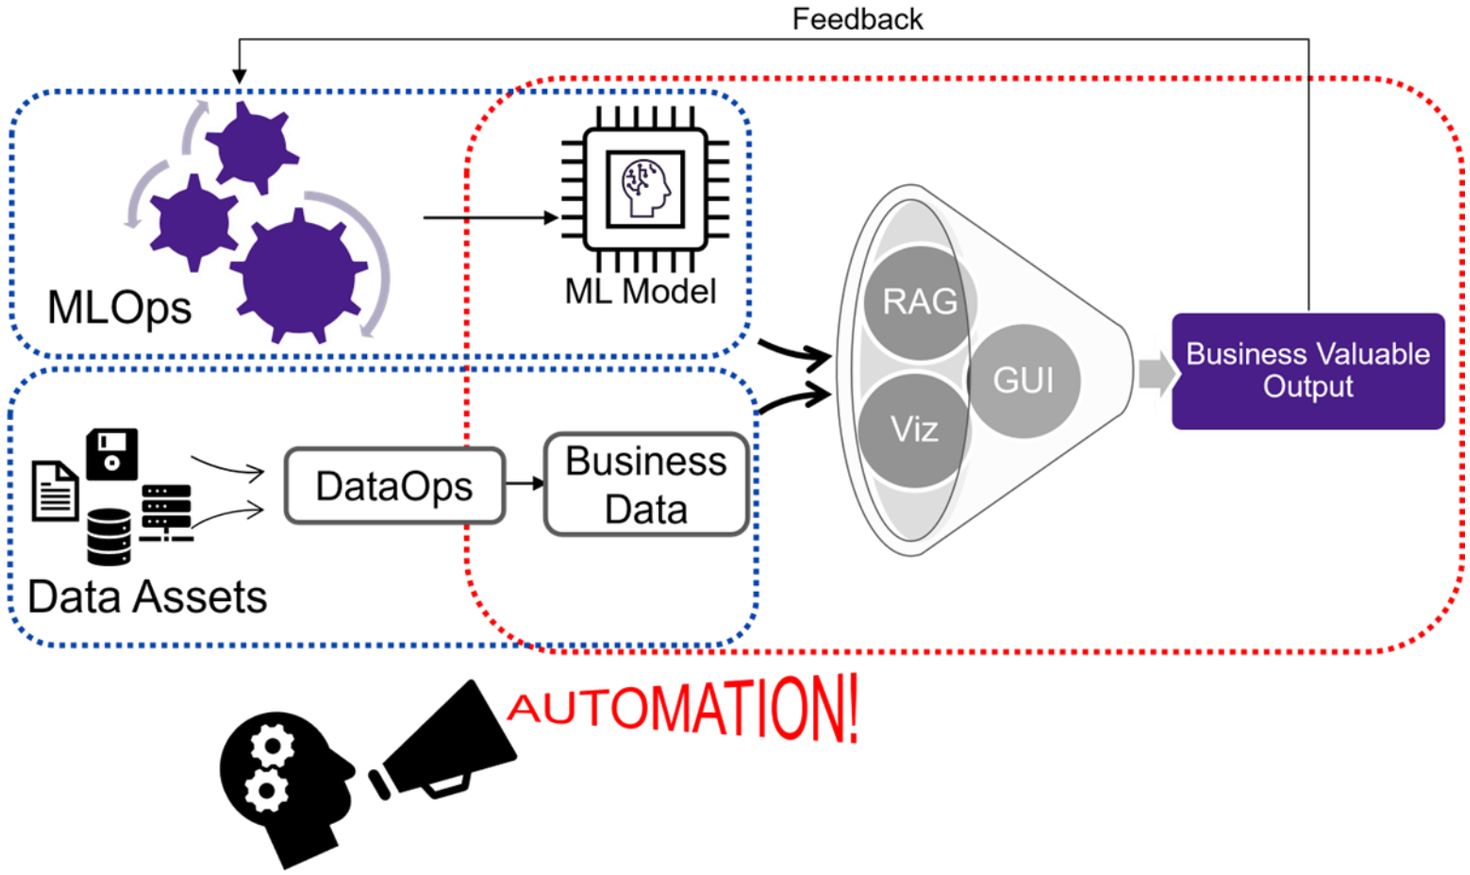

Creating business value with AI/ML ***Source: VIAVI Solutions***

This point is extremely important in telecommunication networks, given the scale they have reached. In telecoms, this translates into delivering the right modalities of data throughout the network elements, down to locations where you need to train an ML model.

The other side of the coin is talent and expertise. This is about knowing how to choose among the AI and ML paradigms and techniques that research has shown to be generalizable and implementable, as well as the evaluation metrics that can measure impact. We need to understand if pre-trained models can generalize beyond a training dataset into an operational environment reliably enough to support our telco use cases. This is particularly relevant in the telecommunication industry, where pre-training that uses large-scale, realistic domain-specific data, such as RAN-specific datasets, is limited relative to other AI/ML application domains.

This is something to be aware of when we compare the use of AI/ML to other industries. In other areas, you may enjoy the luxury of doing research using millions of images that you crawl

Go back to the top

from the web. In telecoms, it is nearly impossible for university labs to have access to similar quantities of subscriber records.

There are ways to remedy this limitation in the availability of training data. Consortiums are being formed; 3GPP is addressing AI/ML as a part of its 6G work.

Good data and good models are the foundation for the successful adoption of AI in telecoms. We need to be able to provision and deliver good-quality data to AI models. We need proven models that can deliver at scale and at an acceptable cost from the processing, complexity and maintainability perspectives.

Per: Everybody is doing various things now. Initial funding may come from special project allocations, the CTO, CIO or other parts of the organization; usually, they don't need a robust or immediate ROI.

When they try to expand and productize these trials, operators start looking at the costs of maintaining well-trained models with good hygiene and expandability. At this point, they realize that AI models have a structural cost.

Operators are probably able to manage one of them. But costs start to rack up if they have five separate AI environments with separate models, even if they use similar data sets. Will they select only one use case if this is all they can afford?

To address the cost and unnecessary duplication associated with multiple AI models, at VIAVI we propose to adopt a decomposable AI framework with reusable components and reusable capabilities. Eventually, we will be able to provide fundamental AI-driven services that can support a multitude of automation use cases. Each model is based on a reusable model. All models have the same high security, reliability and observability, and are managed and maintained in the same way.

If you want to replace one of those models, you can do that once without having to do it five times (or any number of times for the use cases you have). If you find a more efficient or cost-effective model, you can replace that one, not all the five models you have.

As the industry is moving beyond the hype peak, we need rationalization. People have stopped building AI use cases just because they are AI use cases; they now build them because they are the right things to do. And at this point, we need to sober up and face the real costs. It is time to have this discussion today.

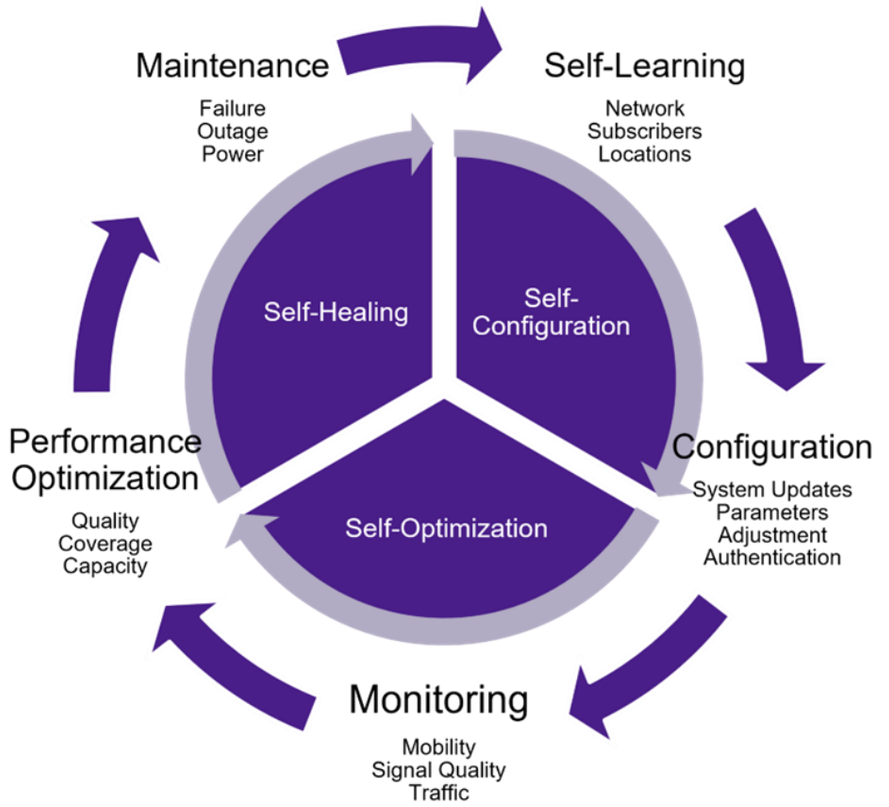

Source: VIAVI Solutions

Monica: Yes, it is time to inject reality and rationality into the discussion. The concept of decomposability and reuse is a new one for AI models. How will it fit within the ecosystem as we move to more disaggregated and open networks?

Takai: The growth in AI adoption is unavoidable and a good thing. We have a mature open-source ecosystem where AI is taking the world by storm. In recent years, we have accumulated experience that we can leverage to deliver good value to telco customers.

To meet the customers' need for automation, scalability, intelligence and optimization, we need to operate in an AI/ML ecosystem with multiple third-party providers that will come with pre-trained capabilities, using data sets from different sources, and applying different ML modeling and training techniques.

They will all need to interoperate, fit together nicely, and provide value to the customer. The old issue of interoperability will have an impact here. If the business value depends on four

orchestrated capabilities, how do you vet and test them—not just when deployed, but continuously to ensure operations run smoothly? You need to think about the output quality of a single model on one hand, and the quality of the output of a pipeline of different models in delivering the final business outcome in an operational environment on the other hand.

AI testing AI emerges as a concept that encapsulates techniques that address this issue in a scalable and automated manner. We can use intelligent systems to ensure other intelligent systems are bound in their behavior, and they do not hallucinate or go overboard. It is about evaluating and providing evidence that an AI/ML component does what it is supposed to do. It is also about identifying and generating test cases for identified failure modes and ensuring these get coverage in the CI/CD cycle.

There is a separation of concern that is decomposable that arises from the deployment of pre-trained, third-party models, and we need to address that. For instance, an ML-based geolocation function may provide input to an ML-based traffic forecasting function, which provides input to an ML-based energy-saving function. Each function has a clear concern, may be provided by a different source, and has undergone development in a different environment before integration into the operational network. Each of these has to undergo a neutral AI-driven testing process that focuses on validating each function's readiness and the whole composed pipeline to deliver the expected business outcome.

Monica: There is also a concern that AI will require more energy and effort than it is worth. A decomposable approach may help reduce both energy consumption and effort. Are there any use cases where this approach works and others where it does not?

Per: The best use case to answer this question is energy saving in the RAN. It is a straightforward use case: where can you turn on energy savings parameters in the network?

A human can easily look at it and come up with the right settings based on historical behavior and forecasts. The human process is going to be relatively slow and work for one or ten cells, and it does not scale. This is an impossible task for humans to complete in a network of hundreds or thousands of cells.

AI/ML can do it, but is the required computational effort too high to deliver a good ROI? Is the amount of CPU the forecasting model uses for the next 12 or 24 hours going to be more or less than the energy savings it enables in the RAN? Are you going to end up with a net reduction in power consumption?

If your network has a high load as a starting point, your energy-saving opportunity is limited, and the forecasting model may consume more energy than you can save in the RAN.

However, if you already run the forecasting model for anomaly detection for some other tasks, there is no additional power consumption requirement for the power-saving use case and the overall ROI becomes interesting.

This is one of the decomposable capabilities we look at VIAVI. You may be using the same functions for multiple things, and this changes the business case. The same logic we used for the energy savings use case works in other cases, too.

In a composable AI approach, we do not look at the business case for a single use case but at the combined one across use cases.

At the same time, AI may help you come up with the answers but give you more than you can actuate. You may get 10,000 network change recommendations, but your organization may be able to implement the top 10. You need to pick what works for you to avoid the bottleneck you get if you try to execute a task list that is too long.

Monica: What about other use cases?

Per: We have been working on the Radio Intelligent Controller (RIC) for a while. At the same time, we have decomposed our geolocation capability as an rApp which uses AI and ML algorithms to find the subscriber location and creates a geo-located event for that subscriber. Other apps, such as spectrum efficiency optimization or energy savings, can now reuse this data so they do not have to compute the subscriber location multiple times.

Another use case is anomaly detection as a service. You identify an anomaly, and the app can share that finding with other apps that have a capability that requires such anomaly data.

With decomposability, not only can you reuse models, but you can also share the artifacts that these models deliver, for instance, through inferencing. But to be able to do so, you need to

Go back to the top

train the model in a way that makes sharing and reuse possible.

Similarly, LLMs may be trained with the local natural language in various countries. The model is the same; the training data and the output are different.

If we go one step further, the country-specific model can be reused for specific industries within that country.

Creating a generator model that can be freely reused across different stakeholders is an efficient way to share the burden of building and training LLMs.

On the telco side, there are discussions to create a unified model that is telco-specific and that multiple operators can use. There are incentives to create such a model in Europe, and other regions and countries may join.

Monica: Who should build the underlying reusable model—and pay for its development and training? Should operators modify existing LLMs? If so, should they use open-source or proprietary ones? If they build a common one, will they be able to fully fit their data and requirements?

Takai: It is difficult to make a recommendation that works for all operators. Each should look at whatever makes the best business sense.

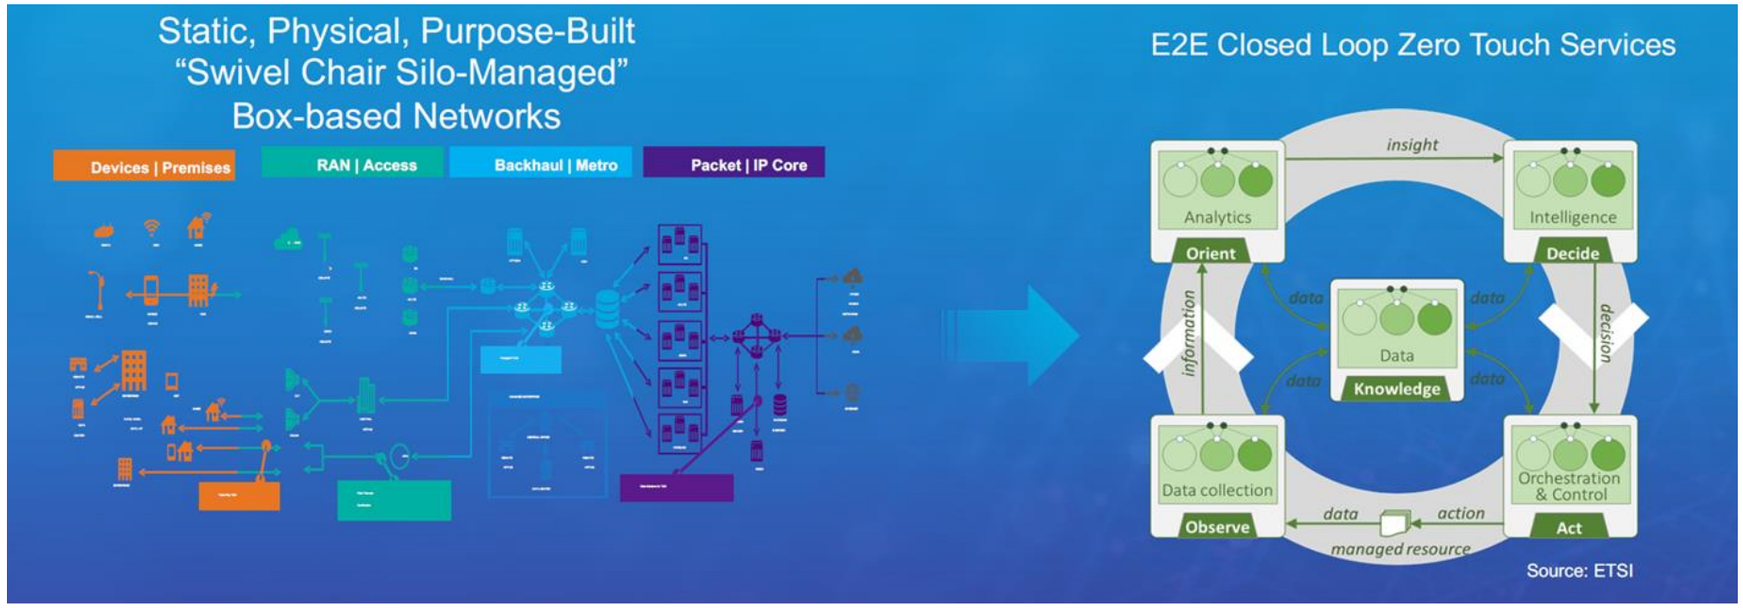

Networks evolution from purpose-built to open, zero-touch automation ***Source: VIAVI Solutions***

If you go open source, the training costs will be lower, but you may have to tolerate some biases or other issues in the model that you have to correct internally with your engineers, assuming the skills exist.

In some cases, the amount of correction that the model needs requires too much focus, effort and expertise to justify the choice of an open-source model, or the model does not deliver the same performance as their closed-source counterparts.

The cost of developing an LLM is so high that I would be surprised if an operator decides to build its own. Therefore, partnering with an LLM provider or a smaller third-party model developer to justify the investment and improve the ROI makes sense. It is still too early to know which model will be most adopted in telecoms.

Monica: How important are LLMs within the bigger picture of AI adoption?

Takai: In most cases, operators do not require generic LLMs to bring value to the adoption of AI/ML. What is crucial is to build the right data foundation, and the right system-level architecture to deploy and orchestrate ML capabilities and extract value from proprietary data. This is how operators can benefit and differentiate themselves.

Per: Within the RIC and Open RAN area, and the wider 3GPP framework, VIAVI is one of the companies that provides a solution that allows operators to generate training data for their models.

The RIC tester, for instance, can generate all the data needed to test and validate xApps and rApps

Go back to the top

before you switch to your current operational model. There are going to be many developers who need to be able to demonstrate their apps work within the operators' network.

In more mature areas, such as service assurance, we do not yet have anybody who can generate traffic data at a scale that includes all the northbound, eastbound and westbound interfaces to train an OSS or BSS model.

This raises the issue of building data assets with the right labeling. This is a prerequisite for reusing data sets and models. If an operator has five use cases, it should not have to support five teams, each with its own data maturity and data gathering project.

Monica: You need to generate data to test scenarios for which you do not have historical data, right? How do you decide which data to generate? Is the data network specific? How do you decide which scenarios to cover?

Per: Generating data requires you to fully understand what you are trying to model, what you want your model to do, and the future state of your network. For instance, you may want to fully define what a future slice network looks like when using an iPhone 20, along with other devices. If you can define this, you can get a model that matches.

You may want to generate a model with enough variations to cover a range of scenarios. I then need to generate data to test my applications and capabilities across a wide range of environments. This requires some additional computational cost, but it prevents humans from making mistakes and gets them to define what they want.

Generating data also allows you to ensure that your data is appropriately labeled. We have been using this in customer care for a long time when we ask for customer feedback at the end of a call. We use the customer score on a scale of, say, one to five to rate the service as part of the training. We can do the same to train our models.

Takai: Incorporating the human in the loop in an automated system has always been a challenge. With AI and automation, we want to reduce costs, free humans from mundane day-to-day tasks of maintaining the network and let them do creative work. Still, human feedback is fundamental to getting the models to learn.

We are not going to have a big generic solution that is going to drive our networks autonomously. Instead, we will have a well-architected system with a solid data foundation that still needs a human in the loop. The way the human provides the feedback is itself an automated data-driven process, as is the case with reinforcement learning from human feedback to fine tune LLMs.

The human role is a fundamental and disruptive element in automation—one that will enable real-world benefits.

Monica: AI requires a major cultural change in the organizations—both vendors and operators. Are they ready for it?

Takai: In telecoms, we have good reasons to be cautious, skeptical and, at times, paranoid about change, about adding cool new things to our networks. Bad things may happen.

There are well-documented and well-understood risks in running autonomous AI models trained on large data sets. We have examples in the autonomous car industry, with cars hitting pedestrians or getting out of the road. We want to avoid the equivalent of this in a telco network.

But our networks are getting more complex and have more subscribers and interconnected devices. We must move forward with automation to extract business value from our networks and manage the increasing complexity. To do so, telco players need to go through the necessary cultural change and have the internal introspection to understand how to change to introduce AI/ML and automation and leverage the existing expertise and human element.

Organizations may have to find how to best use their subject matter expertise to train AI/ML-driven automation models. They may want to introduce systems to capture metrics on administration and management patterns in an automated fashion. Their experts may need to acquire new skills, or they may have to hire additional experts. This requires a forward-looking, courageous strategy.

At VIAVI, we help our customers implement this transition efficiently by continuously delivering business outcomes through increasing network autonomy. We deliver comprehensive observability by collecting the right data, such as topology, that empowers well-defined business needs. It is a fascinating journey.

Per: We are at peak hype cycle, so much of the activity is in response to requests from top management to understand what AI can do for them that have to be fulfilled right away. But when you engage in an authentic conversation, everybody acknowledges that the organization

Go back to the top

must respond in a visible way and that they see value, but that value has to be translated across the organization. Big organizational changes are necessary to realize some of this value, and we are capable of them.

However, because of the hype, right now, many do AI because it is AI, or GenAI because it is GenAI, not because it is what they need. But, we are now seeing a shift toward choosing the type of AI that addresses a need and delivers value.

Monica: This may translate into a deeper awareness of what the needs are and what type of AI can meet them. In the short term, what do you think the most compelling uses of AI are, other than energy efficiency?

Per: Automation, anomaly detection, root-cause analysis and forecasting are the most promising areas for AI. We see GenAI used to articulate intent between functions.

Takai: I would add that event detection, correlation, and causal inference are other areas where AI can deliver value today.

AI can also help mitigate conflicts in disaggregated Open RAN deployments.

For this to work, we need a reliable CI/CD/CT cycle to ensure that these systems continuously deliver value and mitigate issues inherent in data-driven learning. Adversarial testing and digital twin are powerful frameworks for deploying AI and ML in specific, value-generating use cases.

Adversarial testing is an iterative process that can enhance the robustness and security of AI models. By intentionally introducing adversarial examples, i.e., inputs designed to cause the model to fail, developers can identify vulnerabilities and improve the model's ability to handle unexpected or malicious data.

Per: Self-governance is also important as telecoms is a highly regulated industry. This may limit the ability of operators to use hyperscalers' cloud services, especially for deep networking functions, even if they find them attractive.

Monica: But even if regulation is not an issue, operators may worry about lock-in from hyperscalers.

Per: Some operators are making very strong commitments to hyperscalers. There is no one-size-fits-all blueprint that works for everyone, and each organization has to decide based on its use cases.

Takai: The GenAI trend illustrates what lock-in risks operators may face. Telco companies need to carefully assess the ramifications of delegating the intelligence of their networks to LLMs built, maintained and deployed on third-party infrastructure and with external data sources.

Monica: What is the most important thing operators should do to get to a composable AI landscape?

Per: First, you need to have an idea of what is the possible value you want to realize.

Second, you need to prove that that value can technically be realized.

Third, you need to identify and sort out any impediments to implementation. It could be cultural change or other things.

At the end of this journey, you should be able to realize long-term viable value. If you do not follow this journey, you risk having a solution that delivers interesting things but is useless, or one that requires investment but does not deliver value.

Takai: Change management is the key. And there are no shortcuts. Embracing a native AI/ML approach as a core value-generating element of a larger set of network automation capabilities requires serious work, and deep architectural and organizational changes to succeed.

Watch the conversation

Go back to the top

About VIAVI Solutions

VIAVI Solutions is a global provider of network test, monitoring and assurance solutions to communications service providers, enterprises and their ecosystems. Our solutions deliver end-to-end visibility across physical, virtual and hybrid networks, enabling customers to optimize connectivity, quality of experience and profitability. VIAVI Cloud-native Artificial Intelligence for Network Operations (AIOps) portfolio helps CSPs, DSPs, ISPs, Utilities, and Enterprises accelerate their digital OSS transformation journey by evolving their legacy Network Operations Center (NOC) to a Dark NOC.

About Takai Eddine Kennouche

Takai Eddine Kennouche is a highly accomplished Telco AI specialist at VIAVI Solutions. With a PhD in Electrical and Computer Engineering and expertise that spans cognitive radio, network automation, anomaly detection, predictive maintenance, and intelligent resource allocation, Takai Eddine Kennouche has accumulated a deep understanding of AI/ML applications for network testing and optimization. He is currently leading the development of scalable data-driven solutions within VIAVI's CTO office. He focuses on enhancing network performance, reliability, and testability through cutting-edge technologies. Passionate about building autonomous networks and realizing Telco AI use cases, he leads initiatives that explore AI/ML's transformative potential in telecommunications.

About Per Kangru

For more than 20 years, Per Kangru has been at the forefront of the telecom evolution, driving innovation and focusing on allowing operators and equipment vendors to develop, deploy and operate 3G, 4G and 5G networks. Per has been engaged in industry forums and standardization organizations over the years, contributing to the development of new methodologies and solutions to allow vendors and operators to better develop and deliver innovative new services. Most lately, Per has been engaged in the ETSI ZSM ISG and the evolution to the fully autonomous network. This includes zero-touch service assurance in the development of AIOps capability, linking all of Day0 to Day2 into a fully automated framework for both green and brownfield scenarios. Over the years, Per has been a co-author of several books on 4G and 5G, a single innovator on several patents, and has incubated and developed several new product lines. Per has been involved in M&A transactions and developed strategic partnerships with industry players. Per is based in Sweden, and when he is not thinking about technology, he enjoys downhill skiing and cooking food. Per studied atomic physics at Uppsala University in Sweden before engaging in telecom.

Go back to the top

AI everywhere, from the edge to the cloud

A conversation with Cristina Rodriguez, VP Network & Edge Group and GM Wireless Access Network Division at Intel

AI-driven automation will change how we run our networks end-to-end, from the access in the RAN all the way to the core, from the edge infrastructure to the cloud. To get there, we need the right silicon and platforms. Are we there yet?

In this conversation, Cristina Rodriguez at Intel and Monica Paolini at Senza Fili talked about the network transformation that AI and automation enable and how the industry ecosystem is already working on multiple use cases to make wireless networks more efficient and sustainable.

Monica: AI is going to be a pervasive presence in our networks, with its role expanding with time, but we already have many tools we have started to deploy. Cristina, you have been involved in many trials and deployments that involved AI-driven automation and optimization. And what is your role at Intel?

Cristina: I am Vice President and General Manager of the Wireless Access Network division at Intel. I own the RAN business and bring together 5G, the edge and AI in an open virtualized RAN.

Monica: That's an impressive remit, especially with the transformation towards open and virtualized networks. When we talk about AI, some people refer to it as an old technology being replaced by the more recent GenAI, while others consider it something that will become widely used only in the future. You take a different view: AI will be everywhere, but you focus on AI in the present and what it can deliver today.

Cristina: For sure, AI is going to be everywhere in the future. From the cloud to the core, to the edge, to the RAN – everywhere in the network, everywhere in our lives. But we have already made a lot of progress and have much technology we can start using and continue to evolve to meet the needs of society and industry.

Monica: We are going to introduce AI and ML gradually. We will not get AI everywhere overnight.

Cristina: They are complex technologies. AI will replace the way we do things today. In networking and in the RAN, there are algorithms that have been used for a long time and operational practices that have been around forever. AI will replace them, but it will take time. You have to test what you have created, you have to get the data, you have to create and train the models, and, finally, apply and execute these models.

You need to develop trust in these models. This will require testing and fine tuning before you switch them on.

Monica: One of the challenges of AI is that it forces us to do things differently. A big cultural change accompanies the technological change.

There are many use cases for AI in wireless networks, and more specifically in the RAN. Which ones do you think are the top ones?

Cristina: There are use cases that we can deploy today. We have done proof-of-concept trials and demos with operators to demonstrate the benefits of AI. We can apply AI at two levels in the network infrastructure.

The first level is for functions that are not specific to the RAN. Power management and predictive maintenance are two areas where AI can improve efficiency.

A second area where we can apply AI is radio-specific algorithms, such as link adaptation, channel estimation, and beam weight forming, which are typically at Layer 1 and Layer 2.

At MWC24, we showcased work we have done with Vodafone to demonstrate AI-assisted network slicing for radio resource management. We will be talking more about network slicing in

Watch the conversation

Go back to the top

the coming year because it is important for the industry. Not only were we able to meet the SLAs, but we also showed an impressive 2.5 times improvement in radio resource efficiency over manual slice management.

Also, at MWC24, with Deutsche Telekom, we demonstrated a significant 2.47x throughput gain and 3.5x delay reduction at the cell edge with AI-assisted beam management.

With AT&T, we showed a 16% reduction in energy consumption with a frequency-scaling AI system. These energy savings were remarkable because they were over a policy-based approach that already implemented frequency scaling based on a typical historical traffic pattern.

We also worked on AI-driven energy savings with Ericsson. We demonstrated up to 20% energy savings using Intel Xeon's dynamic power management using the Xeon. AI takes the best capabilities that we have in our CPU to the next level.

Monica: Increasing network energy efficiency and sustainability are possibly AI's most impressive use cases.

One thing you need to do when adopting AI, whether for energy efficiency or other use cases, is to ensure that the model delivers the benefits you are looking for. For instance, you may use a digital twin to see how the model will behave in a network before you deploy it.

Cristina: That is a very important part of this journey. You create and train the models, then test, fine-tune and re-test them in the lab. You can collect and feed more data to the model. As the model's performance improves, you will be ready to deploy it in a commercial environment.

Monica: Could AI harm the network rather than improve it?

Cristina: You want to prevent AI models from harming the network. We have conversations about this with many operators. Today, networks work very well. This is a reason why change takes longer than we might expect. You want the network to work when you want to make a phone call or use your data connection.

Operators do not want to jeopardize this and have to meet very stringent performance and quality of service KPIs. So, they repeatedly keep testing and collecting data in every part of the network infrastructure, deploying a new solution.

Monica: What are the use cases where the adoption of AI is most promising?

Cristina: Dynamic network configuration is the first that comes to mind. Operators can use software to allocate hardware resources in real time. It gives them a more granular control of their resources.

With energy management—possibly my favorite use case—we have all the real-time telemetry data to get visibility into power consumption today. You can automatically control core frequency voltage or dormancy levels and adjust them according to your consumption targets.

In traffic steering, AI can help correct coverage and capacity problems in the RAN and, as a result, improve customer experience and spectrum efficiency. This helps operators improve the ROI on their spectrum investment and increase network efficiency and reliability.

Predictive analytics and controls are another case where AI can be used. It helps operators predict and react to changing network conditions before they occur. AI models can recommend actions based on historical KPIs to retain network reliability. Improved anomaly detection will also help detect network issues.

Finally, network slicing will become increasingly important as an opportunity for operators to get new revenue streams and for users to get new services. Here, too, AI will improve performance, efficiency, reliability and security.

Monica: Operators have always had vast amounts of data, but they were unable to use it effectively and in real time to optimize network performance. Is AI going to change that?

Cristina: Today, much of what happens in the network is based on the experience of operation engineers. They make decisions on how to configure or provision the network, when to do maintenance, and how to reallocate resources. We also have access to massive data from years of network operations that we can analyze and use to make better decisions in conjunction with the expertise of network engineers.

Monica: The combination of staff expertise and historical data can be very powerful in driving not only real-time network optimization but also in predicting network performance and issues. Without AI, it may be hard to discern trends and learn from massive data sets.

Go back to the top

Cristina: Predictive maintenance is one example of that. You see that maintenance work is needed, and you can do it before you have a network issue. But you may also see that you do not need any maintenance work yet and can wait.

Or, if you see congestion coming, you can manage resources to reduce or eliminate its impact. You look at how you dealt with the issue in the past and pick the best course of action.

Monica: You correctly pointed out earlier that networks work very well. Does this create not just caution but also risk aversion and delays among operators?

Traditionally, operators ensured network performance and reliability with overprovisioning, but this is no longer a practice they can afford, and AI and automation allow them to avoid this wasteful approach.

Cristina: With AI, operators can right-size the network based on the conditions they face and improve efficiency without compromising reliability and security.

Operators want to improve the TCO after making a tremendous investment in 5G spectrum and infrastructure and get new revenues with network slicing or other new technologies. AI will play a crucial role in these efforts.

Monica: Is AI necessary to get network slicing off the ground?

Cristina: With AI, you can automatically allocate network resources to meet the required SLAs without overprovisioning. You can do this in the background in a cost-efficient way that makes network slicing affordable.

Monica: AI will help us manage increasing network complexity, but it is itself a complex technology we introduce in our toolbox. How do we manage potential conflicts across solutions that use AI?

Cristina: AI is a complex technology, and we are only starting to scratch the surface of this complexity now. As it becomes part of our daily lives, it will become more natural and less mysterious.

As we move forward and have AI more widely deployed, we need to test to make sure there are no conflicts among AI models deployed in different areas. In cases such as managing dormancy in the Intel Xeon energy efficiency use case, the risk of conflict with other areas is limited. Still, there is a wider potential scope for conflicts in other cases.

Monica: AI gives you better tools to manage complexity across the network, but specifically in the RAN, which is your focus area.

Cristina: Open RAN and AI are meant for each other. AI is a critical technology that assists us in the deployment of Open RAN because it is based on an end-to-end—from the core to the edge and the RAN—software-defined network supported by the work of many developers in the ecosystem. This software-defined architecture, underpinned by our CPUs, allows you to port, upgrade, move and refine faster. AI is built in and gives you a further boost. We embraced the open virtual RAN architecture and are committed participants in a rich ecosystem to accelerate innovation in the network with AI in collaboration with our partners.

It is an ecosystem that feeds itself. With the large participation and open interfaces, we can bring value to operators, which widens the adoption of new techniques and ideas.

Monica: The adoption of AI can then progress at its own pace, without a dependency on the development of 5G or 6G.

It is an opportunity for innovation and improvement in network efficiency that will, however, require much work, starting from POCs to large commercial rollouts.

Cristina: Yes, but there is a lot excitement right now, and it is fantastic to see what we can show in the labs, in POCs and all the way to commercial deployments. We have demonstrated that the opportunity is real.

We also have to be careful about what we add to the network and, specifically, to the RAN. The RAN has tight performance-per-watt requirements, and we need to be very careful about the resources we use.

At Intel, we are aware of this and add components only when strictly necessary. Our fourth-generation Xeon, codenamed Sapphire Rapids, is very successful. At MWC, we announced the sixth-generation Xeon, codenamed Granite Rapids. They allow the industry to run AI in the RAN from Layer 1 in a virtualized way and support built-in AI acceleration in the CPU. Critically, we provide the instruction set and the reference software for everybody to run AI in an energy-efficient way. In our POCs, we have demonstrated that, with Intel Xeon scalable processors, AI does not add a burden to the power consumption in

Go back to the top

the RAN, and you do not need GPUs or additional PCI cards.

Monica: How did you keep complexity to a minimum at the component level?

Cristina: We want to reduce component complexity from a supply chain point of view. You have different models, different software tools, and different approaches depending on what you want to do. With Xeon in their server architecture, our customers have what hyperscalers have used for many years to scale up successfully, and they do not need anything else. In telecoms, they can write and run their software, as everybody in the world does.

And generation after generation, their investment is protected because they can keep running their software. I am very proud of our customers who tell us that moving to the next generation was an out-of-the-box experience. They only had to recompile; everything worked as before with improved performance.

This is super powerful because you don't want to start over every year and a half or every two years. You want to be able to take advantage of everything that you invested in the previous generation.

Monica: You made the argument that having all the AI components built in is an advantage because you don't need anything else. Could this approach preclude using other AI optimization solutions that may present conflicts? Or could you be stuck with a solution that may not meet your needs?

Cristina: This goes back to what you need. In the case of video, for instance, you are not going to train a model on video content; you do it somewhere else. If you want to reduce latency and are sensitive to performance-per-watt, you need to optimize the RAN, and there our AI acceleration will help you.

At Intel, we understand the requirements because we have been doing this for a long time. We understand different use cases; we have been working closely with operators and their teams and many partners. We know the RAN and AI, and we know how to use AI in the RAN and how to dimension for it.

You can choose to train the models, but you don't need to do that because you have already what you need. And you don't want to pay for the power consumption training models require. We provide what you need when you need it.

Monica: Building and training a model requires not only power consumption but also money and effort to build. So, you want to avoid reinventing the wheel and customize what is already available.

Where in the network should AI models be?

Cristina: Everywhere. You have to have AI everywhere. If latency is crucial, you need to run AI models at the edge on the vDU. You cannot go back to the data center or the vCU. Because you need real-time data and decisions, you cannot afford to send data to a centralized unit or the data center. You need to make Layer 1 and Layer 2 decisions at the edge and dimension your equipment accordingly.

Monica: Isn't this going to raise costs at the edge? And what do you need at the edge? You may train your network in a centralized location and then fine-tune instances at the edge.

Cristina: You will typically do the model training in the data center, but once you have done this, you can run the models at the edge.

You then have to dimension and provision the edge network based on your target KPIs. You don't want to throw money away to get expensive resources you do not need, and that increases cost, energy consumption, and complexity. This is where Intel Xeon helps you because AI is built-in.

Monica: Understanding an operator's specific needs is crucial to successfully deploying AI. How should operators go about identifying these needs? What best practices are emerging?

Cristina: At Intel, we have been in a unique position in the open RAN deployments for a few years. Today, nearly every deployment of a virtualized RAN is happening on Intel Xeon processors. We have a lot of first-hand experience in software-defined networks and close relationships with operators' teams and the RAN and networking industry.

We see a step-by-step approach. For instance, in the energy use case, you start with an AI application you run in the lab and use for a POC. You then feed all your data, create, train and fine-tune the model. You keep testing the model until you trust it, are comfortable with it, and are ready to move on to a commercial or field trial. You need to trust testing data and adopt a new mindset because it is a big step forward.

Go back to the top

Monica: It will be a gradual process, no doubt. Operators, vendors and other ecosystem players all need to learn to trust AI models.

We mostly talked about the benefits that AI can deliver today. What's your vision for the long-term impact of AI on networks?

Cristina: AI will be everywhere in our networks and enable zero-touch automation. Networks will be able to self-learn and self-heal and become self-sufficient. Autonomous networks will require no human intervention. It will not happen overnight but gradually, in stages. And I look forward to that.

Watch the conversation

Go back to the top

About Intel

Intel is an industry leader, creating world-changing technology that enables global progress and enriches lives. Inspired by Moore's Law, we continuously work to advance the design and manufacturing of semiconductors to help address our customers' greatest challenges. By embedding intelligence in the cloud, network, edge and every kind of computing device, we unleash the potential of data to transform business and society for the better.

About Cristina Rodriguez

Cristina Rodriguez is Vice President in the Network & Edge Group, General Manager of the group's Wireless Access Network Division (WAND), and General Manager of the Austin Design Center at Intel Corporation. She leads Intel's efforts to provide innovative wireless access solutions in both traditional and cloud-native networks, enabling the RAN of the future by bringing together 5G, and the build out of the Edge and Artificial Intelligence. Rodriguez has full P&L ownership, responsible for strategy and business plans, technical product marketing, roadmaps, product definition, forecasting, customer relationships, investment decisions and GTM. Cristina has more than 20 years of experience in the networking industry. Her vast experience covers startups through large corporations like Intel. Her educational background is in electrical engineering and computer science.

Go back to the top

Rethinking AIOps – It's all about the data

A conversation with Rick Fulwiler, Chief Solutions Architect at NETSCOUT

In network operations, we are always looking for areas for improvement. How can we manage applications to best serve end customers? How can we use end-to-end visibility of application and service performance to improve customer experience? Will anomaly detection for automated pattern discovery and prediction improve the quality of experience? What is the best way to use software-driven networks to support new services and monetization strategies? Rick Fulwiler, Chief Solutions Architect at NETSCOUT, addressed these questions in this discussion on how AIOps is becoming the focal point driving transformation in network operations across all lines of business—IT, enterprise, cable, security and wireless—and how its success requires rich and reliable network data.

Monica: AIOps is emerging as a key component in driving successful AI-driven models to optimize network operations. It is becoming increasingly clear that the foundation of its effectiveness depends on a traditional, foundational source: good and reliable network data. Rick combines expertise and experience in both the data and the AIOps areas. Rick, can you tell us how you got to focus on these areas?

Rick: As the Chief Solution Architect working for the CTO at NETSCOUT, I focus on data curation for AI pipelining, and all things 5G standalone and analytics. I came to NETSCOUT in 2015 via the acquisition of Tektronix Communications, where I was vice president of the EMEA sales in the UK office.

Monica: How does AIOps fit in in telecoms?

Rick: AIOps is relevant for many reasons. In telecoms, we are going through a rough patch. Revenue, growth and profitability are in decline. Operations costs have started to rise substantially, especially with the rollout of 5G standalone.

How can CSPs operate networks more efficiently, drive down costs and look for more efficient revenue pathways?

CSPs are looking at automation technologies to cut costs and boost productivity. AIOps provides a new competitive edge that the industry needs to tackle some of the automation challenges and can drive untapped data monetization revenue streams.

Monica: This is an area where we can learn from other sectors. How are we doing in comparison to other industries?

Rick: The initial adoption in telecoms has been slow, but AIOps got its start in the traditional enterprise market space and is ahead in leveraging it, along with GenAI.

With the migration to 5G standalone, increased network complexity and cloud-native environments, telecoms have started to leverage AI and AIOps, and the need for new revenue pathways is accelerating this change.

Monica: CSPs have tight requirements because they operate critical infrastructure. But there is also a sense of urgency to deploy AI, despite the concerns for reliability and security. Is the technology sufficiently mature?

Rick: Over the last few years, new AI use cases have emerged, especially using GenAI. When we think of AIOps, we generally gravitate our thinking toward the modern IT organization that combines machine learning, big data analytics, and other AI/ML techniques to automate and improve IT operational efficiency.

In telecoms, AIOps is not just a buzzword; I believe it is going to be a lifeline. It will empower CSPs to navigate uncharted waters, adapt to new network architectures, and redefine organizational innovation.

AIOps use cases include improvement of network operations efficiency in the RAN, MEC and core;

Watch the conversation

Go back to the top

enhancement of customer experience; insights into customer behavior and preferences; improvement of Net Promoter Scores (NPS); identification of unique traffic patterns and anomalies (e.g., heavy users on FWA who cause spectrum exhaustion via hosting or illegal usage sharing); driving data monetization; network slicing assurance for enterprise customers; and roadway traffic analytics for usage and pattern prediction and planning.

In the last six to seven months, more carriers have started to look at data exposure via open APIs, using CAMARA gateways and partnerships to drive revenue adjacencies.

There are many avenues where carriers can start applying AIOps and GenAI to explore new pathways for revenue creation.

Monica: AIOps provides a framework for AI to work in different use cases. One thing that all use cases share is that they need data from the network. No matter what the use case is, data is fundamental. Don't operators have the network data they need?

Rick: At NETSCOUT, we call network data the hidden gold, which must be mined and refined. We all become enamored with the cool AI parts of AIOps, but we must avoid losing sight of the critical data that drives the AIOps engines.

In 5G standalone networks, the hidden gold needs not only to be mined but also refined to become consumable by the AIOps pipeline. That is critical.

We don't just take network data as is. We add context and meaning to it, and define the relationship to other data being collected. We are not throwing packet-level data in the pipeline or into a data lake and expect magic to happen.

We saw this a few years back, with carriers grabbing the network-level data and throwing it into a big data lake with some data scientists. Did cool things start to happen? We have seen that movie over and over again, and it always ends the same way. They end up utilizing lots of resources, both labor and hardware, without the expected results.

Network data also needs an important refinement process, just like gold does; it needs context to define what network quality is being provided to each subscriber—the who, what, when and where along with the network conditions that affect service quality. You also need that end-to-end view—the RAN-MEC-core visibility—all packaged together and available for consumption for the AIOps pipeline.

Monica: As the cost of dealing with large amounts of data decreases, we may get lazy and feed all available data to AI models. But we need to find data that is not just accurate, but also relevant. Filtering out data that is not relevant is not trivial.

Rick: At NETSCOUT, we have over 40 years of telecom experience to help us understand how to solve network and subscriber problems, and which areas and levels of engineering expertise you need for specific issues. Traditionally, you would have a team of people with core competencies to solve hard network and subscriber problems. We leverage our telecom experience to identify the curated aggregations and correlated data relevant to the AIOps pipeline.

Classic traditional service assurance requires a lot of data—for deep-packet decode, for instance. AIOps does not need all the packet-level information; it only needs the golden nuggets in those packets that are relevant to solving problems. The secret sauce is understanding what data is relevant and correlating the data required to solve specific use cases.

Monica: Relevance depends on the context of what you are trying to do. The same data may be relevant in one context but not in another. And as networks become more complex with virtualization, disaggregation and 5G standalone, selecting the relevant data will be an even bigger challenge.

Rick: From SS7 to 2G to VoLTE in 4G, all the data you needed was available through a packet broker via physical tapping mechanisms, along with the trace-port data from the RAN elements.

In a 5G standalone network, the migration to containerized network functions and Kubernetes clusters makes data collection and packet acquisition much trickier. This is a challenge we addressed. At NETSCOUT, we have been working on this for the last three-plus years with multiple NEMs.

We had a press release in October of 2022 with Ericsson and one of our customers, Swisscom. This was the first example of a true virtualized 5G standalone network utilizing Ericsson virtual taps (vTAPS). Ericsson provided the 5G network and their vTAP and linked them directly to our virtual probing architecture. It was not as easy as you may expect. There were a lot of conversations among the engineering teams to ensure we retained the integrity of the 3GPP packets and did

Go back to the top

not lose any data and timing, especially as we were moving to a virtualized network, where the plumbing is much more complex with network address translation and encryption on the service-based interface (SBI).

We are working with other NEMs, too, on the same process. We also developed a service mesh plug-in that can be used in cloud-native networks.

Along with the service mesh plug-in, we also released extended Berkeley Packet Filter (ePBF) support that allows us to get packet data at the Linux kernel level. This provides visibility to the 3GPP packets from the container network functions and lets us see packet traffic from other applications running in the same environment. This approach is better suited to a cloud-native environment where you rely less on vTAPs from individual vendors.

Unfortunately, there is no vTAP standardization, so we have to coordinate and negotiate with all the vendors we work with.

I believe eBPF is the way to go. It has been a journey with barriers to overcome and a breakthrough at the end. We are now taking this native-cloud path with customers.

Monica: At the beginning, you may need to work more with the vendors your client uses, but eventually, you can have a solution that works smoothly across networks with different vendors.

Rick: Sometimes, NEMs are nervous to turn on eBPF. They can still rely on their virtual taps, and we provide a universal packet converter because each NEM is a bit different. For instance, Nokia's vTAPS may be a bit different from the Mavenir ones.

Our 5G packet converter can convert and send the data to our COTs or virtual probe, and support a true multi-vendor vTAP approach.

Monica: How important is standardization to accelerate the deployment of multi-vendor networks?

Rick: There are 3GPP specifications that all vendors use. These specifications allowed us to support 2G, 3G and 4G networks directly off a packet broker. But NEMs could add proprietary extensions, sometimes in response to requests from carriers, and we also support them.

A cloud-native architecture allows unique ways to transport and acquire packets and feed them to an external probing architecture. This is where we are dealing with non-standard processes, and it would be great to have a standard for vTAPS with set header information, for instance.

ePBF may be a good way to get around some of the unique nuances of different versions of virtual taps.

3GPP specifications do a good job of defining KPIs, but some carriers want to use their own definitions. For instance, they may want to modify the equation for the VoLTE SEER KPI. In such cases, we made our solution flexible so carriers can create their own KPIs/KQIs, formulas and metrics.

Monica: Let's go back to curating data to feed to AIOps. What approach do you suggest to network operators to curate their data using best practices?

Rick: We look at packet-level data that may come from 2G, 3G, 4G, 5G non-standalone or 5G standalone, from the MEC, from a fixed or a wireless network. Raw packet data comes at a very high volume: gigabytes of traffic, especially in the user plane.

We built sensors that collect data that can be curated for AIOps [see graph on the next page]. Unlike traditional service assurance, we do not need to provide all the information; we provide only the right information for the specific use cases. We use a filtering process that reduces the volume of the data that the AIOps engine receives.

Different use cases may need different data. So, we stream data through Kafka topics or other mechanisms to feed the data to the AIOps engines. We think of this process as adaptive data curation and correlation using embedded knowledge about the data set.

We don't just provide the data; we also provide the data with context that we define before we stream the data to the AIOps engine. We can also look for specific outliers and anomalies as part of the pre-work before hitting the AIOps engine.

For instance, we may look at RTP media in Voice over New Radio (VoNR) or even VoLTE to identify and analyze call gaps. We can see each gap's beginning, middle, end and duration, and if it is happening on the uplink or downlink. We can do a cause code analysis on it that can point to different places where we can look for the root issue–IMS signaling, packet gateway, handset firmware issues, RAN coverage, RAN ping-pong handovers, or other issues.

Go back to the top

All of this pre-analysis can be packaged up as an intelligent outcome and sent to the AIOps engine for processing with other data sets. Because AIOps algorithms thrive on well-curated data, this allows them to generate more precise insights, leading to better anomaly detection, classification, and prediction.

Monica: Sensors look at the real-time data. Can you also use them to predict future data?

Rick: The sensor itself does not predict future outcomes. You can get predictions as the sensor data goes through the AIOps pipeline. This is not something we do at NETSCOUT, but many AIOps vendors do that by leveraging the data we provide.

Sensors give us a highly scalable, programmable software component that can run on COTS hardware or in virtual environments. They collect packet-level data that is specifically tailored to the AIOps pipeline.

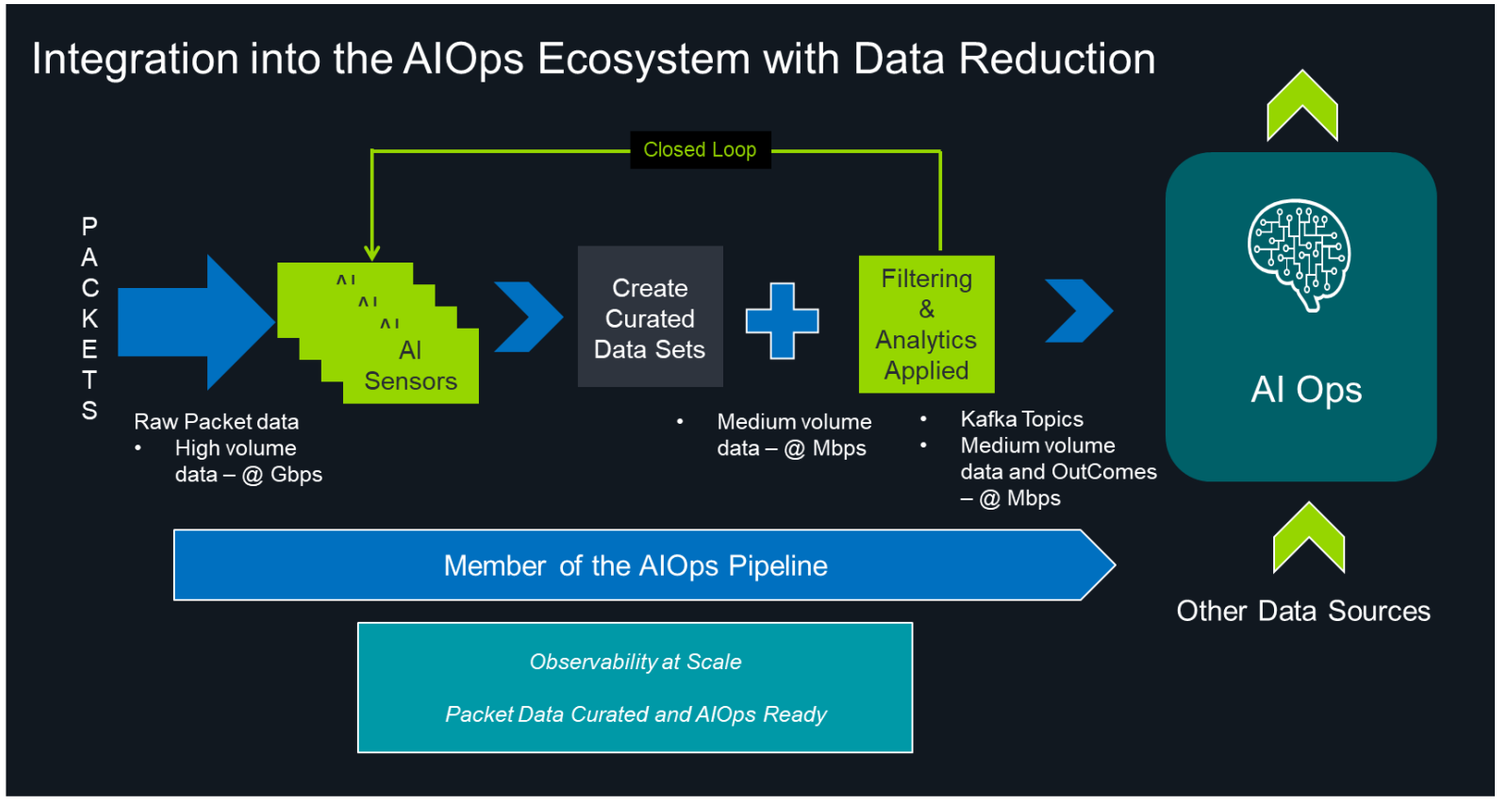

Analytics and Advanced Feature Extraction with AIOps *Source: NETSCOUT*

The sensor data is more refined and driven by specific use cases defined within the AIOps pipeline. This approach creates a filtered and curated data set, sent to the AIOps environment via a stream engine that can reprogram the sensors to collect additional or different data for a specific set of use cases. This gives us the capability to continuously refine what we are looking for in the network.

Again, the sensors do not collect every possible data set; they only collect the right data for the relevant use cases. We can also have a REST-based feedback loop from the AIOps engine back to our streaming interface to reprogram the sensors to request new data they need for any number of use cases. As a result, the AIOps engine could generate a request for more granular data to increase the AIOps' efficiency and speed or drive different use cases.

Monica: Where do you put the sensors within the network?

Rick: The sensors are placed at the standard 3GPP tapping points. In a 5G network, you can tap both the user plane and the control plane traffic off the SBI using eBPF or NEMs vTAPs. You also can use traditional tapping points for 4G. These software-based sensors are much smaller than traditional probes and monitor the interfaces that the use cases require.

Monica: What is the best use case that illustrates this new approach for the collection of real-time network data?

Rick: 5G network data analytics functionality (NWDAF) is a good example of an AIOps use case.

NWDAF was introduced in 3GPP Release 15 and later deconstructed in five sub-functions in Release 17. As defined, NWDAF gets data directly from each 5G network function by queries. These queries place an additional resource strain on the Container Network Functions. The second graph shows how AIOps can help fulfill NWDAF and additional use cases by providing curated data via the AIOps pipeline.

Go back to the top

AIOps places no resource constraints on the 5G Network Functions, nor does its implementation require a complex architecture, while having a full view of the network from RAN to MEC to core. As a result, the AIOps environment can adapt to NWDAF use cases with a much more pliable architecture while under the service provider's command and control.

It is not an off-the-shelf box; it utilizes carrier AIOps as the brains, running business-relevant use cases to the carrier. It allows carriers to still have tight control over the 5G SA network and provides a platform to migrate to automated control as they see best.

AIOps sets up a nice crawl, walk, and run strategy that allows carriers to go through an adoption period to learn how to leverage the derived statistical network resource usage data the AIOps calculates to improve the service operation efficiency of the 5G SA network. This adoption period can have multiple phases as carriers tune the AI algorithms within the AIOps engine to fit their business parameters and their level of comfortability for automated network configurations based upon dynamic network or service conditions.

This process could start in the RAN domain, where the AIOps engine can drive dynamic radio topology, power control setting on mmWave, interference mitigation techniques, energy efficiency improvements, SLA slice violations, orchestration mitigation and so on.

Later, carriers can introduce additional use cases, bringing in dynamic resource allocation in the core.

The key aspect is to provide carriers with the highest resolution of smart packet data through the AIOps pipeline through a platform that leverages the AIOps engine for more flexible NWDAF and other use cases, including monetization ones.

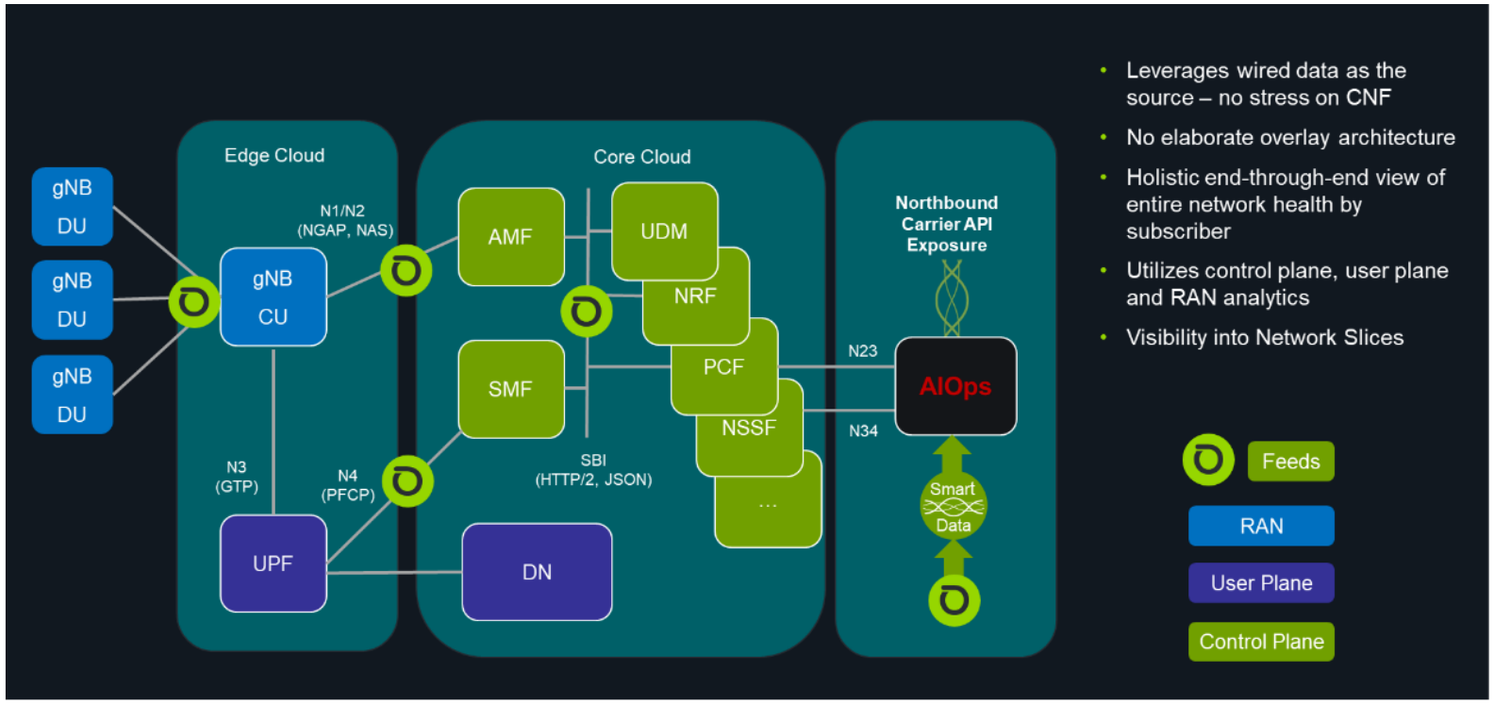

Could AIOps be the next NWDAF? *Source: NETSCOUT*

Monica: What are the use cases where AIOps shines?

Rick: I don't want to just throw AIOps out there if there are no business-relevant use cases. The key questions are: What are the business drivers? How can I secure my network? How can I make my customers happy? How can I improve my NPSs? How can I leverage AIOps for other revenue opportunities?

To answer these questions, we need to learn how to use AIOps to make the network more efficient, enable automation, generate new revenue pathways, and build use cases behind this. But we also need to be able to collect the data required to drive this process through the sensors that feed the AIOps pipeline.

Monica: How can operators estimate and measure the value they should realistically expect?

Rick: Operators may want to see whether their NPSs are improving, whether they can identify network anomalies faster, or whether their average mean-time-to-repair (MTTR) drops. These are all KPIs that VPs in the organization look at, and their bonuses often depend on them trending down. Also, we see new monetization use cases starting to generate revenue for the carriers.

Go back to the top

Monica: How much scope do you see for improving revenues?

Rick: Carriers always try to go after different adjacencies. Beyond reducing churn, raising NPSs, and lowering MTTR, they are also looking at new marketing use cases. They want to know how their subscribers are doing or how to market new services to them. They want to have open APIs and develop new partnerships. Carriers invested a lot into spectrum, core and slicing for 5G and want to see an ROI. They can start leveraging their 5G—and even their 4G—network data and develop partnerships outside their core competence areas to this end.

Monica: What made you decide at NETSCOUT to support AIOps?

Rick: At NETSCOUT, we take a strategic focus on being the AI data provider, and on collecting, carrying and streaming the data to the AIOps engine. We believe we have the industry's best quality and high-fidelity data, which is essential to the AIOps' success.

We don't just move data around in a garbage-in-garbage-out model; we curate it and feed it to the AIOPs engine. We deliver the right, real-time metrics but avoid a data bloat environment. We collect the data that the AIOps engine needs to run efficiently and make decisions in almost real-time. This cannot happen if the AIOps engine has to sift through terabytes of data and does not have telecom-level knowledge of what is relevant. This is why we came up with this approach: we want to provide a clean data stream based on what we heard about the need for visibility and to drive new initiatives from our customers.

It is not about NETSCOUT use cases. We work to support customer-driven use cases through collaborative partnerships with our customers. We work together to find out what level of data, what contextual information, what pre-processing, and which additional metrics they need to make AIOps more efficient and deliver the value they expect.

Monica: How do you see the progress in AIOps adoption over the next few years?

Rick: This is a new offering from us, and we have started having great conversations and see customers beginning to leverage this technology. We are hopeful that AIOps will become commonplace over the next few years. And we truly believe that it is all about the data, so this is where we need to start.

Watch the conversation

Go back to the top

About NETSCOUT

NETSCOUT SYSTEMS assures digital business services against disruptions in availability, performance, and security. Our market and technology leadership stems from combining our patented smart data technology with smart analytics. We provide real-time, pervasive visibility, and insights customers need to accelerate, and secure their digital transformation. Our approach transforms the way organizations plan, deliver, integrate, test, and deploy services and applications. Our mission is protecting the global leaders of industry from the risks of disruption, allowing them to solve their most challenging network performance and security problems, ensuring the connected world runs safely and smoothly.

About Rick Fulwiler

Rick leverages almost three decades of technology experience in the Telecom industry to his position in our CTO Office as a Chief Solutions Architect at NETSCOUT working closely with our global customers and strategic partners on 5G, Analytics, and Virtualization. Before moving to this role in 2016, Rick spent three years as Vice President of Sales- EMEA for Tektronix Communications in their United Kingdom office. Before his international role, he served as Tektronix Communications Director of Sales for the Americas since 2006. Before Tektronix Communications, Richard spent five years as the Director of North American Sales for Ixia and seven years as the Director of North American Sales for Tekelec. Over 10+ years, Rick has also held senior Product Marketing and Application Engineering leadership roles with a Telecom OSS/EMS startup now a part of JDSU.

Go back to the top

Will AI change how humans operate in telecoms?

A conversation with Azita Arvani, formerly CEO at Rakuten Symphony NA, and Paul Patras, CEO at Net AI and Associate Professor at the University of Edinburgh

Most of the time, we talk about how AI, ML and GenAI change how we run and operate networks, enable automation, tame complexity, and improve efficiency. They also change the way we – the humans in telecoms – operate, the skills we need, and how we deal with AI-based platforms. This will require profound changes on our end, and our ability to meet the cultural and technological challenges will be crucial to benefiting from AI.

In this conversation with Azita Arvani, formerly at Rakuten Symphony NA, and Paul Patras at Net AI and the University of Edinburgh, we did not focus on what AI can do for us but on what we can do to make it successful, how this is going to change all of us working in telecoms—from vendors to service providers, regulators, researchers and, well, analysts too.

Monica: We are all trying to come to terms and learn what AI can do for us in telecoms, so we often leave on the back burner questions on how humans' role will change and how we can direct the trajectory of AI adoption. For this reason, we have two guests who got into telecom AI from operators, vendors, academia, and startups, and they bring a wide view of the role of humans in an increasingly AI-driven environment.

Paul Patras is the CEO at startup Net AI and an Associate Professor at the University of Edinburgh. Azita Arvani has been the CEO at Rakuten Symphony NA and has worked with many vendors and network operators. Could you tell us how you got into AI in telecoms and what you see as your role?

Paul: I run Net AI, a spin-out from Edinburgh University, where I still have a faculty position. It is a network intelligence company specializing in software that provides real-time and predictive insights into mobile network usage and performance, which helps operators reduce their expenditure and generate new revenue streams.

Azita: I am the former CEO of Rakuten Symphony in North America, where we deployed an Open-RAN cloud-native network in which intelligent end-to-end automation improved functionality and reduced costs. I also serve on the board of public companies and had the privilege of working with both nimble startups and large Fortune 500 companies such as Microsoft, Nokia and Docomo. AI is not just an enabler of digital transformation but also a force multiplier.

Monica: For me, this is a compelling topic because I did not get into telecom through engineering but through neuroscience, cognitive science and philosophy, where we used AI to understand humans. In a way, talking about humans completes the circle.

In telecoms—or other industries, for that matter—AI models are just as good as the data you feed them or as the way you train them. What do you see the role of humans as we increasingly use AI to operate our networks?

Azita: The process starts with data collection. Before that, of course, you need to figure out what AI priorities you have for your company, and that has to be based on a holistic view from the top. Once you have done that and get to the data collection, humans play a large role in collecting, labeling, curating and possibly augmenting the data.

In anomaly detection, for example, you may want to add anomalies that don't usually show up in your data. Or you may want to select features that are relevant in your domain—they may be packet loss or signal strength at the network level; at the customer level, they may be the usage pattern or call duration. You may need to clean the data before you feed it to your model.

Watch the conversation

Go back to the top

The model selection is super important. There are many kinds of AI that we use in telecoms and they are a workhorse. Picking the right model and doing a lot of trials and experimentation with that is crucial to getting the cost versus performance tradeoffs right.

We all want great performance, but it comes at a cost. Training, fine-tuning and alignment give us the right optimization for the objective we want in the budget that we have.

Paul: To move one step further, beyond model selection, you need to identify the unique characteristics of a network deployment—be it mobile or broadband—and design from the ground up neural models that capture those characteristics and some special abstract features that are unique to your mobile networks. The model can then be tailored for whatever the task at hand is. For instance, the human element is important in deciding what loss function to use to train a model.

Model lifecycle is next. How do you continuously train and test the model? How do you update the model to reflect changes in your deployment or the distribution of your data? How do you go about maintaining the accuracy and performance of your existing model that has already been deployed? How do you extend the model to include other things?

Monica: Some think about AI as something you plug into your operations platforms from an off-the-shelf model. This is not to say that there are no good models in the market, but you need to customize them and ensure they meet your needs. Prior to that, you need to know what your objectives are and that you have reliable data to feed to the model. You may pick a good model, but if this is not the right model for your target, you will be disappointed.

How do we go about this process? What can go wrong? How can we navigate all this successfully?

Azita: I will return to the cost because cost is super important.

With GenAI, for instance, there is a cost for the number of tokens that come in as input, and there is a cost for the number of tokens that get generated as output. When we bring GenAI to solve a problem, we have to be cost conscious of this fact. ChatGPT-4 could be very expensive.

But there are ways to reduce costs. For instance, tightening up the prompt you put into GenAI can cut both the input and output costs.

There are many hybrid ways of working with the models to drastically reduce costs, such as using or combining different models. You can use cheaper models to reduce the noise in the input and create prompts that you serve to expensive models that use more parameters.

To optimize AI costs, you have to select an appropriately sized model (or combination of models) that meets the performance requirements of the task. Additionally, you can optimize the token size for further efficiency.

When customizing AI for specific domains, there are two primary methods: fine-tuning and Retrieval-Augmented Generation (RAG). RAG is often seen as an easy way to customize models by incorporating domain-specific files to enhance the responses of large language models (LLMs).

However, we should not forget about fine-tuning. Particularly in the telecom sector, extensive data can be used to fine-tune pre-trained models, making them more specific to telecom needs. After fine-tuning, RAG can incorporate real-time data, further improving the responses.

RAG increases the number of tokens, adding to the cost.

Monica: Cost is a top consideration in any assessment. But how do you assess cost tradeoffs? How do you do a cost-benefit analysis to decide what is the best model for you? For instance, how do you decide if building a model from scratch, as some operators plan to do, is cost effective than using a publicly available one that is not telecom specific?

And more generally, what range of capex and opex reduction—if any—can we expect from AI in the short and long term?

Paul: The choice may depend on the use case. There are many facets in each network.

To reduce energy consumption in the RAN, you may want to use AI to predict traffic and allocate resources. There is a saving opportunity but also a cost. You need to initially select and train a model—you want to avoid continuously training it. Do you need to have the right compute infrastructure to train that model? Or are you happy to train it in a public cloud and accept that you will have to move data from your deployment to a data center outside your network?

Operators will have different perspectives. Some will train their models in the central cloud, while others are skeptical. Some have regulatory

Go back to the top

constraints to meet. The public cloud is quicker and does not require the purchase of super-expensive hardware to train a model only for a limited time.

The size of the deployment matters, too. It is a numbers game. The bigger the deployment, the more energy you can save, but also the more models you need to train, and there are costs associated with that. If you have a small operation, the saving margins from optimization are low, especially if you need to use GPUs to train the model.

We need to be careful when assessing these factors because the tradeoffs are complex.

Azita: To put some perspective, opex is, on average, around 69% of operator revenues. That is fertile ground for AI to help reduce that, either through energy efficiency, as Paul mentioned, or by reducing the workforce or making the workforce more productive.

This can be done today. In the future, AI agents can become part of your workforce and generate even bigger cost savings.

Monica: This brings up the issue of attracting and retaining talent in telecom. Initially, AI and automation were seen as threats to the workforce because of the risk of major layoffs. But now we have a different problem. In telecoms, there is a shortage of new talent. We need AI and automation to keep our network going and possibly to make the job of people in the sector less boring and repetitive—and this, in turn, may motivate them to stay on.

In addition, the increased network complexity means we would need too many people to manage networks without automation. In this context, AI is a necessary tool to tame complexity.

Paul: Sometimes, we still need to explain that AI is not here to kill jobs. It almost feels like we are facing something like the industrial revolution when everybody thought that machines would do everything and that humans had nothing left to do.

If you look at anomaly detection or prediction, for instance, you no longer need someone monitoring the dashboard 24/7. Instead, you can have some degree of automation, whereby you use AI to identify suspicious events, filter them and only bring the critical ones to the attention of the attending human operator. And for the rest of the time, that person can work on more subtle tasks that require a deeper sector knowledge which, so far, an AI doesn't possess.

But there is a second side of the coin, too. Traditionally, there has been a disconnect between the telco and mobile networking community, and the pure ML/AI community. In the early days, many were skeptical about bringing AI into networking.

Things are now changing. Ultimately, it is a matter of having the right mindset and understanding. AI is another tool that can help us solve some problems, which we could solve with traditional tools that we already have, but we can now do so much faster and more efficiently while also improving employee satisfaction—which are important factors.

Monica: Most of the operations workforce may need new skills to use AI-based automation tools. How should we think about reskilling? Is that going to be a solution? Or do we need to hire new people? And how do we deal with these AI-proficient new hires' potential lack of network experience?

Azita: We faced a similar situation when we went to a cloud-native architecture for telcos. You want people that understand cloud and people who have domain knowledge, and you don't necessarily get both in the same person. You need to bring them together and get them to collaborate.

With cloud, people thought telco is just another industry, and that they would apply there what they already knew. And telco people thought they could learn cloud as you learn anything else—not a big deal. It turns out that you need experience in both areas.

In AI, we have a similar situation. You need data scientists, and ML and AI engineers that understand that field. And telecom is not just another industry. In a way, it is similar to the financial and healthcare industries in that it is very complex and highly regulated. Because of this, you need domain expertise.

The best approach is to get both types of experts to collaborate and develop an appreciation for what the others bring to the table. At the same time, telecom companies should invest in upskilling their workforce in AI and ML, preparing them to independently manage these technologies in the future.

Go back to the top

This collaborative learning environment will enhance skills across the board. Over time, telecom companies must become self-sufficient in network management, relying on a combination of in-house talent and external experts for initial setup, ongoing maintenance, and upgrades.

Monica: How do you get the two sets of experts—network and AI—to collaborate? You can imagine scenarios where the RF person has been doing things in a certain way for a long time, and the newly hired AI expert develops a model that comes up with entirely different recommended actions. Collaboration may be difficult in this environment.

Paul: Governments recognize there is a skill gap that is not telecom-specific and affects other fields, such as healthcare and finance.

So, you do start to see more interdisciplinary, postgraduate research programs with PhD students trained in the application of AI in a specific sector. This pipeline is still in the early days, but then there is a question: at the end of that journey, how do we make working in telco sexy again for these students as they get ready to enter the workforce or a research career?

Young generations now have, let's say, technical heroes who are outside the telecom domain. At the same time, networking has been increasingly perceived as a sort of utility. How do you convince people there is a lot of exciting work to do in telecom?

We see thousands of new applications being launched and a 20% year-on-year increase in mobile traffic. Yet managing the rising complexity in an efficient and profitable way is challenging.

The young generation has an important part to play here, but it isn't easy to get the message across.

Monica: Maybe we should be paying people more.

Paul: That's one factor.

Azita: There is more than pay. I view networking as AI's third pillar. We all agree: Data is the new oil. Compute is the currency of the future. We then need networking to bring the two together.

Look at Nvidia. One of the most productive acquisitions was Mellanox, which was all about networking.

In terms of prestige, networking should rank high, but the industry needs to do a better job of branding itself to make it sexier.

We also have edge clouds where data and compute resources can converge more flexibly than in a central cloud. This allows AI to operate in a private network at the edge, ensuring that sensitive data remains secure.

Additionally, this can be done at the far edge or near edge of a network. For the next generation of AI, this approach offers greater latitude to achieve desired latency, privacy, security, and other benefits.

Monica: I tried to make the argument that working in telecoms is cool and exciting to college students, but they did not seem too impressed. Do you think universities, other educational institutions, and governments are moving in the right direction to address the perception of what it is like to work in telecom? Is it something they should be doing?

Paul: In the UK, the government funds doctoral training in AI and machine learning, for instance, with a focus on how we build compute systems that run AI models efficiently.

Do we need huge compute power to train a massive LLM? Are there more cost-effective ways to do that? These are all questions that need to be answered, and the government is investing in addressing them.

Also, to stay ahead of the curve, you also need the right infrastructure. We have seen investment in high-performance compute clusters that are being rolled out now.

To succeed, we need multiple parties—government, private sector and universities—to come together. The doctoral programs I mentioned earlier are an example of that.

Monica: Is the US doing enough?

Azita: In the telecom industry, we have to take the responsibility to push this forward because we need this young blood to come in and bring innovation into the telco ecosystem.

This is why I have always been a big proponent of ensuring startups can easily work with the telco industry and telco ecosystem and encouraging young people to get into telecom.

Of course, we should improve the branding of this industry, as we discussed. Security within communication networks has become increasingly important for the US and its allies in recent years. Governments are now more involved

Go back to the top

in supporting the communications industry by providing more resources, such as grants for Open RAN, and raising the exposure of key players in the industry. Telecommunications is a dynamic field that integrates digital and analog, as well as virtual and physical networking elements, making it incredibly exciting.

Government efforts to enhance network security will help attract a higher caliber of talent to the telecom sector.

Monica: AI is also going to affect organizational culture. Cultural issues play a major role in the adoption of new technologies, but we often don't give them the attention they deserve.

It is going to take time for us to get used to the pervasive role of AI in managing our networks (or our lives, for that matter).

AI models are black boxes, and we don't have a good understanding of how they work. To use them, we need to give up direct control; and for that to happen, we need to gain trust in the models. And we need to be able to develop, train and test models that indeed do what we want them to do.

Back to culture, we need the organizational changes, the workforce capabilities and the company vision to support this transformation.

Do you think this is happening? What challenges do you see ahead?

Paul: Mindset is an important factor. There is a lot of excitement about deploying AI. We are ready to adopt it in a meaningful way and to commit resources to do so. We are ready to upgrade our existing infrastructure, hire new people, and upskill our people. External organizations may have more expertise in this area and are not there to put jobs at risk but to make employees' jobs easier.

Trust is another crucial factor. Most AI models are black boxes. Why should we trust them to do what we expect them to do? There is much ongoing work on the explainability of neural networks—to tell us what makes a model come to a certain decision. In classification, is it certain features of its input data or something else? Can we hold the model accountable for its behavior and audit it?

Networks are part of our critical infrastructure. The consequences can be nefarious if you have a fleet of autonomous cars that rely on connectivity and the network does not allocate sufficient capacity to support them. If this happens, you want to be able to go back and see why that happened and make sure that it doesn't happen in the future. Or you may want to run a digital twin system ahead of time and verify that you will not run into this situation by testing the model before you deploy it and to understand how it deals with these situations.

Azita: AI can be applied to various parts of the telecom business. If we just look at the networking part, it involves planning, building and operating the network. From personal experience, I have seen that the network architecture determines the organizational architecture. Traditionally, networks are made up of black boxes, which, on the network operations side, results in siloed organizations with experts on various boxes.

As we move to a more open, cloud-native architecture that integrates AI throughout the network, the operations organization must inherently become more open. Naturally, there is tension between established methods and the shift to a new, uncertain state. This transition can indeed be challenging.

I experienced this at Rakuten Symphony. For instance, when we introduced automation for network builds, end users were initially resistant to change. The success of this initiative was driven by a clear directive from the customer's top management to proceed with automation, regardless of staff preferences. There was no way to go back to the old ways; they had to move forward, which motivated them to embrace the shift.

An example with another customer for managing operational changes was to bring in a new team for a trial instead of using the existing organization. This new team had no preconceived notions about network operations. They had an open mind and were more accepting of what automation could do.

The combination of management commitment to motivating the teams and bringing in people outside the current organization to experiment can be powerful in getting the new AI system up and running and building trust.

Monica: This brings us back to the fact that no amount of technology and compute power will suffice to get AI off the ground. You need commitment from each part of the organization, starting at the top, as well as the right culture and structure in place.

Go back to the top

Azita: I see gradients of automation going all the way to autonomy, similar to self-driving cars. At lower levels, automation can happen with the traditional rule-based methods. As you move into autonomy, more advanced AI models will play a role.

Monica: Our networks are getting more open and disaggregated, and operators have to manage a network with more vendors. How can they ensure they get what they want from AI models? How do they avoid or identify hallucinations in their models? How do they define their targets? And how do they test reliability against those targets? Are digital twins useful in this context?

Azita: You start with a trial and proof of concept. You may use a pre-trained model that you want to fine-tune and incorporate reinforcement learning with a human in the loop. You would need to evaluate the information it generates and continually refine it until it meets your needs.

In addition to AI experts, domain experts, and committed executives, you also need specialists to establish guardrails to ensure the model generates the right responses and detects and prevents biases in the data or model. Various tests are essential to fine-tune the model for your specific environment.

You may also consider augmenting the data with synthetic data. For instance, you can generate synthetic data to simulate security breaches and check whether the network can handle them, should they occur. Since potential security breaches keep changing, you need to keep the model updated as well by repeating the exercise.

There is no question there is the notion of AI for security where AI will help us address security issues with things like anomaly detection to deal with fraud detection, for instance. But also, there is the notice of security for AI as AI also increases your attack surface, so you must protect the network accordingly.

Monica: Similarly, AI can test the network, but you can use AI to test AI models or to prevent undesired outcomes.

Paul: We have not reached a point where AI models learn abstract behaviors. They largely learn to reason about the distribution of the input data on which they have been trained. When we put them in production for inference purposes, they will look at the inputs, with the expectation that those inputs more or less fall within the same distribution of the data the models have been trained on.

Whenever you have variations and the inputs fall outside the distribution of the training data, problems may arise. These problems may be accidental, but they may also be deliberate.

In the case of network slicing, for instance, you may be training a model to distinguish different traffic types from aggregates to preserve privacy. Depending on the application, you mark some traffic types as suitable for a certain slice, and you monetize slices accordingly. Then, a malicious actor may try to manipulate the features of the traffic so that it gets mapped to a premium slice without having to pay for it.

That is a kind of adversarial machine learning. How do you prevent that from happening? There is a lot of ongoing research work on that. Some does not necessarily involve AI, but at least some pre-processing of the inputs to the AI model. To avoid adversarial manipulation, you need to sanitize the inputs.

You can also introduce noise to see how it impacts the model. This is similar to adding noise to the image of a cat to see at what point the image no longer looks like that of a cat to a machine. Or, for instance, if kids put stickers on a stop sign, an autonomous car may no longer recognize it as a stop sign and keep driving.

You can extrapolate this to networking, where people may manipulate the traffic they generate to be treated in a certain way or to mask a cyber-attack.

Monica: As we wrap up our conversation, let me ask you what you think is the biggest opportunity for humans to make AI succeed.

Azita: The biggest opportunity is looking at AI as a powerful tool already making an impact and will only become more influential. AI agents are an exciting new frontier, similar to acquiring an intelligent workforce at a low cost. Staying abreast of developments in this area is crucial to avoid missing out. Personally, I would have a significant FOMO around AI.

Paul: It is not a question of when this will happen. AI must succeed for operators to run their networks cost-efficiently and remain relevant. Hyperscalers have already started to move into the telco space. Incumbents may be threatened by more agile companies that are quicker to innovate. It is all about adoption, being open-minded, and being willing to try new paradigms to remain relevant.

Watch the conversation

Go back to the top

About Paul Patras

Paul is an Associate Professor/Reader in the School of Informatics at the University of Edinburgh and a co-founder and the CEO of Net AI, a university spin-out whose mission is to revolutionize network management via AI-driven analytics. He leads the Mobile Intelligence Lab and the Informatics Internet of Things research program, is a member of the Institute for Computing Systems Architecture (ICSA) and is affiliated with the Security & Privacy group. Before joining Edinburgh, he was a research fellow at the Hamilton Institute of the National University of Ireland, Maynooth, and a research assistant at IMDEA Networks (Madrid Institute for Advanced Studies in Networks). He held visiting research positions at the University of Brescia (2017, 2015), Northeastern University (2016), Technical University Darmstadt (2016), and Rice University (2010). He holds a Ph.D. and an M.Sc. in Telematics Engineering from University Carlos III of Madrid, and a Dipl.Eng. degree from the Technical University of Cluj-Napoca, Romania. His research seeks to bridge the gap between fundamental mathematical models and real-world applications of networked systems. He focuses on problems related to artificial intelligence in mobile networks, traffic analytics, performance optimization, security and privacy, prototyping and test beds.

About Azita Arvani

Azita Arvani is a dynamic and high-energy global technology executive renowned for her leadership in innovative transformations within the telecommunications and high-tech sectors. As the former CEO of Rakuten Symphony North America, she spearheaded initiatives integrating AI, cloud, and advanced connectivity, revolutionizing the communications landscape. Her leadership extends to roles as an independent director on public company boards. Azita's distinguished career includes groundbreaking work in both startups and Fortune 500 companies, where she excels in navigating uncharted business territories and opening new markets. She has driven growth through pioneering digital transformations and leveraged cutting-edge technologies such as AI, 5G, IoT, cloud, and AR/VR to create impactful, customer-centric solutions. Recognized for her leadership excellence, she was honored with the Woman in Tech Leader award in 2023 and was a nominee for the Woman in Telecoms at the World Communications Awards. Prior to Rakuten, Azita held strategic roles at Nokia, where she founded and led global innovation scouting teams and forged strategic partnerships that significantly contributed to Nokia's growth. Azita holds a Master of Science in Management from Stanford University (Sloan Fellow), a Master of Science in Computer Science from the University of Southern California, and a Bachelor of Science in Math/Computer Science from the University of California, Los Angeles (Magna Cum Laude). As an esteemed thought leader and speaker at prestigious international trade shows, Azita continues to champion true digital transformation across telcos and other industries, sharing her insights on automation and AI to drive the future of telecommunications.

Go back to the top

About Senza Fili

Senza Fili provides advisory support on wireless technologies and services. At Senza Fili, we have in-depth expertise in financial modeling, market forecasts and research, strategy, business plan support, and due diligence. Our client base is international and spans the entire value chain: clients include wireline, fixed wireless, and mobile operators, enterprises and other vertical players, vendors, system integrators, investors, regulators, and industry associations. We provide a bridge between technologies and services, helping our clients assess established and emerging technologies, use these technologies to support new or existing services, and build solid, profitable business models. Independent advice, a strong quantitative orientation, and an international perspective are the hallmarks of our work. For additional information, visit www.senzafili.com.

About Monica Paolini

Monica Paolini, PhD, founded Senza Fili in 2003. She is an expert in wireless technologies and has helped clients worldwide to understand technology and customer requirements, evaluate business plan opportunities, market their services and products, and estimate the market size and revenue opportunity of new and established wireless technologies. She frequently gives presentations at conferences, and she has written many reports and articles on wireless technologies and services. She has a PhD in cognitive science from the University of California, San Diego (US), an MBA from the University of Oxford (UK), and a BA/MA in philosophy from the University of Bologna (Italy). You can contact Monica at monica.paolini@senzafili.com.