STL PARTNERS

TURNING AUTONOMY INTO MARGIN: AGENTIC AI AND THE AUTONOMOUS RAN

Webinar: Questions and Answers

Questions and Responses:

This document outlines the questions and responses received from the STL Partners webinar, **Turning autonomy into margin: Agentic AI and the autonomous RAN** which was delivered on Wednesday 15th April 2026.

You can watch the recording of the session, and also access the slides, using the link here. In this document, we seek to address the questions raised in the webinar that we were unable to address in the time available.

If you have any questions not addressed in the webinar or this Q&A document, or want to hear more about our latest research or from our speaker and panellists, please contact:

• *Kirsten Rachman, STL Partners– kirsten.rachman@stlpartners.com*

• *Philip Laidler, STL Partners – philip.laidler@stlpartners.com*

• *Neil Coleman, Amdocs – neil.coleman@amdocs.com*

• *Darwin Janz, SaskTel – darwin.janz@sasktel.com*

• *Terje Jensen, Telenor – terje.jensen1@telenor.com*

Questions

1. **Our operator already has an external tool connected to OSS KPIs, mainly for optimisation, parameter consistency, integration scripts, and correction scripts. AI automation could perhaps also help with analysing large volumes of data, such as cell-level traces, and identifying device issues more quickly. What are the actual goals of automation in optimisation?**

   STL Partners: Automation in optimisation is most valuable when it moves beyond isolated, intra-domain tools and translates into cross-domain, intent-driven execution – systems that can understand a high-level business or customer objective and orchestrate the necessary actions across the RAN, transport, core and OSS layers. The next step for automation is to connect domain-level tools through integrated data and intelligence, so that cell-level trace analysis, for example, can be correlated with customer experience and service assurance data to drive faster, more targeted interventions. STL's analysis suggests that operators progressing through these stages can unlock meaningful opex and capex savings alongside new revenue opportunities, with the majority of financial value sitting in Day 2 operational excellence rather than in planning or deployment. But, don't overlook the full lifecycle and remember that for hardware it is a cycle.

   Amdocs: Automation is many things – we take a slightly narrower definition that network automation should result in closed-loop network actions. AI can help greatly reduce data volumes and complexity, but analysis of cell-level traces as a foundational capability to support closed-loop automation is a mature space with capable software already available, rather than something that necessarily requires ad-hoc AI.

   SaskTel: A greater focus on goals – grounded in genuine operational needs – should precede decisions about automation, analytics and AI. Networks are becoming too complex, and the volume of data too large, to manage through manual effort alone: operators cannot move fast enough or sustain the resource intensity that purely manual approaches require. Automation, applied with clear objectives, is what allows CSPs to address and manage optimisation at the scale the modern network demands.

   Telenor: Automation should be able to do the optimization tasks quicker and processing more relevant input data (as you point out). In that way, it can also be forward-looking and the optimization is expected to be done for the future. It may also be looking into a larger part of the network to explore even greater improvements. In some cases, optimization is triggered by certain events such as concerts and larger gatherings of people.

   To lift the optimization objectives to a higher level, customer experience and business value could be used as parameters. This would, for example imply considering aspects of different user applications and an end-to-end perspective like several network domains. For business perspectives, also revenue and other value estimates of different customers could be included in the optimization.

If brought to autonomous levels, it could also be running more frequently or even continuously for some cases.

2. **Is there a risk of losing knowledge as automation increases? Across the EU, many newer engineers already struggle to cope with the volume of information across 2G, 3G, 4G, and 5G. If we move further towards AI agents, is there a risk of losing critical know-how altogether, especially if engineers no longer know how to perform tasks when the tools are not working?**

   STL Partners: Knowledge loss is a genuine risk, and one that grows more acute as network complexity spans multiple generations and newer engineers carry less accumulated experience across 2G through 5G. STL's work on trusted intelligence argues that progressing to Level 4+ autonomous networks depends on capturing that knowledge in forms that AI and automation can use reliably – through structured data products, ontologies, network digital twins and repeatable workflows. Without that foundation, operators risk building teams that can operate tools but cannot validate or recover autonomous workflows when systems fail. The practical response is not to slow automation but to invest in parallel knowledge codification: upgrading legacy inventory data, adopting modern data architectures, breaking down organisational silos and developing skills for supervising increasingly autonomous systems. STL's research on multi-agent systems notes that mastering agentic approaches in telecoms may take the better part of a decade – which makes investment in knowledge assets now a prerequisite for the autonomy ambitions of the late 2020s.

   Amdocs: Agents are typically associated with full automation, but we think they will ultimately function more like co-workers to whom processes – such as L2/L3 diagnosis – can be delegated, making them a force multiplier rather than a replacement. Engineers should remain responsible for managing agent instructions and context, almost like maintaining their own partial digital twin of the network.

   SaskTel: This shift in knowledge and skills is already under way. Applying AI and agentic AI at scale will require new capabilities in building, procuring, implementing and managing intelligent tools – but network expertise remains essential for understanding how those tools should be directed. New software engineers joining today often outpace experienced network engineers in open-source and programming skills, but lack deep knowledge of complex network technologies. The most effective response is to operate as collaborative, cross-functional teams and apply deliberate resourcing strategies as the transition unfolds.

   Telenor: One way to look upon AI agents is to address the additional work load that often comes with new systems and technologies. As the number of parameters to be tuned and optimised in e.g. 5G systems has increased, it is by itself a challenge to be able to optimize the settings across these. Hence, in that way, the knowledge and capacity of the experts could rather be used to have an overview and sanity check of the overall set up, which then can be used to address areas related to business and customer perspectives as presented in the session.

3. **Do you have any case studies from Indian operators, with clear numbers showing savings or network improvement achieved through closed-loop workflows? If not, are there similar examples from international operators that you can share?**

   STL Partners: STL has published case studies on operators including China Mobile and Telefonica, which can be found on our website here: China Mobile and Telefonica. We would also add that there are relevant India-linked examples beyond the traditional MNO category. In our Lightstorm case study, we highlight an India-based NaaS player that has used automation as a core part of its proposition.

   Amdocs: Amdocs has relevant case study material with quantified results from closed-loop automation deployments and would be happy to share numbers directly with operators interested in exploring these use cases further.

   SaskTel: Demonstrating return on investment from closed-loop automation clearly remains a challenge across the industry. Establishing a credible means of quantifying and communicating that value is important for building executive support and organisational momentum. Where value can be demonstrated, a proportion of the savings – in time, effort or cost – should ideally be reinvested into further automation development, creating a regenerative cycle of improvement.

   Telenor: Not from India, but a neighbouring country, there are cases for automating the customer trouble ticketing to more than 95% of tickets. That is, the trouble tickets are automatically logged, categorized, forwarded for technical assessments and correlation, actions taken (if possible) and feedback given to the individual customers. This has reduced work loads on customer care and technical staff. However, key improvements is the customer experience throughout the process with fast resolution and personalized feedback. The overall process is truly closed-loop automation.

   Additional cases are within root cause prediction and proactive actions to avoid service degradation, energy savings, optimal battery replacements, network audit checks/testing, and more. These may mainly focus on specific processes, although giving similar benefits as reduced manual work, more precise actions taken and faster execution.

4. **When AI-RAN is discussed in the context of RAN sharing, does this mean applying the approach between a carrier network and a private enterprise network, or between two carriers?**

   STL Partners: AI-RAN sharing covers a broader set of infrastructure models than either carrier-to-carrier or carrier-to-enterprise alone. STL's research on AI-RAN and private networks identifies several distinct configurations: shared RAN only, shared RAN and control plane, shared RAN with partial control plane separation, and slice-based models – spanning both inter-carrier and public-private arrangements. In the AI-RAN context, one of the more commercially compelling framings is a shared compute and orchestration platform at the RAN edge that hosts both network workloads and enterprise AI inference tasks, effectively turning

distributed RAN infrastructure into a revenue-generating edge platform. That model is inherently a carrier-to-enterprise or hybrid public-private construct.

SaskTel: RAN sharing between carriers was referenced in the context of Multiple Operator Core Network (MOCN) arrangements, which already exist as a form of RAN-as-a-service between partner wireless carriers. The comment on private wireless networks was intended to highlight the potential for operators to consider RAN sharing via network slicing as an alternative to full private wireless networks for customers who do not need or want to manage a complete private network deployment.

Telenor: RAN sharing typically refers to several operators/carriers using the same RAN.

In some of the private network set ups, there can be a requirement for multiple operators/carriers to offer services (similar to a neutral host set up).

How AI-RAN is to be used for such set ups, in particular if working on optimization of radio resources that could affect operators'/carriers' services differently, may still need to be worked out. However, running AI on similar platform as for RAN, and working on technical radio optimization, could still apply.

5. **Are operators exploring rApps and AI mainly on top of traditional RAN solutions, such as a RIC-like system, or are they building on top of O-RAN architectures?**

STL Partners: Operators are pursuing both paths, and the near-term picture is hybrid. STL's tracking of telco cloud and RAN deployments shows that vRAN – which uses O-RAN-compliant open interfaces but is typically deployed as a pre-integrated solution – is currently growing faster than full open RAN disaggregation. This reflects a pragmatic middle ground: operators gain openness and improved data portability without taking on the integration complexity of fully multi-vendor open RAN at scale. For AI and rApps specifically, O-RAN's value lies in standardised interfaces and the RIC framework, which create better conditions for third-party intelligence to act on the RAN. However, STL's research is clear that operators expect AI to deliver value in traditional, vendor-integrated RAN estates too – AI-RAN should not be treated as contingent on full O-RAN rollout. In practice, AI-driven optimisation will be applied on top of existing vendor ecosystems through proprietary or semi-open interfaces, with full open RAN reserved for greenfield deployments or specific use cases where disaggregation is commercially justified.

Amdocs: The reference architecture appears to be settling on an O-RAN SMO/RIC layer sitting on top of traditional radio access networks. The operators we work with value the modularity that rApps provide but are not, in most cases, pursuing full open RAN disaggregation.

SaskTel: SaskTel is exploring rApps and AI on top of traditional and vRAN architectures, as this is more immediately relevant to our network environment.

Telenor: Some use cases may only address Open RAN equipment (or RAN domains), while several use cases would address several technologies and work with traditional RAN solutions.

For example, energy saving cases – as deployed by several operators – works across 5G, 4G as well as 2G/3G. This could then be 'rApps-like' implementations run over a controller or management systems that controls the power of different RAN units.

6. **Many of these RAN autonomy proposals sound compelling in principle, but how do you actually translate them into a deployable architecture within an operator's existing environment?**

   STL Partners: There is no clean architectural path from today's fragmented OSS/BSS and RAN management estates to a fully autonomous network, and STL's recommendation is to build incrementally rather than wait for a complete reset. The first practical step is to expose existing systems through standardised APIs and integration layers, so that data from disparate domains can be brought together without a full platform replacement. The second is to combine digital twin capabilities with live network telemetry: digital twins allow automation logic to be tested safely before actuation, building the trust needed to progressively reduce human oversight. The third is to layer intelligence from centralised cross-domain visibility, through distributed near-real-time decisions, to hierarchical agent coordination across the stack. Data coordination is consistently the most underestimated enabler – answering a question like "is this customer issue caused by a RAN fault or a core problem?" requires data from multiple siloed systems to be accessible in near-real time. The architectural ambition is lifecycle-wide coordination; the deployable starting point is integrated data and progressive trust.

   Amdocs: Deployment is about layering. SMO/RIC is the logical first layer to put in place, sitting on top of existing data infrastructure to provide a standardised scripting and automation programming environment. Agents should then sit above rApps and act as the policy and process layer – for example, a coverage agent that provides business intents to, and monitors, rApps for coverage and capacity optimisation.

   SaskTel: Beyond established monetisation models such as RAN-as-a-service (MOCN) and private wireless, speculative services like network-as-a-sensor appear technically feasible but currently lack the solution definition needed to deploy them with confidence.

   Telenor: Autonomous RAN can be understood in several ways such as i. implementing capabilities for making RAN becoming aware, automatic controlled and adaptive (autonomous network), and ii. supporting independently working RAN with mobile core. On the former it would typically require a controller such as RIC or management system that takes care of the automated tasks to re-configure RAN components. The latter can for example be a private network able to run independently of any external connectivity.

7. **While L4 is the ambition, how much will the journey towards it be shaped, in both substance and timeline, by 3GPP, TM Forum, and other industry bodies?**

   STL Partners: TM Forum's Autonomous Network Framework has given the industry its most operationally useful shared language – the Level 1 to 5 maturity model provides a common

reference for where operators are, where they want to be and what capabilities they need to develop. STL's research uses it as a consistent lens for operator benchmarking and case studies. 3GPP matters for the underlying RAN standards that define what can be exposed and controlled programmatically, and its evolution towards AI-native air interfaces from Release 18 onwards will shape the long-term technical architecture. However, STL's analysis consistently emphasises that the harder challenge is execution rather than compliance: operators genuinely advancing towards Level 4 are doing so by linking autonomy to commercial outcomes, extending it beyond operations into planning and deployment, and building data and governance foundations that frameworks describe but cannot deliver. The pace of progress will ultimately be determined by operator investment decisions and organisational capability, not by industry body release cadences.

Amdocs: Some operators will move ahead of the industry bodies, driven by necessity, and explore the limits of what can be automated. Many others will follow body frameworks, which offer the best opportunity to build consistent, reusable automation products. Regardless of approach, operators will need to fill the gap between the "what" defined by TM Forum and the "how" specified by the O-RAN Alliance, 3GPP and other bodies.

SaskTel: L4 autonomy is already being heavily shaped by industry bodies, and SaskTel expects that achieving it will come through a combination of vendor and third-party solutions alongside operator-led development, all aligned with industry standards. Using standards consistently is essential to ensuring that automation products remain interoperable and maintainable over time.

Telenor: Several of the capabilities/functions will be standardized by 3GPP for RAN, while also using descriptions related to cloud platform (e.g. CNCF). In addition there will be references and best practise sharing such as facilitated by TMforum.

Then, these need to be architected, designed, implemented and operated where also readiness of the organization and partners will influence the time lines and steps.

8. **Is agentic AI-based automation the only viable route to autonomous RAN, especially when RAN vendors are also bringing forward their own isolated autonomous solutions?**

STL Partners: Agentic AI is not the only route to autonomous RAN, but it is becoming the most important enabler of the cross-domain coordination that Level 4 genuinely requires. Deterministic, rules-based automation remains highly effective for well-bounded, repeatable tasks – restarting cells, provisioning network functions, applying known parameter corrections – and will continue to underpin the majority of automated actions. What agentic AI adds is the ability to handle processes requiring multi-step reasoning, cross-domain data synthesis and adaptive decision-making. STL's work on multi-agent systems identifies this technology set – combining agents, knowledge planes, reasoning engines and data mesh – as the most likely bridge between Level 3 and Level 4. The concern about vendor-isolated autonomous solutions is well-founded: self-contained vendor autonomy reduces operators' ability to coordinate across domains and build the cross-domain intelligence Level 4 requires. The strategic

response is to insist on open, standardised interfaces, so that agentic orchestration can sit above and across vendor-specific automation rather than being confined within it.

Amdocs: Agentic AI needs to sit on top of vendor autonomous solutions rather than compete with them. We see agents as the connecting process layer that joins pockets of vendor-specific automation together into a coherent, cross-domain whole.

SaskTel: Agentic AI-based automation is not the only viable route to autonomous networks, but the modular, intelligent agent model does present compelling opportunities. Delivering on those opportunities requires applications capable of executing intelligently within a closed-loop model, underpinned by clearly defined policies, requirements and controls.

Telenor: In principle, there are several ways to implement autonomous RAN, where e.g. competences/capabilities in the organization, support from suppliers/partners and company ambitions regarding needed business agility should be considered when detailing out potential roadmaps.

9. **Building L4 or L5 on top of legacy OSS/BSS systems can seem almost impossible. A unified operations and business system may need to be built from the ground up, with AI-native foundations. Do you agree that this level of transformation is a prerequisite for true autonomy?**

STL Partners: Legacy OSS/BSS can be a constraint but a full rebuild is not a prerequisite for meaningful progress (it is also unrealistic for most operators). The operators advancing most clearly today are not rebuilding an AI-native clean slate: they are either greenfield operators (who may have a relatively clean slate already) or more established operators who are wrapping existing systems with standardised APIs and integration layers, then advancing use case by use case and building trust in automation progressively. STL's recommendation is to avoid frozen transformation initiatives – large, multi-year OSS replacement programmes that defer all autonomy investment until a future state is reached. The alternative is a capability-by-capability approach: identify the highest-value automation opportunities, connect the data those use cases require, and deploy intelligence on top of the existing estate while the longer-term architectural journey continues. AI-native is the right direction of travel, but scalable autonomy can begin well before the overhaul is complete.

Amdocs: Building from the ground up can work in some instances, but it involves replacing automation that already works – including logic that encodes dozens of hard-won edge cases – and making technology decisions that will themselves become legacy within a few years. It is better to focus cutting-edge AI on what cannot already be done, and work backwards with targeted refactoring where needed.

SaskTel: Building from the ground up is always preferable when major transformation is required, but at the scale of large network deployments – given investment levels, customer impacts and operational continuity requirements – some degree of bridging between technology generations will almost always be necessary. In those circumstances, standard

protocols, open interfaces and well-designed adapters are what make the transition manageable.

Telenor: This would depend on the actual capabilities supported by the current OSS/BSS, for example regarding APIs, real-time, and so forth. However, several of the older OSS/BSS component would not be able to work for a full-blown L5 operation.

10. **A great point was raised on outcome-based agentic technology strategy. Telcos are generally strong at managing networks and identifying patterns, but capabilities such as data products, data models, and data governance are newer and often weaker. How big a challenge is this for operators trying to adopt agentic AI effectively?**

    STL Partners: Data products, data models and governance are not supporting issues for agentic AI but the central dependency. Trusted intelligence, which STL identifies as the prerequisite for Level 4+ autonomy, requires data that is cross-domain, high quality, contextualised, well governed and reliably accessible. Most operators today remain at an early stage of data maturity: inventory data is often incomplete or siloed, telemetry is collected but not correlated across domains, and the organisational structures needed to treat data as a reusable asset rather than an operational byproduct are not yet in place. Practical success factors include an empowered central data team, coordinated governance that gives domain teams ownership without creating new silos, and a cultural shift from treating data as a product to treating it as a strategic asset. For agentic AI specifically, agents that reason and act across domains will amplify both the value of good data and the damage caused by poor data.

    Amdocs: Agents are surprisingly effective at operating on specific, pertinent data subsets, which means it is not always necessary to solve the data problem holistically before making progress. That said, we do see the challenge within operators of consistently collecting and storing data, which is sometimes treated as the end goal rather than the enabler. Keeping data close to the systems that generate it and providing direct hooks to those systems tends to produce more robust architectures. In the RAN context, rApps provide well-defined pockets of data and automation that are easier to maintain and provide sufficient context for agents to act on.

    SaskTel: This is a genuine challenge. Applied to a narrow, well-defined use case it is manageable, but achieving consistent data and analytics maturity across domains and across the organisation is significantly harder. The complexity scales quickly as the scope of agentic AI broadens.

    Telenor: Effective and scalable implementation of AI agents would require a solid management of the relevant data domains. For several of the use cases the amount and scope of data is limited. Hence, it would be possible to take several steps before having all data product and management details fully in place.

11. **We should bear in mind that private network roadmaps are often derived from public network roadmaps for the top-tier vendors. Given that, is it realistic to expect vendors to prioritise AI for private networks before they do so for their public network portfolios?**

    STL Partners: Public network portfolios will lead vendor AI roadmaps, for the straightforward reason that scale, installed base and operational complexity create the strongest incentive and the largest data sets for AI development. That pattern holds consistently across network automation, AI in the RAN and agentic use cases. However, the relationship between public and private network AI development is not simply sequential. STL's private network research identifies use cases – particularly those requiring low latency, data locality, predictable SLAs and privacy guarantees – where enterprise demand is strong and distinct from public network requirements. In those environments, private and hybrid networks are more likely to act as targeted proving grounds for specific AI capabilities.

    SaskTel: Vendors should prioritise AI for carrier network solutions, given that public networks represent by far the greater need in terms of volume, scale and investment. For that investment to have lasting value, AI capabilities need to be designed for reuse and structured to generate opportunities across multiple use cases rather than being built for a single deployment context.

    Telenor: Prioritization – when needed – would typically be based on expected value for the different parties. On top of that several vendors support similar features across their different product lines (although this varies across the vendors). Hence, timing of this may be different for the different vendors and their markets.

STL PARTNERS



Research



Consulting



Events

www.stlpartners.com