Whitepaper

June 2026

Monetizing AI-RAN

Delivering 6G commercial potential on 5G networks

Author: Robert Curran, Appledore Research

www.appledoreresearch.com

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

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

Publish date: 12-Jun-26

Cover image: Robert Curran

INTRODUCTION

The unstoppable momentum of AI is driving the pace of change in telecom networks faster than the traditional G-cycle: additional AI compute power at cell sites (in the form of AI-RAN) now looks more likely to be the near-term upgrade path for today's millions of 4G and 5G cell sites. Early AI-RAN evangelists include T-Mobile's Chief Network Officer Ankur Kapoor ("every cell tower should have a brain"), Softbank's Hideyuki Tsukuda ("we will build infrastructure with AI-RAN") and Indosat CTO Desmond Chung ("[Our] AI-RAN Research Center shifts our network strategy from being connectivity-led to being AI-native").

But while much of the immediate focus is on using more AI in the RAN to create better-performing networks, operators looking for growth from additional network investment are asking a critical question: how AI-RAN will unlock new commercial opportunities?

One compelling answer emerges from a new approach to an increasingly challenging problem: how to optimize network performance and spectral efficiency in real-time. The method of achieving that turns out to offer an accelerated route to the sort of "network-as-sensor" capability that has immediate commercial value for operators around the world today.

This paper explores how **Cohere Technologies** has opened the door to a new vista of network data — generated as a by-product of work the operator already pays for — that operators can use both to gain far richer insight and observability of their networks, and to monetize across a range of high-value commercial applications including defense, public safety and more. Along the way it draws on examples from telecom and other industries where data that was initially discarded, ignored, or treated as mere exhaust turned out, when considered from a different perspective, to be of huge economic value in its own right.

In pursuit of better network performance, mobile operators are working to capture more data from their networks — but discovering that the instrumentation they have is nowhere near fast enough. Fifteen-minute polling cycles cannot detect and resolve issues at the speed today's always-connected customers expect. Operators are left with limited, delayed visibility into how their networks are actually performing, which in turn limits their ability to make changes in real-time. However, operators face a trade-off between gathering vastly greater amounts of information and being able to process and respond to it in a timely fashion.

But vastly greater amounts of *the same* data is no longer the only option. And better network performance is not the only prize on offer: the same data that is created to drive real-time optimization turns out to have much greater potential value for operators – addressing a key question about how to derive meaningful net-new revenue "beyond connectivity" from the billions spent on building mobile networks.

Richer, faster network telemetry is not only the key to better performance and observability; in the right form, it can also be the foundation for realizing Integrated Sensing and Communications (ISAC) — a capability the industry has, until now, considered an inevitably expensive network upgrade reserved for the 6G era, and logjammed in standards development. The same data that lets an operator see its own network more clearly can also allow it to perceive the physical environment that network sits inside, and that perception has commercial value well beyond connectivity.

THE OBSERVABILITY PARADOX

Modern mobile networks are extraordinarily complex, and yet the tools operators use to watch them remain coarse and slow. Conventional RAN observability typically polls performance counters at fifteen-minute intervals. For a network carrying latency-sensitive traffic — video calls, cloud gaming, industrial automation, real-time collaboration — a fifteen-minute snapshot is an eternity. A congestion event, an interference source, or a degraded handover can appear and disappear several times over within a single polling window, leaving the operator with an averaged, smeared picture that hides exactly the transient behaviour engineers most need to see.

Exactly as service availability becomes more critical to customers, operators' view of network and service conditions falls behind. Operators make capacity, scheduling and optimisation decisions on the basis of data that is both delayed and aggregated. They cannot easily distinguish a momentary spike from a sustained problem, cannot see the geometry of how devices and reflections are moving through a cell, and cannot feed high-resolution signals into the AI systems they increasingly rely on to run the network. The data that would resolve these questions either is not collected at all, or is collected and immediately averaged away.

This is the paradox at the heart of network operations: the more granular the data an operator tries to gather, the harder it becomes to process and act on it within an actionable timeframe. Historically, operators have resolved that tension by collecting less — sampling slowly, retaining little, and accepting the blind spots – or by building slack in the form of redundant capacity. The alternative, collecting a continuous high-fidelity stream and processing it at the edge, has been out of reach because the underlying systems simply did not expose the data required.

Understanding Congestion

The base station scheduler is the component that decides, every transmission interval, which users get which slices of spectrum; and it can only make good decisions if it knows the current state of each user's radio channel. That knowledge comes from channel estimation, which is inherently limited: pilot or reference signals consume spectrum that could otherwise carry data, estimates are noisy, and the information ages almost immediately as users move or the environment shifts.

When the channel state is stale or uncertain, the scheduler is forced to hedge: choosing conservative modulation and coding, leaving guard margins, and relying on retransmissions when its guesses are wrong. Each of those hedges wastes resource blocks, so the network delivers far less than its theoretical capacity. Under load, imperfect channel knowledge prevents the scheduler from packing users tightly or exploiting favourable channel conditions, so cells saturate well before their spectral limit – operators see dropped throughput, latency spikes, and contention even though radio resources are actually available.

In other words, much of what looks like congestion is actually the cost of the scheduler operating blind—a limit set not by spectrum itself but by how accurately and how far ahead the channel can be known. Alternative methods for estimating the channel (such as delay-Doppler) can be more precise, reducing overhead and effectively recovering the otherwise lost spectral capacity.

The emergence of Open RAN introduced some new structures and standards to allow richer, faster network data, but that comes with much more strategic implications – with many operators not yet ready to make that particular leap. In any case, simply more of *the same* data – more link quality KPIs – overlooks the potential of a *different* form of real-time network data.

A NEW PERSPECTIVE ON THE NETWORK

Cohere's Universal Spectrum Multiplier (USM) is software that runs alongside, but independent of, the gNB scheduler in 4G and 5G base stations. Its primary use is to increase downlink capacity on congested cells. To achieve this, it builds a real-time, high-resolution model of the wireless channel — including how every connected device is moving, where multipath reflections originate, and how the radio environment is changing on a sub-second basis. That alternative model (the **delay-Doppler** representation) allows USM to assign calls to channels in a more intelligent, reliable and efficient way, squeezing extra capacity from the available spectrum without the need for extra network hardware or complex network re-engineering.

The delay-Doppler representation is a geometric coordinate system that resolves, for every signal received at the antenna, the time-of-arrival (delay, which maps to distance) and the rate-of-change of that delay (Doppler, which maps to relative velocity). The essential engine of the scheduler is thus not simply improved but completely replaced.

The raw data used to build the delay-Doppler model is huge: logging over 200 distinct attributes every second, generating over 400 million rows of data per sector per day (around 6GB every 24 hours). Compared to the fifteen-minute polling cycle of conventional observability, this is the difference between a periodic health-check and a continuous heartbeat monitor for the network.

This is not a repackaging of telemetry that exists today. It constitutes an entirely new data set for operators. Where conventional OFDM-based RAN observability captures throughput, latency, error rates and beam-quality metrics — operational signals about how well the link is performing – the delay-Doppler model captures the *geometry of the radio environment*: a continuously updating map of what is moving, where, derived as a by-product of computing the channel model needed to schedule traffic intelligently. Cohere calls this feature "ECHO": Enhanced Channel Insight with Holographic Observability.

In trials with **Bell**, network managers encountering this data for the first time reported being "flabbergasted" by the level of insight it afforded, having previously lacked any comparable visibility into the radio environment — a useful indicator that the data really is novel rather than a re-presentation of what RAN engineers already see.

With this data, for every connected user on a cell site, it is possible to resolve signal latency, motion vector and amplitude. However, for *passive* objects — vehicles, drones, people, anything that scatters or reflects RF energy — it can resolve range and rate of approach, effectively turning a mobile network into a network of environmental sensors. The cost-saving value of addressing a congestion issue thus turns out to offer a upside revenue potential in the form of a new and monetizable capability, increasing the ROI from congested-cell capacity multiplication, MU-MIMO beam-forming optimisation, AI-driven scheduling, and protection of legacy LTE 4T4R investment.

ISAC: THE 6G PROMISE AND THE 6G PROBLEM

Integrated Sensing and Communications (ISAC) is an old idea that has gained a new relevance. ISAC (or JSAC – Joint Sensing and Communications) treats the radio waves a network already transmits as a sensing medium. The network analyses the reflections returned from the physical environment to infer the position, speed, and behaviour of objects — including objects that are not connected to the network, but happen to be in or moving through the space the RF signals occupy. This expands the operator's role from carrying data between endpoints to perceiving the environment itself: a cell site becomes both a base station and a passive radar.

ISAC is formally recognized as one of six key usage scenarios for 6G. 3GPP initiated an ISAC study item in Release 19 in December 2023 and finalised the channel-modelling specification in May 2025; over thirty candidate use cases were captured in TR 22.837. As of early 2026, the Release 19 ISAC study is approximately 58% complete.

The major infrastructure vendors are all engaged. **Ericsson** demonstrated a live ISAC drone-detection proof-of-concept at its Texas headquarters in February 2026, positioned by the US Department of Defense's FutureG Office as a strategic priority. **Nokia**, **Huawei**, **Samsung** and **Qualcomm** have each published architectures or trial results. A small group of software-focused specialists — notably **Tiami Networks**, partnered with Vodafone in Malaga (February 2026), **e&** in the UAE (GITEX Global 2025), and defence-focused **SEMPRE** for passive drone detection near US military installations — is attempting to deliver ISAC-like capability without waiting for native 6G.

The 6G commercial horizon nevertheless carries non-trivial economics. A fully capable native ISAC radio is likely to require expensive hardware upgrades per cell site, alongside self-interference cancellation hardware to support monostatic sensing (the same node transmitting and listening for echoes) in the same band as communications (today's 5G base stations operate in half-duplex mode and cannot easily transmit data and receive weak sensing echoes at the same time). An operator with 50,000 cell sites will spend billions in upgrades before the first sensing-revenue dollar arrives. Compounding the issue is that the 6G radio standard itself is not expected to be finalised until late 2029, and waveform co-design — making a single signal acceptably good for both throughput and sensing — remains an open research problem.

Alongside the technical and standards risks sits an unresolved regulatory environment. ISAC waveforms do not fit conventional spectral allocations; spectrum sharing between civilian sensing signals and federal or military radar systems is technically feasible but politically sensitive. Privacy frameworks for ubiquitous environmental sensing have not been written. ETSI's December 2025 work item on security and privacy is a recognition that these are unsolved problems, not a solution to them.

6G-based ISAC looks like a feature whose marketing has run somewhat ahead of its commercial readiness. There is high vendor and policy enthusiasm; clear strategic relevance for nation states; a maturing standards process; and a yawning gap between the technology's promise and the conditions for at-scale deployment. Against that backdrop, anything that brings ISAC-grade sensing forward — without the hardware bill, on existing infrastructure — deserves close attention.

SENSING THE FUTURE

Despite the timeframe of ISAC-related 6G standards, there remain near-term opportunities for telcos to leverage (in monetizable terms) the sort of sensing capability that is now possible:

- **Defence and national security** represents the most immediately well-funded application. The US Department of Defense's FutureG Office identified ISAC-enabled drone detection as a strategic research priority in late 2023. The fiscal context is unusually large. The US Congress appropriated $24.4bn for the Golden Dome air-and-missile-defence initiative in the 2025 budget, with $43.5bn requested for 2026; Congressional Budget Office estimates run $161bn–$542bn over twenty years for core elements, up to $831bn for full implementation. (For comparison, the prior US administration allocated $1.5bn for Open RAN through the CHIPS and Science Act, of which approximately $850m was subsequently rescinded — a useful reminder that sensing-related defence funding now dwarfs the policy-driven funding for the network-modernization programmes most operators are familiar with.)
- **Public safety and pedestrian protection.** The same primitives — detection of moving objects at known velocities and ranges — apply to vulnerable-road-user protection. Cell sites positioned along urban corridors can detect vehicles approaching pedestrian crossings at unsafe speeds, generating alerts to municipal traffic systems, connected vehicles, or directly to V2X-equipped phones. The application sits at the intersection of operator infrastructure and municipal smart-city programmes, where procurement frameworks, and the willingness to spend, already exist.
- **Airspace surveillance around sensitive locations.** Airports, government buildings, prisons and major events all face the small-UAS threat: drones used for nuisance, surveillance, or worse. Conventional counter-UAS systems are expensive, single-purpose, and typically reactive. A passive, network-resident sensing layer that already covers the airspace around an existing cell site offers a different cost and coverage profile. This is the application most amenable to early commercial monetisation, because the customer — airport authority, government facility, event organiser — has clear willingness-to-pay and a well-understood procurement process.
- **Smart-city, venue analytics and autonomous systems.** Crowd-flow analytics, occupancy detection and intrusion monitoring without cameras — privacy-friendly relative to camera-based alternatives, since no facial-recognition primitives are exposed — are useful for venue operators, transport hubs, and high-footfall retail. The same real-time environmental awareness can augment the onboard sensors of autonomous vehicles, extending perception beyond line-of-sight, and can feed the continuously updated digital twins that operators and cities increasingly want.
- **Sensing-as-a-Service for enterprise.** Operators expose sensing data via APIs to logistics firms, property developers, insurers and other enterprises. This is the model that mirrors historical analogues: recurring B2B revenue, anonymized and aggregated data products, and pricing tied to the value of the use case rather than to the volume of the underlying telemetry.

Not all of these are equally near-term. Defense and airspace surveillance are credible 2026–2027 commercial opportunities given existing procurement and policy momentum; smart-city and pedestrian-protection programmes typically require longer municipal contracting cycles; pure data-product businesses depend on market-development work that operators will need to invest in alongside trials. The addressable application set is broad enough to support staged commercialization rather than a one-time market opportunity.

To be clear, USM and ECHO are not a substitute for native, monostatic 6G ISAC at every performance metric — they do not, for example, deliver the long-range air-domain awareness of a dedicated phased-array sensor at the same energy budget. What they do provide is sensing-grade observability of the cell's neighborhood, in real time, from the network footprint that mobile operators already control. For a wide range of practical applications — close-in drone detection, vehicle and pedestrian tracking, perimeter monitoring around sensitive infrastructure, and the continuous real-world feeds that high-fidelity digital twins require — that is sufficient. Driving value from something other than connectivity services is a perennial concern for modern operators.

SAME DATA, NEW BUSINESS

Telcos have struggled to find significant revenue streams that are not ultimately and directly connected to the transportation of bits. But in the age of AI, the availability of a vast new and distinct dataset represents a genuinely new opportunity. Historical examples - from search, payments, agriculture, logistics, finance in which a regulated or infrastructure-heavy operator discovered that the data it produced as a side-effect of its primary activity was structurally more valuable than that primary activity itself– illustrate the potential.

Google search becomes advertising

Google's stated mission was to organize the world's information; search-query logs were generated as exhaust. The reframe — that the queries themselves were a high-fidelity intent signal monetizable to advertisers — turned a free consumer service into a near-trillion-dollar advertising business. The data did not exist before the primary activity did its job well; the better Google indexed and ranked, the richer the query log it implicitly accumulated. Google didn’t set out to build an advertising data business — it set out to help people find information quickly, and the record of what they were looking for turned out to be the asset.

Satellite imagery as alternative data for hedge funds

Until around 2011, commercial satellite imagery was sold mainly to defense and mapping customers. But a couple of geospatial analytics companies realized that the same imagery, processed through computer vision, could count cars in retail parking lots, oil in floating-roof tanks, and container throughput at ports — “alternative-data” products (i.e. not directly financial data) that hedge fund managers could use to make more richly-informed decisions. Subsequent research found that traders using this data earned better returns in the three days around earnings announcements. The alternative-data market is now estimated by Grand View Research at approximately $11.7bn (2024).

Legacy networks: unlocking the commercial potential of AI-RAN

5G has, for the most part, not delivered significant additional revenue for telcos, and while work on 6G standardisation progresses, few expect that its eventual arrival will result in a meaningful uptick in revenue. In the meantime, however, there is significant interest in (and backing for) leveraging more AI compute in the existing RAN, initially to boost performance and efficiency, but with a view to creating new monetizable services.

Value in AI comes through the combination of large amounts of data with AI processing that delivers actionable insights: precisely the combination that Cohere's geometric model in USM, streaming the ECHO dataset, represents.

Rather than being a depreciating network asset, each AI-RAN-enabled legacy base station can become an environmental sensor – with monetizable service potential for operators and their customers in local government, retail, logistics, public safety. SON-like control could be applied to manage how cells co-ordinate in using this capability.

Cohere's ECHO innovation represents an on-ramp to an AI-RAN-based upgrade path that could deliver incremental value long before and right though an envisaged future 6G refresh cycle.

RISKS, CONSTRAINTS AND OPEN QUESTIONS

Ubiquitous environmental sensing of this kind is unprecedented at this scale, regardless of which technology delivers it. Public, regulatory, and civil-society engagement with the ethics of network-resident sensing remains immature. Operators will need to invest visibly in transparency, governance frameworks and consent mechanisms — not because the technology is inherently surveillance-grade (it is not) but because public acceptance will be a more important constraint on growth than any technical bottleneck.

As with native ISAC, USM-derived sensing operates in the same spectrum as the operator's communications service. Coexistence with federal and military radar in adjacent bands is a real consideration in defence applications. Regulators will need to develop frameworks specific to multifunctional radio uses; today's frameworks are designed for service-specific allocations and provide limited scaffolding for the kind of capability USM enables.

There is, as yet, no established buyer ecosystem for network-sourced sensing data. Pricing benchmarks, demand profiles, and competitive dynamics with adjacent providers — camera systems, dedicated radars, GPS-based services — have not been established. Operators considering network-derived sensing services should invest in market-development work alongside any technical pilot, testing customer willingness-to-pay, procurement-cycle length, data-product format requirements, and the right channel-partner economics.

Finally, the eventual ratification of 3GPP 6G ISAC standards will have consequences for how a software-led sensing capability interoperates with native-6G sensing in the longer term. Operators investing today should treat it as a near-term capability that buys early (and timely) commercial experience, not as a substitute for engagement with the longer-term standardisation process. At the time of writing, Cohere is the principal vendor offering this specific capability profile, although the broader sensing-software space includes specialist players with different architectural models based on xApps and rApps on the RAN Intelligent Controller. Vendor-due-diligence considerations apply: roadmap durability, OEM-partnership posture, and integration paths with the operator's existing RAN supplier relationships.

CONCLUSION

The drive for more frequent telemetry data from mobile networks in pursuit of more real-time insights and better network performance can be seen as an end in itself. But that challenge also represents an opportunity: can the same data be the basis for extracting new value?

The mobile industry's “beyond connectivity” quest is increasingly urgent as the cost and complexity of networks rises, and as revenue-per-bit-transported falls. As other industries have shown (and even more so in an AI era) data can be the key to new value, and to new services for new customers. In telecom, the network is the richest source of data, but its sheer volume represents an obstacle to providing ever more real-time insight into the health of networks and services. Solving this conundrum helps operators meet high-value SLAs, contribute to user satisfaction and high network/service quality ratings.

But harnessed to AI engines, the data used to build the model of real-time channel conditions can reveal valuable information not about the network, but about the real-time physical environment around the network. This could provide a new source of revenue for telcos, in the same way that other companies turned secondary information into primary value: Google's search-query data, Mastercard's transaction analytics, Bloomberg's pricing exhaust, Climate Corporation's hyperlocal weather signals.

The same processing that lifts capacity on congested cells, and closes the operator's long-standing observability gap, exposes a real-time, sub-metre, motion-resolved view of the cell's environment — sensing-grade telemetry generated as a by-product of better scheduling, on installed 4G and 5G infrastructure, without the substantial per-site hardware programme that a native 6G ISAC deployment will require.

This is not the early realization of the full 6G ISAC vision. It does, however, make a meaningful subset of high-value sensing applications — defence, public safety, airspace surveillance, smart-city services, and sensing-as-a-service — commercially addressable several years earlier than the standards-led timeline suggests.

For operator strategy teams, the implication is not to “wait for 6G.” It is to begin building the operational, commercial and governance capabilities that turn an installed-base capability into a recurring revenue line, before peer operators and over-the-top sensing-data aggregators do the same. Cohere’s ECHO is, in effect, the gold hidden in the network: a way to transform discarded engineering data into a high-value commercial asset. It also demonstrates realizable commercial potential of the AI-RAN concept – long before any 6G network refresh – offering an extended economic lifespan for millions of installed 5G base stations.

Insight and analysis for telecom transformation.

@AppledoreVision

Appledore Research

www.appledoreresearch.com

info@appledorerg.com

+1 603 969 2125

44 Summer Street Dover, NH. 03820, USA