



BLOG POST: REAL-TIME DATA FOR AUTONOMOUS NETWORKING

Telcos remain ambitious about network autonomy, but most AI deployments are still narrow, domain-bound and far from real-time. STL Partners' research with 11 senior telecoms stakeholders points to operators' data foundations as the limiting factor. This article highlights that agentic systems need both real-time event data and discoverable context to supervise automated networks, and sets out three priorities for building the data architecture that makes network autonomy possible.

Magnus Jeffery, Consultant

September 2026

© STL Partners

Telcos' ambitions outpace their data foundations


It has been a long and ongoing journey, but telcos are still highly ambitious about what network autonomy can achieve. AI, including agentic systems, promises to help them manage more complex networks, improve service assurance and respond to unfamiliar conditions. Yet most AI deployments remain narrow: they are often confined to one network domain, avoid operationally critical decisions, and are far from real-time. STL Partners' research, based on interviews with 11 senior telecoms stakeholders, suggests that this gap reflects failings in operators' data foundations more than the maturity of AI itself. To increase their level of network autonomy, telcos must develop suitable data architectures which can incorporate both real-time and network context data.

Agents need context to interpret live data

As operators move towards more autonomous networks, deterministic automation of mature and well-understood processes such as load balancing and certain security responses will expand. Additionally, AI agents are expected to sit above this automation, monitoring network conditions and assessing whether automated responses remain consistent with the operator's operational and commercial objectives. They may also recommend interventions when conditions fall outside established rules.

To perform this supervisory role, agents will need two types of information: real-time, detailed data about current network events and reliable context with which to interpret them. This context will include historical performance, network topology, available capacity and service-level commitments. The context must be current, discoverable and accessible at the point of decision. Clear schemas, consistent labels and controlled interfaces allow agents to retrieve information without manual translation. Enriching data with operator-specific knowledge can improve reasoning quality, reduce inference work and make recommendations easier to validate.

Figure 1: Real-time agentic decision-making requires a complete picture of live event data, compared against known context



Source: STL Partners

Why current data practices fall short

Current data practices do not always support these requirements. Firstly, though data is often stored in a central system, it is often inaccessible across domains and insufficiently standardised. These issues restrict the potential of engineers and the agents they create. Telcos must ensure that their data is accessible, discoverable and trusted, overcoming long-standing domain-level data silos. This challenge is primarily an organisational and political one; domains have long controlled their own data (e.g. RAN teams owning the RAN data), and teams may resist wider access if it weakens their authority or risks their data being misinterpreted. Telcos therefore need to maintain clear domain ownership of data, and implement usage policies and controls that show who can access each data product and for what purpose. Without these safeguards, calls for greater accessibility risk introducing friction, rather than trust.

Secondly, due to the petabyte-scale of telco operations, network data is frequently sampled, aggregated and transferred to central platforms to control processing, transport and storage costs. This approach works well for reporting, historical analysis and model training. However, it can remove the detail needed to detect emerging anomalies and introduces latency that make the data less useful for real-time operational decisions. Telcos must therefore address the challenge of providing trusted data across domains, at network scale, and at the speed each decision requires.

Three priorities for telcos

Telcos face the challenge of improving their data architectures without interrupting services, decreasing network resiliency, overwhelming transport and compute capacity, or creating an uneconomical system. To achieve this, our research points to three priorities for telcos:

Treat data as a product. Domain teams should retain ownership of their data while publishing it to common organisational standards. Each data product needs a clear owner, defined quality, consistent metadata and schema, accessible interfaces and service commitments. This makes data easier for both people and AI systems to find, interpret and use. A central function should set and enforce these standards, while catalogues or internal marketplaces make data products visible across the organisation.

Be selective about full-fidelity processing. Not every data stream needs to be processed in complete detail. Operators should identify which agentic use cases depend on unsampled, low-latency events and reserve full-stream processing for those cases. Other data can continue to be aggregated or sampled.

Combine central and distributed processing. Central data platforms remain important for governance, enterprise analytics and long-term knowledge. Real-time operational use cases will increasingly require processing closer to the source. A hybrid architecture can keep high-volume event streams local while retaining shared standards and central oversight. We are starting to see emerging examples of real-time, distributed intelligence in the networks, for instance as operators implement RAN intelligent controllers (RICs) in the radio domain.

Setting the data foundations now

Operators do not need to wait for a fully unified data estate before pursuing network autonomy. They can start by establishing data product disciplines for current, less time-sensitive use cases, then extend the same governed foundation towards real-time operations. As real-time data requirements increase, a flexible data architecture with distributed processing will become increasingly important to control costs and improve agentic AI performance; telcos should act now to set the data foundations that enable future autonomy.

September 2026

© STL Partners

A message from our sponsors

These quotes are from the webinar delivered by STL in June 2026, supported by our sponsors Volt Active Data & Ocient. Find the link to the webinar recording here:

<https://attendee.gotowebinar.com/recording/4698439041509104044>

Quotes from the webinar

Volt Active Data

Andrew Keene, Volt Active Data: "The main barriers to real-time agentic AI identified in STL's research, beyond trust, are inference latency and data currency... it's not going to make an accurate decision if the context it's given is already stale and out of date."

Andrew Keene, Volt Active Data: "Over the last five years or so, walking around MWC or DTW talking to customers, we've gone from hype – AI everything, everywhere, claims on stands that really was absolutely everything AI-enabled or AI-powered or AI-native, whatever that all meant, with little to back it up when you looked under the covers – to a marked shift in focus on agentic AI over the last couple of years."

Ocient

Jai Rajaraman, Ocient: "You can't be asking the agent to constantly look at the entire treasure trove of data that you have, real-time as well as historical, in order to determine the best path to action. You've got to be able to train it to understand what is more important, what is more contextual – and that's going to save you cost, and it's actually going to speed up the time to action"

Jai Rajaraman, Ocient: "You're dealing with a torrent of data – you're drinking from a fire hose. How do you organise that? ... You've got to organise that data in a way that is productised and available to other domains, and you have to do it in real time."

Magnus Jeffery is a Consultant at STL Partners, specialising in AI & network APIs.

Get in touch with the author to learn more

magnus.jeffery@stlpartners.com

Or visit STL Partners' AI Hub

www.stlpartners.com/AI

September 2026

© STL Partners