Research Note

October 2026

AGENT-READY DATA

Is telco data ready for AI agents, and what are operators doing to make it so?

Author: Robert Curran and Patrick Kelly



Sponsored by:

amdocs

Appledore RESEARCH

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

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

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

Figures and projections contained in this report are based on publicly available information only and are produced by the Research Division of Appledore Research LLC independently of any client-specific work within Appledore Research LLC. The opinions expressed are those of the stated authors only.

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

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

Publish date: 27-Oct-26

Cover image: Patrick Kelly

SUMMARY

Operators believe agentic AI is moving fast. In Appledore's March 2026 survey of 54 Tier 1 and Tier 2 communications service providers (CSPs), 53% said agentic AI was already in production and 69% expected meaningful returns within two years. Yet most placed their network autonomy at Level 2 or Level 3, and most of what is described as agentic today is assistive: copilots and recommendations, with a person making the final call.

To understand that gap, we put two questions to network, data, architecture and AI leaders at a selection of operator groups in Asia, Europe and North America. Is your data ready for AI agents? And what are you doing to make it so? This note sets out what they told us.

The short answers are "no", and "a lot". Telco data is largely available, but it is not yet agent-ready. Operators are converging on a recognisable set of principles to close the gap, and most of them concern how data is owned, linked and governed rather than the AI itself.

Question	What operators told us
Is telco data agent-ready?	Mostly not. Data is plentiful but fragmented, poorly linked, often missing history, and exposed to agents one connector at a time. Readiness is being achieved use case by use case, not across the enterprise.
What are operators doing about it?	Starting from use cases; publishing data once as owned data products; engineering the links between domains; fixing data at source; keeping execution deterministic; earning autonomy with evidence; and keeping agent estates simple.
Why does it matter?	Curated data changes agent accuracy dramatically, and better data shortens the road to trusted autonomy.

IS TELCO DATA AGENT-READY?

Operators are not short of data. They hold vast amounts of network, operations, customer and business data, and one European Tier 1 described its raw data as generally of good quality. The problem is that very little of it is in a form an agent can use.

We define agent-ready data as data that an agent can find, understand and safely act on. *Stored* data provides visibility; *agent-ready* data provides understanding.

Data as most operators hold it	Agent-ready data
Stored, often in abundance	Findable through a governed interface
Fragmented across systems, each with its own model of the network	Fresh enough for the decision being made
Relationships implicit, often held in people's heads	Linked to its relationships and dependencies, now and historically
Good enough for dashboards and reports	Governed well enough to act on, with lineage for audit

Five deficiencies reveal why most operator data is not yet agent-ready:

- **Gaps in Meaning.** Putting lots of data in one place does not correlate it. The same labels (customer, site, service…) may have different meanings across OSS, BSS and customer-care systems. Even time is ambiguous: a billing timestamp records when a session started, not when the event of interest happened. A focus on simply bringing data into one place can delay appreciation of the missing meaning. "You don't see this until you start, and we hit the wall," said the architecture lead at an Asian Tier 1 group.
- **Gaps in History**. Operators usually know today's network topology. They rarely know what it was six months ago, which is exactly what an agent needs to learn from past incidents, or to be tested against them.
- **Variation in Freshness at source**. Network elements do not emit useful data by default; what is generated, in what format and how often, must be configured – and may be changed over time. Data streams are broken by reboots and faults, so raw telemetry needs curating before an agent can rely on it.
- **Non-scalable Exposure**. The common early pattern is a bespoke connector or MCP binding for each agent. Every operator we spoke to had done this. Those further along said this exposure model breaks as agent numbers grow.
- **Data debt.** Past OSS consolidation programs sometimes skipped data cleansing as too hard and too expensive. A European operator group's architects said that autonomous network programs are now effectively presenting a bill for that skipped work.

The roots are often organisational. A US Tier 1 was reported to have four incompatible data sets describing what should have been standard IP traffic – a consequence of incomplete integrations of wireless and wireline businesses that grew up separately.

The prize for fixing that problem is hugely valuable for agentic operations. A North American Tier 1 reported that its agents reach roughly 80% accuracy on difficult public benchmarks that mix many messy data sources, but approach 100% when the same techniques run over its own curated data. The figure is self-reported, but its direction matches everything else we heard: much of what looks like a *model* limitation is actually a *data* limitation.

One qualification shapes everything that follows. No operator is trying to make all its data agent-ready. Readiness is judged per use case: is the data this agent needs present, complete, fresh enough, linked and trusted?

 © Appledore Research 2026

SEVEN PRINCIPLES OPERATORS ARE ADOPTING

The operators we spoke to work on different continents, with different toolchains and at different stages of maturity. Their approaches nonetheless converge to a notable degree. We see seven principles.

1. Start from the use case, not a data clean-up

No operator is waiting for clean data before building agents. Each starts from a valuable, bounded use case, works out exactly what data it needs (including history and freshness), checks that the joins exist and the data is complete, and then builds. A North American Tier 2 operator put it simply: AI readiness does not require enterprise-wide data readiness, only readiness for the use case being automated.

The discipline matters as much as the sequence. A European Tier 1 requires every proposed use case to carry a business case and KPIs, and measures them after deployment. That is a useful corrective to counting agents as a measure of progress.

2. Build data once, as owned data products

This was the strongest point of agreement in our research, and it came up largely unprompted. Three of the five operators have independently arrived at the same answer to the exposure problem: curated data products, each with a named owner accountable for quality.

The North American Tier 1 runs an internal catalogue of data products. Each draws on several source systems, is cleansed and certified by a domain owner, and carries service levels and a support team. Teams used to treat pulling raw data straight from source systems as a badge of honour; unwinding that habit is what created single sources of truth. The finance function went furthest, declaring itself the certified owner of core financial data products. Anyone bringing other data into a financial discussion has to justify why.

A European Tier 1 is building the same model on a data virtualisation layer, so that an agent asks for a data product without needing to know which of several data lakes holds it. Lake consolidation continues in parallel and will be invisible to the agents above. A European operator group curates streaming telemetry into data products with owners, quality standards and metric checks.

Two details matter. Products that span domains, such as network performance joined to customer and product data, need an owner who sits across those domains rather than the owner of any one source. And the payoff is reuse: built well, the second use case on a data product costs far less than the first.

3. Correlate, don’t just co-locate

Correlation is designed, not discovered. The Asian Tier 1 named four axes that must be engineered explicitly to join data across domains: subscriber identity, time, location and network topology. Without them, an agent falls back on statistical guesswork.

Operators are taking different routes to that context. The North American Tier 2 layered AI over separate transport and IP metro systems and got only moderate results; for the same operator, converging the data into one lakehouse at the foundation made AI substantially more effective. A European operator

group writes domain-specific ontologies as version-controlled code, starting small and extending them as use cases are added, to feed knowledge graphs and digital twins. The North American Tier 1 builds graph models across fragmented network data to spot deviations from normal, particularly for security and fraud, and argues that every agent needs both a defined way of operating and deep context about its environment.

Our view is that physical convergence of data works well across two domains but does not scale to the number of domains, sources and agents that autonomous operations will involve. Beyond that point, explicit semantics stops being optional. A telecom-trained language model helps interpret telecom data, but it cannot supply the ontological links between an operator's own systems.

4. Fix data at the source

Several operators pushed the readiness problem upstream. The North American Tier 1 described observability as step one: without high-fidelity data that spans network, cloud, application, identity and time, everything built on top "falls down like a deck of cards." A European operator group argued that operators must professionalise how network elements are configured to emit data, and calibrate the acceptable freshness according to the decision being made: minutes for closed-loop fault handling, months for network planning.

Data retention is the other factor. The Asian Tier 1 found completeness, especially retained history, a bigger issue than quality. The North American Tier 2 added a reminder about ground truth: the network element and its live configuration are the truth, not the provisioning record.

5. Let agents reason; let deterministic systems act

Every operator keeps the irreversible act out of the hands of a free-running agent. The Asian Tier 1's rule is that no agent may touch a network element: reasoning can be agentic, but execution goes through deterministic domain controllers. The North American Tier 2 has built around 400 deterministic workflows over six years and placed an AI layer above them. Its agents can diagnose faults, trigger workflows and close tickets, but may not change the network on the strength of an inference. The North American Tier 1 lets agents run root-cause analysis up to the point of decision, then hands live configuration changes to a human.

A European Tier 1 made a point worth repeating: **autonomous does not have to mean agentic**. Much autonomous operation is, and will remain, deterministic automation with no AI in it. Agentic AI makes that automation more valuable rather than replacing it.

6. Earn autonomy with evidence, scaled to blast radius

Operators agree that autonomy must be earned, but differ on how. A European operator group runs agents in shadow mode alongside human or deterministic processes, logging proposed actions and confidence scores against real outcomes. Only when humans have consistently approved decisions at high confidence, of the order of 99% over a year, does a class of decision move towards automated execution. A European Tier 1 adds a digital twin to simulate impact and prove rollback before any agent touches live network elements.

 © Appledore Research 2026

The North American Tier 1 challenged the timescale. In AI, “a year is like a decade,” it argued. The proving period should depend on the complexity of the decision, confidence in the data and the blast radius of a mistake, and high-risk use cases should simply not go first. Its proof point lies outside the network: external auditors approved autonomous agents for a regulated financial-compliance process once the operator showed empirically that agent accuracy matched human performance.

These positions are closer than they sound, and they bring the argument back to data. Curated, owned data raises agent accuracy. Higher accuracy shortens the time needed to prove an agent. A shorter proving period lets autonomy expand. Data readiness is therefore not only a precondition for agents to work; it sets the pace at which they can be trusted. It also explains operators’ scepticism about Level 4 and Level 5 as blanket targets. Autonomy will advance one use case at a time.

7. Keep agent estates simple, and test before you trust

Multi-agent orchestration is the least mature area we examined. The North American Tier 2 found that a supervisor agent coordinating specialised sub-agents worked at small scale, then degraded as agents and tools multiplied and agents became confused about which should act. The Asian Tier 1 had predicted exactly that from first principles, calling super-agent hierarchies “a very nice piece of PowerPoint engineering.” The North American Tier 1, which runs tens of thousands of agents across the enterprise, observed that getting two agents to work together reliably is hard enough, let alone ten.

Standard protocols such as MCP and agent-to-agent (A2A) are seen as the realistic lowest common denominator for connecting agents, including across vendors. But they provide connectivity, not policy, trust or conflict resolution. Nor does every product sold as an agent deserve trust: operators intend to test vendor agents empirically against their own data before connecting them to anything consequential.

A note on data placement

One further consideration came from a single operator but deserves attention. The Asian Tier 1 argued that the cloud is an excellent home for analytics but cannot run network operations at scale: telemetry volumes would overwhelm it, and ingress and egress charges now outweigh storage costs. It distinguishes long-lived “golden” analytics data from high-volume, short-lived operational data, which need different structures and different homes. Where data sits is becoming a readiness decision, not just an infrastructure one.

WHERE OPERATORS STILL DIFFER

Sequencing. Every operator works use case by use case, but they treat the data work differently. The Asian Tier 1 uses each use case as the instrument for fixing data, through a standing task force drawn from its operating companies that selects use cases, establishes joins and audits completeness as one activity. The North American Tier 2 treats data work as the cost of getting a use case live, and ranks data third among its AI spending priorities. The European operators treat data as a programme in its own right, with product owners, a catalogue and a virtualisation layer built alongside delivery. All three approaches deliver; the third is the one designed to scale.

Proving periods. A year of shadow running at high confidence sits at one end; evidence-led approval tiered by blast radius at the other. Expect the second to gain ground as curated data improves agent accuracy.

Context. Physical convergence, data products over virtualisation and ontology as code are all in use. Few operators yet have an enterprise ontology or knowledge graph. Most see standards work, such as **TM Forum**'s AI and data activity and **IETF** work on knowledge graphs, as a useful frame rather than a solution, and expect ontologies to remain domain-specific.

CONCLUSION AND RECOMMENDATIONS

Telco data is not yet agent-ready, but the operators furthest ahead no longer treat that as a reason to wait. They are building data foundations alongside their first use cases: owned data products, explicit links between domains, better data at source, and governance that lets autonomy be earned.

Agentic AI in network operations is a data and operating-model transformation more than a question of AI models. And as agents begin to exchange state across company boundaries (as one North American operator already does with an enterprise customer's agent for fibre fault and restoration updates), shared meaning will matter even more.

- ☐ Available is not ready: telco data is plentiful but rarely findable, linked, complete and governed enough for agents to act on.
- ☐ Data products are the emerging consensus: publish data once, with a named owner, for many agents.
- ☐ Better data sets the pace of autonomy: curated data lifts accuracy, shortens proving periods and lets autonomy expand one use case at a time.

Based on our discussions, we offer the following recommendations in relation to the data component of agentic AI strategies:

For operators

1. Approve agents use case by use case, with data requirements, freshness, history, business case and KPIs written down first.
2. Stop building a connector per agent. Stand up governed data products with named owners now, in parallel with delivery.
3. Engineer identity, time, location and topology joins before expecting cross-domain reasoning.
4. Instrument the network for the data agents need, and retain history, especially topology history.
5. Keep execution deterministic and define in advance which classes of action need human approval.
6. Set proving periods according to blast radius and evidence rather than a set time period, with shadow mode, confidence scoring and rollback built in from the start.

 © Appledore Research 2026

7. Measure data readiness and business outcomes, not the number of agents deployed.
8. Develop an enterprise governance architecture that decides which agents are trusted, what authority each holds and when human approval is mandatory.

For suppliers

1. **Deliver the correlation layer, not just another agent.** Reconciled data and semantics that agents can reason over, ideally as governed data products with named owners and service levels.
2. **Start with a vertical slice and a readiness check.** Value in weeks, with joins and completeness tested up front.
3. **Work with the operator's data estate.** Virtualize, federate and place data sensibly, and address egress cost and residency.
4. **Build guardrails in.** Deterministic execution, digital twins, measurable trust thresholds per agent, and explainability.
5. **Be open and interoperable.** Standards, secure and authorized tool access, and cooperation with other vendors' agents.
6. **Prove value with agreed KPIs.** Mean time to repair, opex and truck rolls, measured before and after.

OSS vendors must show they can expose trustworthy operational context to external agents and take part safely in multi-agent workflows. **Observability and data platform vendors** gain additional importance, since AI cannot reason about infrastructure that it cannot see.

Insight and analysis for telecom transformation.

@AppledoreVision

Appledore Research

www.appledoreresearch.com

info@appledorerg.com

+1 603 969 2125

44 Summer Street Dover, NH. 03820, USA

© Appledore Research LLC 2026

Appledore
RESEARCH