L&T Technology Services

By the time a RAN alarm fires,

the customer already knows.



CASE STUDY

EXECUTIVE SUMMARY

The client is a **Tier-1 US mobile network operator and one of the world's largest telcos.** LTTS ran two programmes for them, both making the same move: replacing human reaction with engineered anticipation.

On the network, LTTS built an **AI advisory decision-support system** to move radio-access support from reactive to proactive: anomaly-detection models over **20 TB of live, historical, contextual and expert data,** engineered against a target of flagging service-impacting failures more than 24 hours ahead. On the commercial side, LTTS automated the network-service testing that gates every wholesale brand launch, cutting regression cycles **from hours and days to minutes.**

THE CHALLENGE

A national radio network fails constantly in small ways, and the operator finds out from its customers. Fault management is alarm-driven, and an alarm is a receipt: it confirms something already broke. Truck rolls follow. Tickets follow.

The same reactive posture ran through the wholesale business. Launching a new MVNO brand on the operator's network meant manual testing across every environment with several vendor teams in the chain. Long lead times, no service-uptime reporting, and customer-experience tickets as the first signal of degradation.

The operator wanted the same thing in both places: **to know first.**

THE ENGINEERING COMPLEXITY

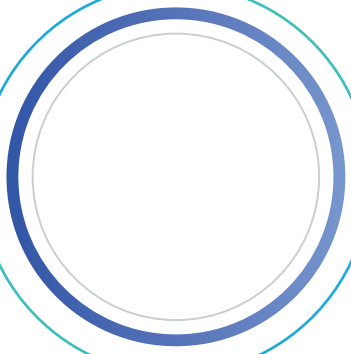
Both problems are hard for reasons that do not show up in a demo.

01

An alarm is a receipt, not a warning.
The alarm channel is also noisy. One published analysis of a seven-day operator dataset of over 5.5 million alarms found roughly 49% were chattering alarms that filtering removed outright. Prediction has to come from performance data, not the alarm stream.

02

The worst failures raise no alarm at all.
The industry calls it a sleeping cell. A granted patent defines it as a cell outage that makes mobile services unavailable to subscribers while, from the network's point of view, the network still appears operable; no alarm reaches the operations centre. Finding it means reasoning over traffic patterns, not faults.



03

Almost nothing fails, which is the problem.
This is a severe class-imbalance problem. A peer-reviewed study of a real operator network of roughly 10,000 base stations, each with about a dozen replaceable components, found only a few hundred faulty components replaced per year. There is no rich labelled failure set to learn from, and the network keeps changing under the model.

04

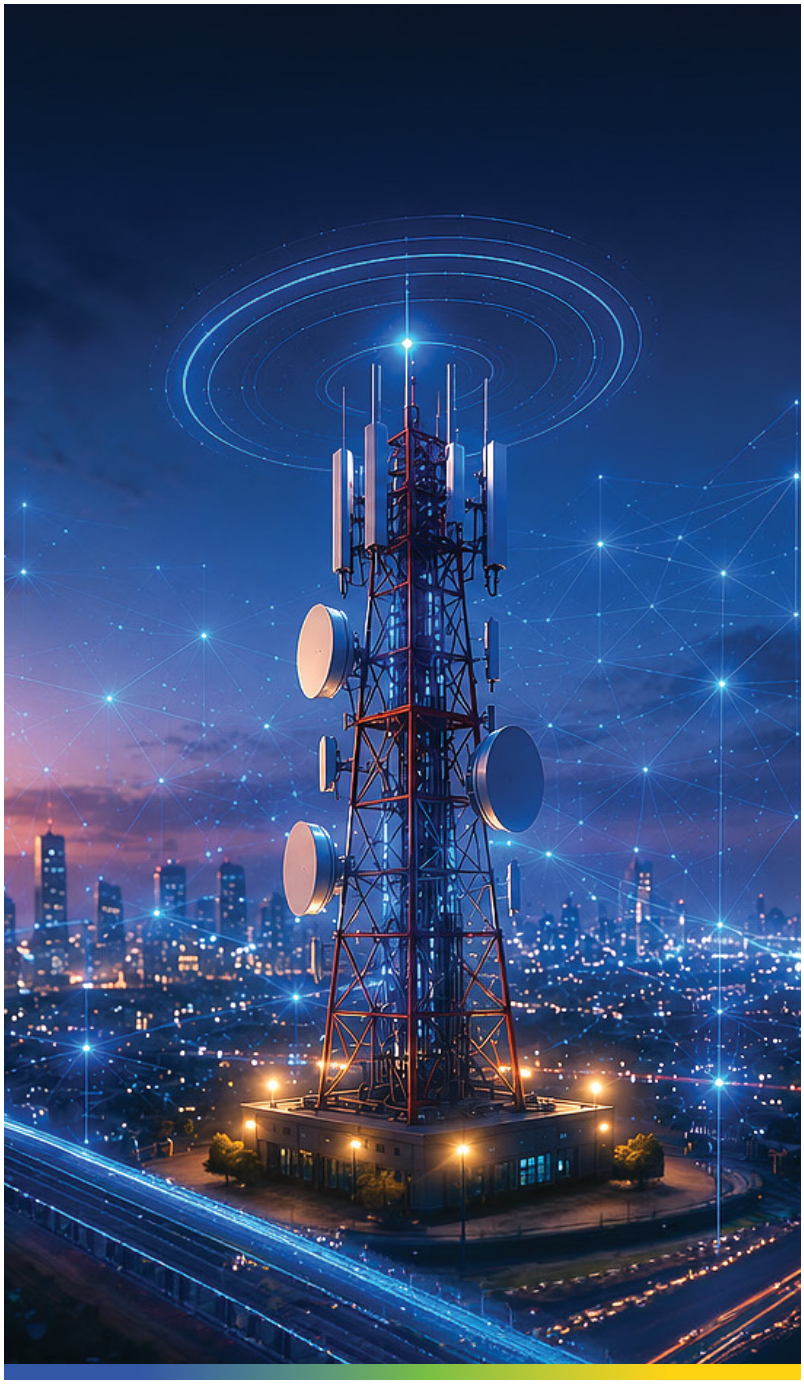
A wholesale launch is a combinatorial test problem.
Every new brand means revalidating voice, SMS, MMS and data, plus policy behaviour such as throttling after a data allowance expires. That runs across environments, devices and OS versions, with third parties in the chain, each interface measured against TM Forum Open API and GSMA specifications.

WHAT LTTS ENGINEERED

Two programmes, engineered against the difficulties above.

Predictive RAN: an AI advisory decision-support system

The system scores live network data continuously, alongside historical, contextual and expert-supplied inputs, rather than studying it after the fact. Hardware faults, backhaul degradation and software faults are modelled separately, because each leaves a different signature on different data, and because labelled failures are too scarce for one approach to cover all three. The system raises anomaly and failure alerts to the operations team. The decision to change the network stays with the people accountable for it.



Wholesale launch assurance: automated network-service testing

The LTTS IntelliAgent platform validates voice, SMS, MMS and data on real devices, unattended, across every environment, and covers the wholesale web and mobile applications in the same suite. Service availability is probed continuously, so "are the services up?" becomes a reported question rather than an unanswered one. Testing also reaches policy states that manual regression rarely does, such as throttling after a data allowance expires.

THE FEAT · ANTICIPATION IN TWO PLACES

A national radio network where the most damaging failures raise no alarm, and a wholesale launch path where every new brand used to mean starting the test cycle by hand. Both moved from reacting to knowing first.

BUSINESS OUTCOMES

Outcome	What changed	Why it is credible
Regression test cycle	Hours and days to minutes	Device-based service tests and UI workflows run unattended across environments, instead of being executed by hand across several vendor teams.
Test effort and efficiency	45 to 50% less test design effort; 75% efficiency gain	A reusable framework built to TM Forum and GSMA specifications means each new brand inherits a validated suite instead of authoring cases from scratch. Automation removes the wait states that dominate a manual cycle: environments, cross-team handoffs, manual result capture.
Fault discovery	Post-validity throttling faults found	A live customer-experience defect, found by a test suite rather than by a subscriber.
Service availability	Reactive to proactive detection	Continuous automated probing surfaces availability gaps before customers raise tickets. Previously there was no service-uptime reporting at all.
RAN prediction lead time	Design target: over 24 hours at 95% confidence	The system was specified and built to that lead time, which is what turns maintenance from a dispatch queue into a schedule.
Data foundation	20 TB over one year, streaming	Kafka ingestion into elastic containers, scoring live data alongside historical, contextual and expert inputs.
Maintenance economics	10 to 15% less downtime; 5 to 8% more productivity	Modelled benchmarks from comparable predictive-maintenance programmes on rotating and legacy equipment. The mechanism is lead time: scheduled work costs less than an emergency.

ENGINEERING INTELLIGENCE IN ACTION

1. **Domain depth, not just data science.** Data science alone would have produced models nobody could act on. A test-automation vendor alone would have automated a broken process faster. The difference was radio-network domain knowledge fused with AI and automation engineering.

2. **Models built around real failure physics.** Three failure families, three signatures, three sets of variables. That is why the output is credible to the engineer who owns the site, and why it arrives as ranked advice rather than an unexplained score.

3. **Tested to the industry's own specifications.** TM Forum Open APIs and GSMA specifications, not one operator's local interpretation. That is what makes the suite reusable for the next wholesale brand instead of the next rewrite.

This collateral includes AI-enhanced imagery

L&T Technology Services

ENGINEERING THE CHANGE

Copyright © L&T Technology Services