

Anyone can put

AI in your network data.

The key is in **understanding your network.**

Six production AI agents for telecom operations- built by engineers who know fronthaul topology, rectifier behaviour and tower drawings. Running today on live operator data.



Why This Matters Now

A single fronthaul fault throws thousands of alarms, every one of them true and none of them the answer. A rectifier degrades for days before it takes the battery, then the site. A cell-site bill of materials costs an engineer days of counting antennas off drawings, and a miscount surfaces only when the crew is already on the tower. **None of this is a data problem. Operators have the data. It arrives faster than anyone can act on it.**

Autonomy is the answer the industry has agreed on; TM Forum Level 4 networks that sense, decide and close their own loops. What has been missing is agents that work on real network data, inside real operator constraints, at real scale.

These six from LTTS do. We are pure-play engineering services company that has spent 25 years building and operating the world's most complex engineering systems. That depth is why these agents understand fronthaul topology, rectifier behaviour and tower drawings, and not just tokens. It is the part an AI-first firm cannot replicate.

The Six Agents

01

Robot Management System (RMS) Agent

Keeps autonomous mobile robots out of hazardous zones, over private 5G, without new tracking hardware.

02

Root Cause Analysis (RCA) Agent

Reads the alarm storm and names the actual fault (and the fix).

03

Computer Use Agent (CUA)

Automates the OSS portals and legacy tools that will never have an API.

04

BoM Creation Agent

Reads site drawings and produces the bill of materials directly.

05

Failure Prediction Agent

Calls service-impacting failures 24 hours before they happen.


06

Interference Prediction Agent

Forecasts atmospheric ducting interference a day ahead, base station by base station.

01 Robot Management System (RMS) Agent

Autonomous robots that stay out of harm's way, guided by the network itself.



About the Robot Management System

On a factory floor, an autonomous mobile robot does not know it is about to drive into a hazardous zone. And the safety system that does know usually cannot see where the robot is. The RMS Agent closes that gap over a private 5G network: it tracks every robot in real time, detects when one approaches a restricted or dangerous area, and routes it to the nearest safe point or charging station before anything happens.

What sets it apart is how it handles permission. **Monitoring and safety are separated into two agents with deliberately different access rights.** The safety agent holds no location data of its own — it must request the robot's position from the monitoring agent through a governed agent-to-agent exchange. Least privilege is enforced by the architecture, not by a policy document.

HOW IT WORKS

- A 5G camera streams the floor live. A computer-vision model identifies each robot and measures its distance from a buffered hazard boundary.
- The monitoring agent retrieves the robot's real-time position from the network's Location API — the network becomes the sensor, so no additional tracking infrastructure is deployed.
- The safety agent requests that position over an agent-to-agent protocol and runs the proximity check, without ever holding location access itself.
- On a breach, the agent issues a navigation instruction and the robot reroutes to the nearest safe location or charging point.
- A Quality-on-Demand API reserves the network profile the control loop needs, so the safety command never queues behind other traffic.



BUSINESS BENEFITS

01

Continuous safety enforcement — hazard prevention runs without a human watching a screen

02

No new hardware — the 5G network supplies location, so tracking infrastructure is not duplicated

03

Access boundaries that hold under automation — regulated environments stay auditable as autonomy scales

04

Guaranteed control quality — safety commands are delivered on a reserved network profile, not best effort

02 Root Cause Analysis (RCA) Agent

Thousands of alarms in. One root cause identified along with the fix.




About the Root Cause Analysis

When a 5G site fails, it does not send one alarm. It sends thousands — cascading across fronthaul, transport, timing and power. Every alarm is accurate. None of them is the answer. Engineers correlate by hand across topology and protocol layers while the service outage clock runs.

The RCA Agent does that correlation inside the network. It builds a live map of how every element actually connects — radio unit to distributed unit to central unit, cell to site, interface to link, primary path to redundant path — then reads alarms and performance counters **against that structure**, rather than in isolation. The output is not a shorter alarm list. It is a named cause and a recommended action.

HOW IT WORKS

- Three live feeds — alarms, performance counters, and topology and inventory — are normalised into a common structure aligned to O-RAN and TM Forum standards.
- Those feeds build a knowledge graph capturing hierarchy, physical and logical connectivity, cell-to-radio mapping and redundancy paths.
- A graph neural network reads the graph in context and classifies the fault into a root-cause category.
- The agent returns the cause and the recommended remedy — for example, a fronthaul link-down alarm resolved to a Layer-2 link failure caused by an MTU mismatch, with “restore MTU” as the action.
- The result is passed to fault management as an intent: automated remediation, a trouble ticket, or a field dispatch — whichever the situation warrants.



BUSINESS BENEFITS

01

Faster service restoration — isolation moves from hours of manual correlation to an immediate, reasoned answer

02

Fewer wasted truck rolls — dispatch decisions are made against the actual cause, not the loudest alarm

03

Expertise made institutional — diagnostic knowledge lives in the graph, not in a handful of senior engineers

04

Drops into existing OSS — standards-aligned to O-RAN O1 and TM Forum Open APIs, so it augments rather than replaces

03 Computer Use Agent (CUA)

Automates the systems no one will ever build an API for.




About the Computer Use Agent

Every operator runs critical work through tools that will never expose an interface — OSS portals, vendor dashboards, server management consoles, PDF manuals. The work is repetitive, high-volume and entirely manual. The integration project that would fix it never clears the business case.

The Computer Use Agent works the way the engineer does: **it looks at the screen, reads what is there, and clicks.** Same portals, same credentials, same screens — no integration, no vendor dependency, no rollout project. Two are already running: pre-installation system checks, which verify server state against the vendor manual before software is installed, and mobility load balancing, which reads live cell KPIs and adjusts load-balancing parameters through the OSS portal.

HOW IT WORKS

- The agent receives a task in plain language — a validation checklist, a configuration change, a report to pull.
- It captures the screen, and a vision-language model determines the next action: click, type, scroll, navigate, upload.
- It executes, re-reads the screen, and repeats — reconciling what it sees across live dashboards, PDF manuals and system screenshots.
- Anything requiring judgement is escalated. The agent proposes; the engineer approves or skips. It never commits a judgement call alone.
- Results are written to a spreadsheet — a complete record of what was checked, what was found and what was changed.



BUSINESS BENEFITS

01

Automation without integration — no API, no middleware, no vendor roadmap dependency

02

Engineering hours reclaimed — skilled engineers move off checklist work and back onto engineering

03

Human-in-the-loop by design — the agent asks before it commits anything consequential

04

Audit-ready by default — every run leaves a complete, exportable record

04 BoM Creation Agent

Reads the site drawings. Writes the bill of materials.


About the BoM Creation Agent

Before a cell site can be built, someone has to count it. Antennas, radios, jumpers, surge protectors, mounts; read off AutoCAD designs, red-line construction drawings, plumbing diagrams and mount analyses, one page at a time, into a spreadsheet. It takes days per site. It does not scale with rollout pace. And a miscount does not surface until a crew is on the tower without the part.

The BoM Creation Agent reads those drawings the way an experienced deployment engineer does — **interpreting the diagram, not just the text on it** — and produces the bill of materials directly. It has been proven on live 5G site packs: antenna models and counts, remote radio heads, surge protection units, RF jumper counts traced from plumbing diagrams, tower elevations and antenna plans.

HOW IT WORKS

- Site design packs are converted page by page into images.
- The agent locates the relevant region on each page — an antenna schedule, a tower elevation, a plumbing diagram — and crops to it.
- A vision-language model answers a structured question for each component category: what part, which model, how many, mounted where.
- It reads what the drawing shows, not only what it states — counting cable runs between blocks on a plumbing diagram, or antenna quantities noted in brackets on an elevation.
- Confirmed quantities are written directly into the bill of materials sheet, ready for engineering review.



BUSINESS BENEFITS

01

Scoping compressed — days of manual page-by-page counting become a reviewed output

02

Fewer material errors — and fewer aborted site visits caused by them

03

Rollout no longer capped by scoping capacity — deployment velocity stops depending on headcount

04

Consistency across every site — the same engineering standard applied to every drawing pack

05 Failure Prediction Agent

The signals appear a day early. This agent calls the failure before it hits.


About the Failure Prediction Agent

Most network failures announce themselves. A rectifier degrades before it fails. A base station's counters drift before the cell drops. A microwave link deteriorates as weather closes in. The signals are all there — buried in tens of terabytes of performance and fault records that no operations team can read in time.

The Failure Prediction Agent reads them, and calls the failure **24 hours before it happens**. It has been validated across four failure modes on live operator data: base station faults at 99% precision, rectifier loss at 96% precision, microwave backhaul radio link failure using weather and KPI correlation, and cell-site power outages during extreme weather — predicted 10 to 16 days ahead, so mobile cell sites can be positioned for emergency responders before the storm makes landfall.

HOW IT WORKS

- Raw fault-management and performance-management data is ingested continuously at 15-minute granularity across the estate.
- A forecasting model projects every key indicator 24 hours forward.
- A classifier reads that forecast and flags the sites heading for a service-impacting failure.
- Explainability output shows which indicators drove each prediction, so an engineer can act on the reasoning rather than trust a score.
- The prediction enters the maintenance workflow as a scheduled intervention — not an outage report.



BUSINESS BENEFITS

01

Unplanned becomes planned — 24 hours of lead time converts an outage into a maintenance window

02

96–99% precision — engineers act on real risk, and are not worn down by false alarms

03

Cascading damage avoided — a predicted rectifier fault is a part swap, not a dead battery and a dead site

04

Predictions engineers will use — because explainability shows them why, not just what

06 Interference Prediction Agent

No fault. No failed site. Just interference - predicted before it hurts service




About the Interference Prediction Agent

Under the right atmospheric conditions, radio signals travel far beyond the cell they were meant for. Tropospheric ducting turns a distant base station into an invisible aggressor on your uplink. Throughput drops, packet loss climbs, subscribers complain — and there is nothing wrong with the site reporting the problem. Conventional fault management has nothing to find.

The Interference Prediction Agent forecasts these events **24 hours ahead, base station by base station, with a probability score** — so the network is retuned before customers notice. It has been proven at scale on roughly 8,000 base stations of live operator network data.

HOW IT WORKS

- Hourly network KPIs are collected across the estate — uplink throughput, received power, connected users, packet loss, signal quality and setup success rate.
- Historic ducting events are matched against those KPIs, so the model learns what the hours before an event look like.
- A forecasting model projects the next 24 hours of network behaviour.
- A classifier scores every base station for the likelihood of a ducting event, and surfaces the result on an operations dashboard.
- Engineering acts pre-emptively — frequency reallocation, transmit power reduction or antenna tilt adjustment.



BUSINESS BENEFITS

01

Reliability protected in advance — interference is addressed before it degrades experience, not after complaints arrive

02

Effort focused where it counts — engineering attention goes to the sites that will actually be hit

03

Spectrum and power used intelligently — rather than defensively over-provisioned against a risk that may not materialise

04

Accuracy that holds — the pipeline retrains and self-evaluates continuously as conditions shift



engineering
intelligence

ENGINEERING **THE CHANGE**