5G Slicing Demands New Tricks From Your Old OSS.

How ephemeral services, edge compute, and 5G slicing are forcing a redesign of telecom operations — proven at TMF DTW Ignite 2026.

Rapax

May 2026

Written by Shawn Ennis, Founder & CEO, Rapax

$127 Billion. Sitting on the Table.

At Mobile World Congress 2025, GSMA's Head of Networks Henry Calvert said the quiet part out loud.

> *"There is $127 billion that we could possibly be leaving on the table at the moment."*
>
> **— Henry Calvert, Head of Networks, GSMA**

He was referring to enterprise revenue from 5G network slicing — the capability that telecom built its entire 5G investment case on. Premium services. Differentiated experiences. New revenue beyond connectivity. The story we have been telling boards and investors for seven years.

The number today is a rounding error against that projection.

Take yourself back six or seven years. The 5G slicing pitches were everywhere. Stadium 5G for sold-out concerts. URLLC for connected cars. Premium tiers for healthcare, manufacturing, public safety. Every analyst chart pointed up and to the right. Every vendor pitch promised the same future. We were finally going to escape the bandwidth commodity trap. We were going to sell experience, not capacity.

Then we did the work. We funded the RAN refresh. We deployed standalone cores. We stood up Network Slice Selection Functions. We integrated Network Exposure Functions. We trained engineers. We presented roadmaps. Across the global industry, hundreds of billions of dollars went into networks that can do exactly what we promised.

The network works. Slicing works. The promise is real.

> The laggard is the OSS, not the network.

Your 5G network is ready to sell services that did not exist five years ago — slices that live for minutes, missions that exist for hours, edge sessions that vanish on completion. Your OSS is still monitoring static MPLS circuits the same way it did in 2010.

You bought the network. You wrote the checks. You sat in the boardroom and defended the capex for the next decade of services. The question every telecom CTO needs to answer before the next budget cycle is this: did you also buy the operations stack to match? Or are you running 2030 services on a 2000 OSS?

This paper is the honest assessment, the operational blueprint, and the proof. It is also the story of a TMF Catalyst at DTW Ignite 2026 — where Wavelo, CGI, BT Group, and Rapax demonstrated what slice-aware operations actually look like in practice.

The Problem: Ephemeral Services and the OSS That Can't See Them

Defining the Term

Before we go further, let me define what I mean by an "ephemeral service." The phrase shows up in conference talks and analyst reports without a clean definition, which is part of why the industry has been slow to address the operational gap.

An ephemeral service is a service that exists for minutes to hours rather than months to years. It is auto-provisioned when needed, operates for a defined window, and is auto-decommissioned when its purpose is complete. It is bound to a specific mission, event, session, or contractual time window.

Examples telecoms can sell today, given the network capability that already exists:

- A 5G slice dedicated to a stadium concert for a four-hour window with SLA-backed performance
- A drone surveillance mission patrolling a port for 18 minutes with guaranteed mission-critical telemetry
- An autonomous vehicle handoff lasting 30 seconds across a cell boundary
- A private 5G campus event running for an eight-hour shift with isolated user groups
- A live broadcast feed from a sports venue for three hours with low-latency video uplink
- An emergency response slice for first responders during an incident — duration unknown, priority absolute

All six of those examples are revenue opportunities. All six are technically possible on the networks already deployed. And all six are nearly invisible to the operations stacks running underneath them.

How We Got Here

The OSS was not designed wrong. It was designed for a different world.

Through the 2000s, the world of telecom services was static. MPLS circuits. Dedicated fiber. SD-WAN deployments. EVPN configurations. Provisioned once, monitored for years, billed monthly. The OSS world was built around three assumptions: services are permanent, topology is fixed, and the time between an event and the response can be measured in hours or days.

Then 3GPP standardized 5G slicing. NSSF, AMF, slice profiles, dynamic instantiation. Through 2020 to 2024 the networks themselves became capable of doing what the slicing promise demanded. But the OSS stack did not evolve at the same pace. Slicing pilots succeeded technically and stalled operationally. The networks could provision slices in seconds. The OSS could not monitor them at that velocity.

By 2025 the industry shifted its narrative from "when will slicing explode" to "why is slicing revenue stuck." The answer was almost always operational, not technical.

Why 5G Slicing Makes This Possible — And Urgent

If you have only heard about slicing in pitch decks, here is the short version. 3GPP network slicing lets the network dedicate guaranteed capacity, latency, and quality-of-service to a specific application, customer,

or mission. Application-level slicing layers a service contract on top of the network slice. Together, they let an operator sell a premium experience as a discrete, billable event — not a permanent subscription.

This is the technical foundation for ephemeral services. The network can now do what the OSS cannot: spin up a customized service, deliver it for a defined window, tear it down, and bill against it. The infrastructure is ready. The operational stack is the bottleneck.

What the Old OSS Can't Do

When I talk to OSS leaders inside major operators, the gap shows up in six specific ways:

Service catalogs require pre-registration.

OSS service catalogs were built to register services as permanent entities. Service ID, customer ID, configuration template, monitoring policy — all defined ahead of provisioning. A slice that exists for 18 minutes has nowhere to live in that catalog. The mission begins, the mission ends, and the catalog never knew it was there.

Inventory systems assume permanence.

A static drone, a fixed camera, a permanent fiber — these have inventory records. A mission that materializes for an hour and disappears does not. Without a record of the resource, no SLA, no billing event, no audit trail.

Correlation engines depend on static topology.

Traditional correlation says "this alarm relates to that device" because the topology between them is fixed. When the topology changes every minute — a new slice anchors to a different edge node, a drone re-associates with a different antenna, an autonomous vehicle hands off cells — correlation degenerates into noise. Every topology change looks like a new failure.

RCA runs on cycles longer than the service exists.

By the time your OSS finishes investigating why a slice degraded, the slice is gone. The customer is gone. The SLA breach is recorded. The revenue is refunded. Root cause analysis that runs on cycles of hours or days is operationally useless against services that live for minutes.

Alarm systems generate noise on every teardown.

Every slice teardown looks like a failure to a legacy alarm system. Multiply that by thousands of slices a day and your NOC drowns in false positives. Real failures get buried in routine teardown noise.

Geospatial visualization does not exist.

Your topology view is a static diagram. Drones move. Vehicles drive. Campus users walk. The OSS cannot show you where the service is right now, let alone where it is going next.

The Business Impact

These six gaps add up to four real business consequences.

Commercial: You can't sell what you can't operate.

GSMA's $127 billion is on the table because operations cannot deliver SLA-grade service for the ephemeral category. Network Product teams have services ready to sell. Sales teams have customers asking. Operations teams cannot commit to delivery. The deals stall.

Strategic: High-margin services migrate elsewhere.

Edge services are the highest-margin growth category in telecom. Every quarter the OSS gap persists, more of that market moves to providers — both inside and outside telecom — that built their operations stacks for dynamic services from the start. The longer the wait, the harder the recapture.

Operational: NOC teams overwhelmed by churn.

Every slice instantiation generates alarms. Every teardown generates more. Every drone handoff, every cell boundary crossing, every mission completion. NOC teams spend their days dismissing notifications that should have been correlated automatically. SLA breaches accumulate. Customer churn on premium services follows.

Personal: Internal blame between sell-side and operate-side.

Heads of Network Products who cannot ship. CTOs who cannot defend the slicing investment. NOC directors taking the blame for services they were never given the tools to operate. The slicing initiative becomes a political flashpoint rather than a revenue engine.

Where This Is Heading

Without intervention, the trajectory is clear. Slicing revenue continues to stagnate. Edge services migrate to adjacent providers. The OSS gap becomes the strategic constraint on telecom revenue growth. Operators that solve this in the next 18 months lead the next decade of telecom services. Operators that do not become wholesale carriers of someone else's services.

The window is not infinite. The CSPs that figured this out first are already in market.

The Solution: An OSS That Learns New Tricks

Defining the Solution

The answer is not to rip out your existing OSS. The answer is an AI-native operations layer that knows how to handle ephemeral services, sits alongside what you already have, and teaches the legacy stack new tricks.

Specifically, the OSS needs to learn four things.

Trick 1: Dynamic Service Modeling

Treat ephemeral services as first-class catalog entities. A drone mission, a stadium slice, an emergency response session — each is a service with its own lifecycle, parent relationships, SLA, and revenue event. Services auto-instantiate when needed and auto-decommission when complete. The catalog grows and shrinks in real time.

Trick 2: Cross-Domain Correlation

Network telemetry from the 5G RAN, transport, and edge compute. Application telemetry from missions, sessions, and slice health. Business context — which customer, which service contract, which SLA. All correlated in real time, in a single view, not in five separate dashboards stitched together by humans.

Trick 3: Mission-Scoped RCA

Root cause analysis that completes before the service ends. Time-bounded analysis windows. Correlation scoped to the active mission, not to the entire network history. When a slice degrades, the RCA engine produces an answer in seconds while the service is still live — in time to actually fix it.

Trick 4: Real-Time Geospatial Topology

Services move geographically. Drones fly. Vehicles drive. Slices anchor to different edge nodes as users move. The OSS topology view updates in real time as the service moves through space. Static topology diagrams are obsolete the moment the mission launches.

Why This Architecture Works

Most existing OSS vendors are bolting ephemeral support onto static architectures. They are adding service-lifecycle endpoints to inventory systems that were never designed for it. The architecture fights itself. Every new ephemeral capability comes with three new operational gotchas.

An AI-native platform — one built for streaming data, dynamic graphs, and real-time correlation from day one — handles ephemeral services natively. There is no impedance mismatch between the network's velocity and the OSS's processing model. The technology stack exists today. This is not a 2030 roadmap problem.

How It Works (Without the Jargon)

Walk with me through the conceptual model. There are five steps.

1. The network emits telemetry, as it always has. But it now also emits slice lifecycle events — instantiated, modified, terminated. These are first-class signals, not afterthoughts.
2. Edge nodes emit application telemetry — mission state, drone location, slice tenant identity, session ID. The data that used to live in silos at the application layer now flows into the operations layer alongside the network data.
3. The BSS emits service catalog events — service created, service modified, service billed, service closed. The commercial reality and the operational reality finally share the same timeline.
4. The AI-native OSS correlates everything in real time. A network event paired with a mission context produces "this slice is degrading, the affected mission is X, the customer is Y, the SLA threshold is Z, time remaining is N minutes." A geographic event paired with a service context produces "this drone moved, the slice handed off to a new edge anchor, the mission continues." A closure event paired with a revenue context produces "this mission ended, the service is closed, the billing event triggered, alarms cleared."
5. Operators see one unified view. NOC teams act on contextualized alerts. RCA completes within the service lifetime. The same engineers who were drowning in alarm storms are now running mission-aware operations.

Market Validation

This is not theoretical. The industry has been converging on this answer in parallel from multiple directions.

GSMA Intelligence has been increasingly explicit about the gap between slicing capability and slicing revenue, and the role of operational readiness in closing it. Their Mobile Economy 2026 report calls out AI as the evolution from cost-optimization tool to core revenue driver for operators. Henry Calvert's $127 billion line is not just an analyst soundbite — it is the industry's quietly acknowledged reality.

On the operator side, BT Group's commercial trajectory tells the story. In January 2025, BT ran the UK's first real-world 5G Standalone slicing deployment at the Belfast Christmas Market, dedicating a slice to mobile payment terminals in Lavery's Beer Tent. In early 2026, BT announced an extended partnership with Ericsson to deploy NSSF and NEF — Network Slice Selection Function and Network Exposure Function — on its Network Cloud. BT is planning a commercial launch of nationwide network slicing by summer 2026, targeting both enterprise and consumer segments.

Vodafone Business beat BT to the first UK SLA-backed slicing offer in early 2026. The competitive race is on. The CSPs that solve operations win it.

And at TMF DTW Ignite 2026, the entire industry conversation centers on AI-native telco operations. The Catalyst program — TM Forum's proof-of-concept showcase — features more than 55 teams demonstrating exactly the kinds of capabilities this paper describes. One of those Catalysts is ours.

Apply It: A TMF Catalyst at DTW Ignite 2026

The Partnership

Four companies have collaborated to demonstrate end-to-end slice-aware operations at TMF DTW Ignite 2026 in Copenhagen. Each owns a different part of the stack.

- **Wavelo — BSS.** Service catalog, dynamic provisioning, and billing for ephemeral services. Wavelo's modern BSS treats a 30-minute drone mission as a first-class billable service, not a footnote on a monthly invoice.
- **CGI — Systems integration.** The multi-domain orchestration backbone connecting BSS, OSS, and network. CGI's integration expertise threads the components together in a way that an operator could actually deploy.
- **BT Group — Customer profile.** The operator perspective. The use case scenarios. The standards of service quality, security, and SLA expectation that a Tier 1 CSP brings to a demonstration like this.
- **Rapax — Observability, correlation, service analytics, and visualization.** Cross-domain telemetry ingest from network, edge, application, and mission. Real-time correlation. Mission-scoped RCA. The geospatial map that makes the abstract concrete.

The Scenario: A Three-Site Drone Surveillance Network

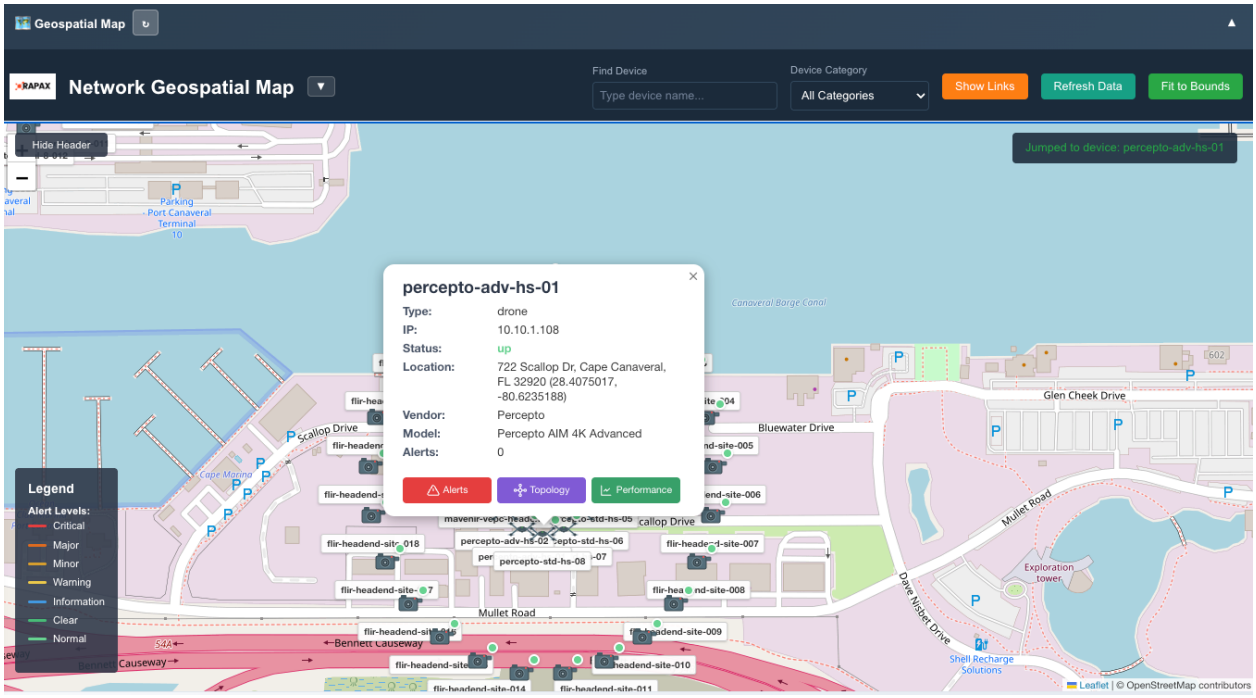

The demonstration scenario is a drone-based port surveillance network — physically grounded, visually striking, and operationally complex enough to exercise every aspect of ephemeral service management. The architecture is modeled on real port surveillance deployments and the kind of mission-critical 5G use case BT and other operators are building toward.

Three sites. 113 devices. A heterogeneous mix of vendors and technology — the messy reality of a real telecom deployment, not a sanitized lab.

- 75 FLIR fixed cameras distributed across three sites (25 per site) on GigabitEthernet
- Three Cisco Nexus switches providing site aggregation
- Three Fujitsu 5G ORAN antennas providing radio access to mobile assets
- 30 Percepto drones — 24 standard 1080p surveillance drones and 6 advanced 4K observation drones (10 per site)
- One Dell PowerEdge server and one Mavenir vEPC at the headend site

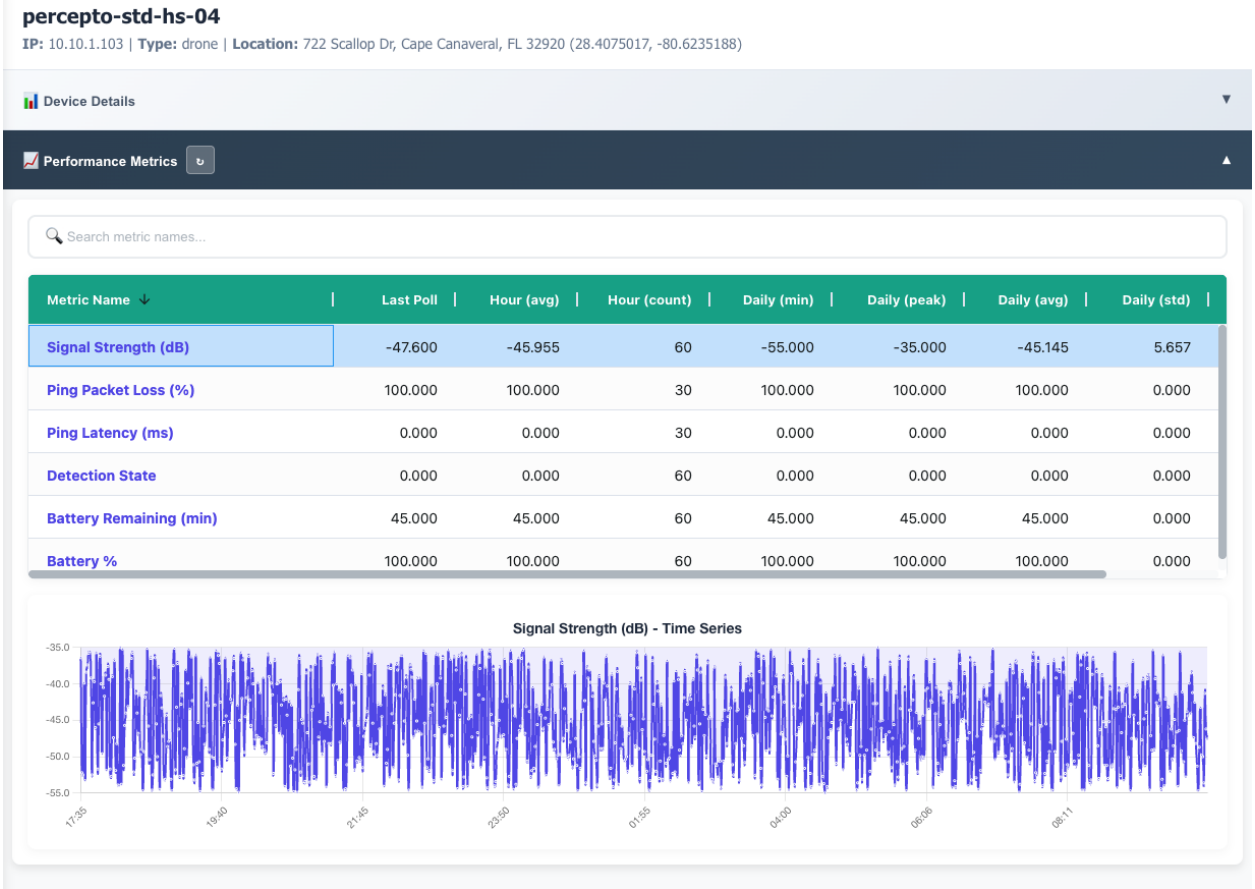

Each drone has three radio interfaces connecting to all three site antennas. The drones move through space and re-anchor to different antennas in real time. This is the dynamic topology the OSS must handle — the topology graph reshapes itself every 60 seconds as missions execute.

The scenario supports two mission types. A surveillance patrol uses a standard 1080p drone for a 15 to 20 minute mission, sweeping multi-site waypoints, auto-scheduled at sunset. An observation mission uses an advanced 4K drone for a 30 to 45 minute mission, auto-triggered by intrusion detection during a patrol. Both create billable, monitorable, SLA-backed services that exist for the duration of the mission and vanish on completion.

The Five Demonstrations

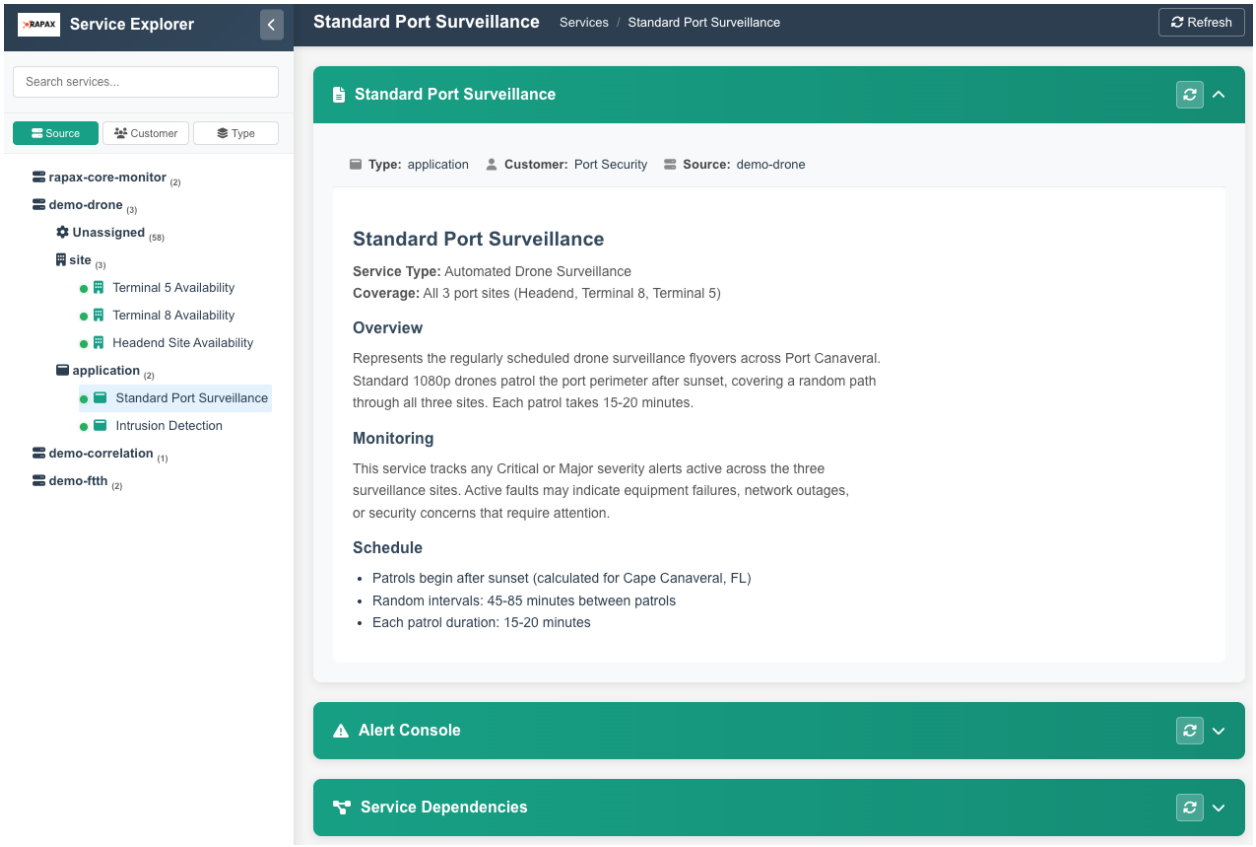

At the Mission Garage at DTW Ignite, visitors see five specific capabilities demonstrated end-to-end.

Demonstration 1: Watch the Drone Move

This is the visceral one. A visitor clicks Launch Surveillance on the demo console. A new mission service appears in the catalog. A 5G slice instantiates. A drone icon lights up on the geospatial map. Over the next 18 minutes the drone traces its waypoint path across the map in real time, updating GPS, battery percentage, and signal strength every 60 seconds. Visitors see the service exist as a moving object in space — not as a row in a static dashboard.

This is the I-can-see-it moment. The mental model of ephemeral services finally has a visual referent.

Demonstration 2: Mission Equals Billable Service

The moment the patrol launches, a service appears in the catalog with a name like "2026-06-23 19:42 Surveillance Trip." Its parent is the Standard Port Surveillance service. It has its own lifecycle, its own SLA — surveillance coverage of three sites for the duration of the mission — and a billing event tied to its successful completion. When the mission ends, the service closes. The revenue event fires. The operations stack and the BSS share the same timeline.

Demonstration 3: Auto-Observation as Closed-Loop Operations

At roughly 50% patrol progress, the simulator triggers an intrusion detection event with 2% probability. The system responds without human intervention. An advanced 4K observation drone auto-launches from

the nearest site. A new mission service spins up. A new slice instantiates. A new SLA goes into effect. The patrol drone continues its mission. The observation drone flies a direct path to the detection coordinates.

This is the closed-loop operations model: the operations layer not only sees the event, it acts on it, and creates new ephemeral services as a result. The OSS is no longer a passive observer.

Demonstration 4: Cross-Domain Root Cause Analysis

A Cisco switch port flaps. In a traditional OSS, that single physical event generates dozens of independent alarms — every camera that lost connectivity, every drone that lost backhaul, every interface that timed out, every device-down event. The NOC operator drowns.

The Rapax correlation engine sees one root cause and one consolidated impact picture: 25 cameras at Terminal 5 affected, one antenna degraded, three drones with reduced backhaul, one active surveillance mission at risk, SLA breach probability estimated at 73%. One alert. One context. One RCA. Completes in under 30 seconds — while the mission is still in progress.

Demonstration 5: The Teardown

The mission completes. The drone returns to base. The service catalog auto-closes the mission service. The slice tears down. Alarms are cleared. The topology graph reverts to its baseline. A billing event fires to the BSS. Revenue is recognized. Total elapsed time from drone-on-pad to revenue-recognized: under 60 seconds.

That last 60 seconds is the entire story. Every step of it is happening today, in software, on hardware that exists, with partners that ship product.

The Business Metrics This Demo Proves

< 5s	< 30s	< 60s
Service Instantiation	**Mission-Scoped RCA**	**Provision to Revenue**

What This Proves About the Broader Pattern

Drones at a port are a tangible, visual, easy-to-grok example. They were chosen for the demo because they are physical, mobile, and immediately understandable. But the architecture is general.

The same operational model handles a stadium 5G slice for a sold-out concert. It handles a private 5G campus for an enterprise event. It handles connected vehicle handoffs across cell sites. It handles a live broadcast slice from a sports venue. It handles an emergency response slice when a disaster strikes.

> If the OSS can operate this, it can operate any ephemeral service the network can sell. Drones are the demo. Slice-aware operations is the product.

What to Do About It

Three Questions for Your Current OSS Vendor

Before you write your next OSS check, ask three questions. The answers are diagnostic.

1. Can your service catalog instantiate a service that lasts 18 minutes? If the answer is no, you cannot sell mission-bound services.
2. Can your correlation engine scope RCA to a mission context? If the answer is no, your SLA breaches accumulate faster than you can investigate.
3. Can your topology view update in real time as services move geographically? If the answer is no, your NOC is blind to half your service catalog.

> "We're working on it" is a no.

The Six-Month Adoption Roadmap

You do not need to rip and replace. You need an overlay that teaches the existing stack the new tricks. The path is four phases over six months.

Phase 1 (Months 1-2): Observability

You cannot fix what you cannot see. Start by collecting cross-domain telemetry — network, edge, application, mission context — into a single platform. This is the foundation. Everything else depends on it. In two months you should have unified visibility into the data sources that previously lived in different silos.

Phase 2 (Months 2-4): Correlation

Layer in cross-domain correlation. Mission-scoped RCA. Geographic context. By the end of month 4 your NOC starts seeing one consolidated alert per incident instead of 50 independent alarms. The first cultural shift happens here — engineers stop spending their day dismissing notifications.

Phase 3 (Months 4-5): Service Modeling

Add ephemeral service catalog entities. Auto-instantiation, auto-decommission. The service lifecycle aligns to the mission lifecycle. The catalog becomes a living view of what is happening on the network right now, not a static inventory of what was provisioned last quarter.

Phase 4 (Month 5-6): BSS Integration

Connect the service lifecycle to billing. Revenue recognition fires on teardown. SLA enforcement runs throughout. By the end of month 6 you have production-grade slicing operations. Not three years. Not a wholesale replacement. An overlay that learns the new tricks without ripping out the old stack.

Risks and Challenges

Risk 1: Organizational change.

Network Product teams and Operations teams have to actually talk to each other. Ephemeral services require shared accountability for SLA, revenue, and operational health. The fix is joint OKRs, shared metric dashboards, and integrated service reviews. This is harder than the technology.

Risk 2: Vendor lock-in.

Some incumbent OSS vendors will resist overlays. They will push proprietary extensions and tightly coupled architectures. The fix is open standards — TMF Open APIs, 3GPP service interfaces — and clear data ownership. Your data, your platform.

Risk 3: Cultural inertia.

"We've always done it this way" is the most expensive sentence in telecom operations. The fix is to start with a single ephemeral service use case. A stadium event. A campus pilot. A drone trial. Prove the model, capture the data, then scale. Show, don't tell.

Risk 4: Talent gap.

Your existing OSS team knows MIBs, YANG models, and SNMP. They may not know streaming graphs, AI correlation, geospatial telemetry. Plan training and hiring accordingly. The shift from packet-level engineering to service-level engineering is real and the people who navigate it are the most valuable in the industry.

So What?

The 5G slicing revenue opportunity is real. GSMA's $127 billion is not a fantasy number — it is what the industry could be capturing if operations could deliver what the network can already do.

The network can deliver it. The OSS is the laggard.

> CSPs that solve the ephemeral service operations problem in the next 12 to 18 months lead the next decade of telecom services. CSPs that don't watch their high-margin edge services migrate to adjacent providers — and end up as wholesale carriers of someone else's offerings.

This is not optional. This is not a 2030 problem. The slicing investment you have already made is at risk if the OSS cannot catch up to what the network can already do. Every quarter spent waiting is a quarter of revenue not captured and customers not won.

Come See It Yourself

If you are at TMF DTW Ignite 2026 in Copenhagen, visit our Catalyst demonstration in the Mission Garage. Watch a drone fly across a real-time geospatial map. See cross-domain root cause analysis complete in

seconds. See the entire ephemeral service lifecycle — provision to revenue — happen in under 60 seconds.

If you cannot make it to Copenhagen, book a follow-up conversation with us directly. Thirty minutes is enough to walk through your specific scenarios, share the dashboard, and answer the questions this paper raises.

Email: sales@rapax.app **Book directly:** cal.com/shawn-ennis

About the Author

Shawn Ennis is the Founder and CEO of Rapax and Citus Technologies. With 25+ years in telecom operations and network assurance, Shawn previously founded Assure1 (later Federos), which Oracle acquired in 2021. He led AI-first initiatives at Concentrix before returning full-time to founding Rapax. He holds 12 patents in telecom OSS/BSS and hosts the Transformation Leaders Podcast, where he interviews the operators and operators-turned-vendors building the next generation of telecom infrastructure.

Connect: linkedin.com/in/shawnennis

About Rapax

Rapax is an AI-native service assurance and observability platform built for Tier 2 and Tier 3 telecom operators — and for the ephemeral, slice-aware operations the industry is moving toward. The platform delivers cross-domain correlation, dynamic service modeling, real-time geospatial topology, and AI-augmented operations including Frank (the AI Tier 1 support agent) and Wade (the AI knowledge engine). Rapax is deployed in production today and is engineered for the operational reality of 5G slicing and edge services.

Learn more: rapax.app

About the Catalyst Partners

Wavelo is a modern BSS/OSS platform, part of Tucows, building software-driven infrastructure for the next generation of communications service providers. CGI is a global IT and business consulting services firm with deep telecom systems integration expertise across North America, Europe, and Asia-Pacific. BT Group is the UK's largest telecommunications operator and one of the most advanced operators in the world on 5G Standalone and network slicing deployment.