


NEVER GO DARK

Predictive make-before-break for encrypted tunnels on the move.

EXECUTIVE SUMMARY

Teams on the move lose their encrypted tunnel for seconds when a link drops. Never Go Dark predicts the collapse and finishes the backup handshake **before** the link fails, leaving the crypto core untouched. It earns the best composite of every existing mechanism we emulated, ATSSS-level continuity at **42% less** standby airtime, and is confirmed on real WireGuard. It is a software-only adaptation of the documented KATIM (EDGE Group) 9001-R.

One real 5G drive in Cork: the bearer collapses at t+44 s and four mechanisms meet the same collapse. Hatched = voice dark.



When the bearer dies, the encrypted tunnel goes dark.

Every existing fix trades dark gaps for an always-on standby.

HOW IT WORKS, AND WHERE IT GOES DARK

KATIM (EDGE Group) 9001-R · rugged network encryptor for vehicles and field teams

Wi-Fi



[D] post-quantum handshake [D] up to 8 encrypted QoS planes, own keys each [D] handshake retry every 10 s [I] one cipher port: the encryptor cannot see the radio

■ Documented reactive emulated 9001-R design Recovers only after the peer is declared dead, retrying every 10 s. No standby. 2.52 s6.7 s outage / transitionstandby airtime	▲ MOBIKE-style roaming IETF RFC 4555 Moves the live session in one round trip, but only after dead-peer detection. 1.23 s43.3 s outage / transitionstandby airtime	◆ ATSSS active-standby 3GPP TS 23.501 §5.32 Lowest outage of the three, but the standby transmits all the time. 0.87 s98.6 s outage / transitionstandby airtime
--	---	--

PROBLEM STATEMENT

An in-motion encryptor that cannot see the radio can only react after the bearer fails: it waits to declare the peer dead, then rebuilds the handshake, so each handover costs seconds of dark voice and stale positions. Today's alternative, a standby that is always on, closes the gap only by transmitting all the time.

It predicts the collapse from the gateway's own telemetry, and warms the standby **before** the break.

Predict, don't react


Warms the standby from risk, seconds before the active link stalls.

Only the timing changes

Standard handshakes, fresh keys; the crypto core is untouched.

Blind to the radio, by design

Uses only the gateway's own encrypted-side telemetry.



Two predicted collapses: each time, the standby is ready before the active link stalls, so each costs one 0.1 s stall tick. At 66.4 s bearer B fails with no warning and the tunnel is dark for 2.1 s.

THE PREDICTOR

Inputs are what the gateway already measures on its encrypted side (**keepalive RTT and its trend, loss, staleness, delivered goodput, queue backlog**) and each tunnel's configured bearer class. No modem counters, no router integration. Gradient-boosted trees compiled to plain Python, an **adaptive per-link threshold**, and an early-warm gate for slow satellite handshakes. Only the **timing** of standard handshakes changes.

The prediction is blind to the radio. A pre-registered modem-hint variant **failed on unseen rail** (score -0.123, 6.0 s outage on S5), so the no-radio design shipped.

SHIPPED PREDICTOR · VALIDATION SPLIT (BY TRACE)

0.832

ROC-AUC · logistic 0.640

4.9 s


median lead, bounded by the 5 s horizon

120/132

caught (~91%) at 0.65


0.177

false-alarm tick fraction



N = 132 collapse events over 40,073 validation ticks

A small, auditable policy times standard handshakes. The crypto core never changes.



Seven real datasets, eight scenarios, scored **once** on fresh seeds the loop never saw.

7 MEASURED DATASETS · NO INTERPOLATION, EVER

● UCC 5G · 676 real switches

● UCC 4G LTE

● Lumos5G mmWave · 1,718 blockages

● Starlink on the Road

● Meireles V2I Wi-Fi

● HSR-MPTCP · simultaneous

● Berlin V2X · event truth

● GEO = measurement-parameterised model

8 SCENARIOS · HOW EACH STANDBY IS BUILT

S1

REAL

LTE↔5G walk

single real replay

S2

STITCHED

5G→LEO drive

UCC 5G + Starlink LEO, concatenated

S3

MODEL

5G→GEO drive

UCC 5G + GEO model

S4

STITCHED

Wi-Fi edge

Meireles Wi-Fi + UCC 4G, time-aligned

S5

SIMULTANEOUS

HSR rail

two carrier subflows, no stitching

S6

MODEL

collective 5G→GEO

UCC 5G + GEO model · 15 users

S7

STITCHED

Lumos walk

Lumos walk + UCC 4G

S8

STITCHED

Lumos drive

Lumos drive + Starlink LEO

HOW THE POLICY WAS FOUND, AND AUDITED

1 · LOOP 1

An LLM agent edits policy.py

The frozen evaluator scores each change on training traces; it is kept only if the score rises.

2 · INDEPENDENT AUDIT

Over-fitted constants found

Loop 1 had tuned per-scenario constants. We replaced them with an adaptive per-link threshold.

3 · LOOP 2 · FOLD-GUARDED

No change generalised

7 iterations across trace groups kept no change, so the simpler policy shipped.

Protocol

100 ms tick · zero-order hold, never interpolated · transitions anchored on real degradation events, with randomised timing · static isolation checker on every policy · split **by trace** · oracle = perfect-foresight ceiling · one run on **seeds 300–319**.

Score, frozen before the test

$$\begin{aligned} \text{cost} &= 1.0 \cdot \text{AoI p95} + 1.0 \cdot \text{outage} \\ &\quad + 0.3 \cdot \text{overhead} + 0.3 \cdot \text{airtime} - 0.2 \cdot \text{video} \\ \text{score} &= \text{cost}(\text{documented}) - \text{cost}(\text{policy}) \\ &\quad - 10 \times \text{gate shortfall} \end{aligned}$$

Aol, outage and video are scaled by the documented design, overhead and airtime by the always-warm reference, so **the documented design = 0**. The evaluator's SHA-256 (a78f791c...) is recorded in the one-shot file and unchanged.

Disclosure

The held-out traces were evaluated once earlier (seeds 100–119) by an independent auditor on an older policy. That audit led us to remove per-scenario constants. **Nothing was tuned on held-out data**. The final test ran once on fresh seeds 300–319.

Best composite of every existing mechanism we emulated, on unseen traces, scored once.

OUTAGE VS STANDBY AIRTIME · NGD LANDS IN THE EMPTY REGION



+0.373

Best composite score of every existing mechanism we emulated. Next best: MOBIKE at +0.166.

0.83 s at 56.7 s

ATSSS-level continuity (0.87 s) with 42% less standby airtime than ATSSS (98.6 s).

3x · 1.5x

Less outage than documented reactive and MOBIKE, at more standby airtime than those two (6.7 s and 43.3 s).

MECHANISM	SCORE	OUTAGE	AIRTIME	STALE P95	ALERTS	VOICE	VIDEO
■ Documented reactive emulated 9001-R design	-1.269	2.52	6.7	3.88	0.952	0.912	0.779
▲ MOBIKE-style RFC 4555	+0.166	1.23	43.3	2.26	0.986	0.962	0.803
◆ ATSSS-style TS 23.501 §5.32	+0.097	0.87	98.6	2.17	0.991	0.970	0.783
○ Always-warm naive make-before-break	+0.142	0.86	98.6	2.16	0.992	0.972	0.789
● Never Go Dark shipped policy	+0.373	0.83	56.7	2.00	0.992	0.981	0.844
⦿ Oracle ceiling, not a mechanism	+0.759	0.54	30.0	1.89	0.995	0.982	0.841

Seconds: outage per transition; airtime = transmit on the non-primary bearer per episode – a pre-warmed standby, or the active tunnel after a make-before-break roam (so MOBIKE, which pre-warms no standby, still shows airtime once it roams onto B). Rates over feasible messages: alerts on time, voice within 400 ms, video goodput. Replay reproduces outage and gate rates exactly; airtime within one 100 ms tick.

Clear wins on Wi-Fi edge and rail, one loss on the GEO model, within noise elsewhere. Always less standby airtime than ATSSS.

PER-SCENARIO HONESTY · NEVER GO DARK – ATSSS, PAIRED 95% CI

S1REAL LTE↔5G walk 0.00 s outage [0.00, 0.00] identical airtime 0.0 s	S2STITCHED 5G→LEO drive +0.07 s outage [-0.10, +0.23] within noise airtime -27.8 s	S3MODEL 5G→GEO drive +0.15 s outage [+0.05, +0.23] ATSSS better airtime -22.4 s	S4STITCHED Wi-Fi edge -0.51 s outage [-1.18, -0.07] NGD better airtime -96.8 s	S5SIMUL. HSR rail -0.89 s outage [-1.27, -0.55] NGD better airtime -53.8 s	S6MODEL collective 5G→GEO +0.02 s outage [-0.17, +0.18] within noise airtime -20.3 s	S7STITCHED Lumos walk +0.12 s outage [-0.20, +0.45] within noise airtime -74.2 s	S8STITCHED Lumos drive +0.01 s outage [-0.14, +0.16] within noise airtime -73.1 s
---	--	---	--	--	--	--	---

We ran the shipped policy, unmodified, on real WireGuard tunnels.



The interactive console and the 1:58 film are submitted alongside this document.

REPEATED 60× · 3 SETUPS × 5 RUNS

least-dark in all 3 setups · mean ± sd

Cork 5G + GEO MODEL 0.42 ± 0.24 s documented 3.60 ± 0.69 s	Cork 5G + Starlink LEO STITCHED 0.50 ± 0.07 s documented 2.75 ± 0.01 s	HSR rail REAL SIMULTANEOUS 0.00 ± 0.00 s documented 6.10 ± 0.75 s
--	--	---

ONE MEASURED RUN · A = REAL 5G · B = GEO MODEL

Documented reactive	3.90 s
voice-dark · standby 2.4 s · ≤400 ms 0.910	
MOBIKE	0.76 s
voice-dark · standby 0.0 s · ≤400 ms 0.966	
ATSSS	1.04 s
voice-dark · standby 89.5 s · ≤400 ms 0.982	
Never Go Dark	0.44 s
voice-dark · standby 63.2 s · ≤400 ms 0.992	

The shipped `policy.py` runs read-only on the real tunnels and drives real warm, handshake and cutover operations. It had the smallest dark gap of the four: **8.9×** less than reactive, **2.4×** less than ATSSS and **1.7×** less than MOBIKE, with 29% less standby airtime than ATSSS (this run; -31% over 5 runs).

single measured run · `prototype/logs/result*_final.json`

A software rollout in three phases. No new hardware, no change to the keys.

PHASE 1

Shadow

The policy runs and logs p and the actions it would take. No actuation.

Advance when the field catch rate and false-warm rate match the validation operating point: ~91% caught at a 0.177 false-alarm tick fraction.

PHASE 2

Advisory

Predicted collapses are shown to operators in the management console `I`.

Operators see the lead time and the standby's readiness; the policy still does not act on its own.

PHASE 3

Active

Per-link warm, cutover and teardown, under a per-profile airtime and emission budget.

The budget caps standby airtime per bearer, so emissions stay inside each mission profile.

WHAT SHIPS

No modem or router integration: the only bearer fact it uses is each tunnel's configured class. Keys and crypto are untouched. The policy is **about 200 lines of readable logic plus a generated decision-tree block (about 1,300 lines of plain Python)**, delivered as a policy update through the management suite.

LIMITS WE STATE

An emulation built only from public documents. GEO is a model, not a measured link. The prototype is 60 real-tunnel runs across 3 bearer setups. Always-on wins on the GEO model. The oracle ceiling (0.54 s) shows headroom for better prediction.

Disclosure

Nothing was tuned on held-out data; the final test ran once on fresh seeds 300–319.

Fit with the 9001-R


Software only. No new hardware, no new radios, no change to keys or cipher. It uses the documented QoS planes and handshake as they are.

PROPOSED FIRST STEP

Shadow mode on vehicles that already carry 9001-R units, logging the predicted risk against real handovers, with no change to their behaviour.

Continuity for teams on the move, at less airtime, with QoS guidance any encryptor can reuse.

EVIDENCE · ONE SHARED QUEUE VS THE DOCUMENTED 8 PLANES



Merge the emulated 9001-R's planes into one shared queue and voice within 400 ms falls from **0.89 to 0.42**; alerts on time fall from **0.91 to 0.73**. So keep per-class planes.

INTENDED IMPACT

- 01 Continuity for moving field teams and command nodes**
Voice and positions stay up through bearer changes: **0.83 s** outage per transition vs **2.52 s** documented.
- 02 Lower emissions and airtime**
56.7 s vs **98.6 s** of standby airtime against ATSSS, at the same continuity: a smaller emission signature.
- 03 Reusable QoS guidance**
Six measured rules, below. The headline: keep per-class planes (one queue: voice **0.89** → **0.42**).

QOS / QOE GUIDANCE · FROM THE MEASUREMENTS

- G1** Keep per-class encryption planes · one queue: voice 0.42, alerts 0.73
- G2** Never discard in-flight data under staleness · dropping it thrashed freshness to 17–20 s p95
- G3** Warm on predicted risk, not observed loss · median lead 4.9 s, 120 of 132 caught
- G4** Warm satellite standbys early · a GEO handshake takes ~1.2–1.8 s at 600 ms RTT
- G5** Report airtime next to outage · **56.7 s** vs **98.6 s** at **0.83** vs **0.87 s** outage
- G6** Hold the voice targets (TS 22.179, G.114) · within 400 ms: 0.912 documented, 0.981 Never Go Dark

HOW WE WOULD DEMONSTRATE IF SHORTLISTED

At Demo Days we would show the interactive continuity console with its **live mode** driving the real WireGuard prototype, a live run of the shipped policy, and the 1:58 film.

Who built it, what we assumed, and what you can check.

TEAM

Mohammad Al Ali	coder
Aya Elsayed	researcher · reviewer / coder
Thani Almarzooqi	coder

SUPPORTING MATERIALS

Interactive continuity console · replay of every result, plus live mode on the WireGuard prototype

The 1:58 film · the problem, one measured drive, the result

The evidence pack · result files, claims register, simulator and evaluator

All built from public documents · not affiliated with EDGE.

ASSUMPTIONS, STATED PLAINLY

- A1 GEO is a measurement-parameterised model, not a live link.
- A2 Stitched scenarios are independently measured and only time-aligned.
- A3 The KATIM internals we build on are derived from public documents and tagged documented D or inferred I.
- A4 One measured drive and the WireGuard prototype illustrate the design; the 28 held-out instances are the statistical evidence.

Disclosure

Nothing was tuned on held-out data; the final test ran once on fresh seeds 300–319.

AI DECLARATION

AI assisted tools aided in the research and design of this project.

NEVER GO DARK

ATP2026 · EDGE track · predictive make-before-break for encrypted tunnels