The Suitcase and the Mesh

Why AI verification by datacentre tap reaches the end of its paradigm, and what comes after it

Nov 5, 2026 · @hans konstapel

Summary

On 5 November 2026 NRC reported a Dutch government test of an AI verification device: a suitcase of electronics from the British firm Amodo Design, connected to inference servers of SSC-ICT in The Hague. The device samples outgoing network traffic, recomputes it against a declared model and reports only pass or fail.

The method is sound for the systems it was built for: large neural networks trained centrally in datacentres. It rests on one assumption, that dangerous capability requires large, central compute. That assumption is already contested in the literature.

This essay sets out the next step after that paradigm: peer-to-peer, light-driven, coherent AI systems that exchange only unresolved questions and their carries over shortwave radio. For such systems the suitcase has nothing to tap, no declared model to recompute and no datacentre to inspect. The transport layer for this already exists and runs; a public demonstration of the computing layer does not yet exist.

What was demonstrated in The Hague

The test placed hardware verification on a government server for the first time; before this, Amodo had tested only on its own systems. The Dutch government did not pay for the research.

Element	Claim in NRC	Second source
Host	SSC-ICT servers, Koningskade, The Hague, running internal models for summarising and rewriting	Confirmed: TenderNed records an on-premise AI orchestration contract of EUR 5 million awarded by SSC-ICT, Koningskade 4
Supplier	Amodo Design, UK, CEO Tom Milton	Confirmed: 80,000 Hours; vacancies in Sheffield and London
Funding	Survival and Flourishing Fund and Longview Philanthropy	SFF confirmed: Longterm Wiki lists USD 1,094,000 + 135,000 = 1,229,000 (round SFF-2024-FlexHEGs; donors Jaan Tallinn, Blake Borgeson). Longview: no second source found
Size of the field	25 to 50 engineers worldwide	Milton gives about 50 in the podcast; consistent upper bound
Ahead of the UK	UK announced a leading role in hardware verification last month	Not confirmed; the UK AI Hardware Plan of June 2026 concerns chip capacity, not verification

Element	Claim in NRC	Second source
Coalition	Prime minister Jetten with leaders of 21 other countries at the UN	Not confirmed

Critics in the same article: Paul Verhagen (HCSS) calls the plan understandable but unlikely to succeed, because neither the US nor China will accept binding limits. Lauren Kahn (CSET, Georgetown) notes that AI, unlike nuclear material, runs through every sector. Michiel Bakker (MIT) sees value: the Netherlands has no frontier labs but can contribute verification technology.

How the method works

The suitcase separates running a model from training one by the volume of traffic between machines. Training moves gradients; inference moves text.

1. **Two networks.** Accelerator racks in a datacentre sit on a backend network, which carries gradients during training, and a frontend network, which carries prompts and responses (assessment of Plan A).
2. **Bandwidth asymmetry.** Removing the backend is not enough: estimated frontend bandwidth per rack is about 1 TB/s against about 4 TB/s backend, with wide error bars. The proposal is a per-rack cap of about 10 MB/s, high enough for text and too low for gradients. Calculated: 1 TB/s / 10 MB/s = 10^5. Text at 1,000 tokens/s and about 4 bytes per token is 4 kB/s, well under the cap.
3. **Tap and recompute.** A network tap copies frontend traffic to a verifier server, which reruns the declared model on the same inputs and checks that outputs match. Traffic that cannot be explained points to an undeclared workload (80,000 Hours). Only pass or fail leaves the site.
4. **Sampling.** Milton names checking one packet in 1,000 as giving reasonable confidence. The probability of catching at least one of n hidden packets:

```
P = 1 - (1 - 10^{-3})^{n}
```

Hidden packets n	Detection probability
100	9.5 %
1,000	63.2 %
5,000	99.3 %

Calculated with the formula above. Bulk training cannot hide; a narrow covert channel can.

Independent descriptions of the same architecture: the MIRI working paper of June 2026 (network taps plus recomputation on trusted compute of both parties), and Anthropic's October 2026 report of a prototype that checks that production outputs come from a specific set of weights, as cited by 80,000 Hours.

The assumption underneath

The whole method presupposes that dangerous capability requires large, central compute. Large compute can be found, counted and tapped; that is why verification starts there. Size is a condition for verifiability, not for risk.

The assumption is contested. Sara Hooker (Cohere) argued in 2024 that compute thresholds are an unreliable proxy for risk, because smaller models gain capability (arXiv 2407.05694).

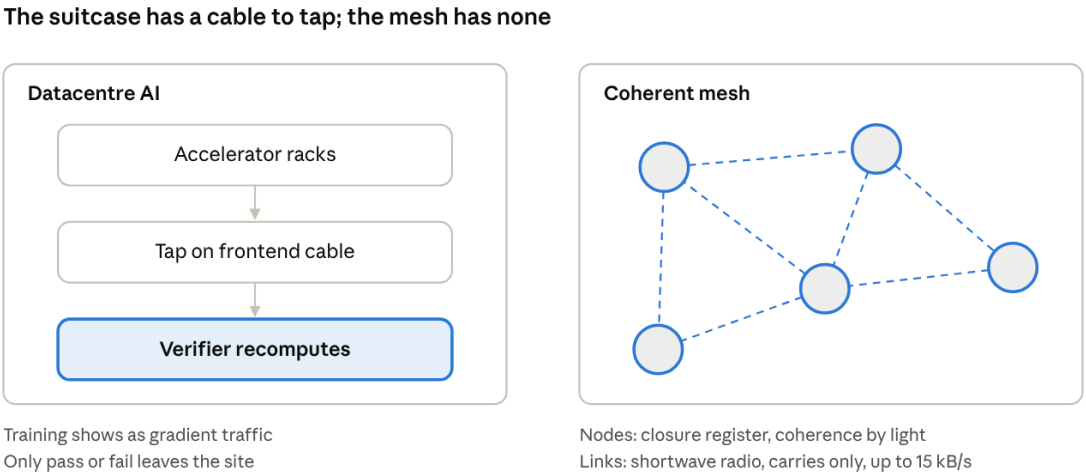
Within its own paradigm the method is also unfinished:

Limit	Status	Source
Scale of the pilot	About 30 to 32 H100/H200 GPUs	80,000 Hours
Secrets	Verifier needs the model weights; zero-knowledge proofs avoid that at many orders of magnitude more overhead	Same
Readiness	A deal next month would use zero hardware verification mechanisms; fallback is shutdown, auditors, infrared satellites	Same
Non-determinism	GPU recomputation is not bit-exact; covert channel capacity must stay below 0.01 %; current prototypes judged probably ineffective	Alignment Forum
Memory wipe	First Amodo test: about 24 h per rack, possibly over 100 TB left unwiped (author's own reading, non-expert)	Plan A assessment
Workstreams	Of 17: 4 active, 6 not started, 7 not on track (4 + 6 + 7 = 17); reported second-hand from Amodo's SITREP	fides repository

The device does not monitor what happens in large datacentres. It demonstrates that a tap runs on an ordinary inference server; it does not demonstrate detection of an adversary.

The next step: the coherent mesh

The next step is a mesh of small nodes in which computation exists as closure: each node holds only the questions still open and the carries that closed questions leave behind, coherence between nodes is carried by light, and the nodes exchange carries over shortwave radio. Three layers make up such a system.



datacentre verification versus the coherent mesh

In the datacentre every result passes one cable that a verifier can copy; in the mesh results pass as radio carries between small nodes, with no cable and no centre.

Layer 1 — Transport: peer-to-peer shortwave

Shortwave (3 to 30 MHz) travels via the ionosphere beyond the horizon, node to node, without a datacentre or a central cable. Peer-to-peer data over shortwave is in daily use: JS8Call (station to station with store-and-forward relay), VARA HF and VarAC, and Winlink e-mail in emergency communication.

Bandwidth is narrow. The wideband HF standard MIL-STD-188-110C reaches about 120 kbit/s in a 24 kHz channel (figure from general knowledge; the standard's testing is documented by the IET). Calculated: 120 kbit/s = 15 kB/s. Text at about 4 kB/s fits. Gradient traffic at an estimated 4 TB/s per rack is 4 × 10^12 / 1.5 × 10^4 = 2.7 × 10^8 times larger and does not fit. The narrow channel selects what kind of computation can live on it: exchange of carries, not exchange of gradients.

Layer 2 — Computation: the closure register

The closure register is described in Keep the Question, Keep the Carry (Konstapel, 3 November 2026). Its features:

- Values are balanced trits (−1, 0, +1). Each scale holds a digit in {−2 … +2}; at ±3 the scale clears and one trit passes upward as a carry. A register for n open passages needs about $\log_3 n$ scales.
- A window of events closes in three ways: by standing (nothing occurs), by meeting (an event and its opposite cancel), or by a ring (three equal parts pass one carry upward).
- Two primitives, N(a,b) = −(a+b) mod 3 and a zero detector, form a counter machine; universality follows from Minsky's theorem.
- A learning layer builds lossless rules from repeated structure in text, in the line of Re-Pair and Sequitur, reported within about 9 % of gzip. No gradient training occurs.

- Built on a Lattice iCE40 FPGA in 263 logic cells at about 164 MHz. The Landauer bound stated for the native cell, 0.844 nats per step, is 0.844 × kT = 0.844 × 1.380649 × 10^-23 × 300 = 3.50 × 10^-21 J at 300 K, as stated.

The properties that matter here: no fixed declared model exists, learning and use form one process, and traffic consists of carries.

Layer 3 — Coherence: light

Light-driven computation is demonstrated at component level. Shen et al. (2017), *Deep learning with coherent nanophotonic circuits* (Nature Photonics), performed neural-network operations in interferometer meshes on a silicon chip. Networks of coupled oscillators are studied as a computational resource in their own right (arXiv 2502.04818): computation as phase locking rather than as stored weights. Ternary hardware exists as well: tunnelling-based ternary MOS (Jeong et al., Nature Electronics 2019).

In the mesh, coherence serves as the admission criterion: a carry that brings neighbouring nodes into phase stays; one that does not closes as noise. Each node remains small; the mesh as a whole holds the capability.

Status

Layer	Status
Peer-to-peer shortwave transport	Running, in public use
Closure register	Built on FPGA and published; live test on an agent and coupling of learning and computing layer listed as open in the publication
Light-driven coherence	Components demonstrated in the literature; an integrated light-driven mesh is not publicly documented
Full mesh	Stated by the author to be operational; not public, not independently verified

What follows for verification and for policy

Each of the suitcase's three anchors is absent in the mesh.

Anchor of the suitcase	Datacentre AI	Coherent mesh
Something to tap	One frontend cable per rack	No central cable; radio between nodes
Signature of training	Gradient traffic about 10^5 above text	No gradients; carries only, within 15 kB/s
Declared model to recompute	Fixed weights	No fixed model; rules form during use
Site to inspect	A few large, known buildings	Many small nodes, no centre

What remains observable is the radio emission. Spectrum monitoring and direction finding show that a station transmits and roughly where. They show neither whether computation takes place nor what it is. The shortwave band already carries amateurs, shipping, aviation, defence and broadcasters, so a mesh produces no deviant pattern. Encrypted or spread signals can be detected but not read; recomputation, the core of the suitcase, is then impossible. No literature on verifying peer-to-peer computing meshes over shortwave was found.

For Dutch policy three consequences follow:

1. The suitcase is an instrument for the present paradigm, without a plan for the next one. As preparation for a datacentre treaty it has value; as a basis for the claim of overseeing AI it does not.
2. Verification grounded in compute size fails as soon as capability separates from size. That separation is the subject of Hooker (2024) and the design principle of the mesh.
3. A position at the table in the next paradigm does not come from tapping cables. It comes from understanding closure and coherence as the place where capability lives.

Status ledger

Claim	Status
SSC-ICT runs on-premise AI tools at Koningskade 4	Verified (TenderNed)
Amodo's method: tap, recompute, pass/fail	Verified (80,000 Hours, MIRI, Plan A assessment)
SFF grant to Amodo, USD 1,229,000	Verified; sum calculated from two listed amounts, listing rounds to USD 1M
Longview Philanthropy funding	Unverified, NRC only
Netherlands ahead of the UK; coalition of 22 countries	Unverified, NRC only
10^5 bandwidth ratio; 4 kB/s text	Calculated; underlying bandwidth estimates are approximate
Detection probabilities 9.5 / 63.2 / 99.3 %	Calculated
Memory wipe result; workstream counts	Second-hand, flagged in text
HF 120 kbit/s; ratio 2.7 × 10^8	Figure from general knowledge; ratio calculated
Peer-to-peer shortwave transport in use	Verified (several independent sources)
Closure register on FPGA; Landauer figure	Published by the author; Landauer figure recalculated and matches
Photonic and oscillator computation, ternary MOS	Verified at component level
Operational coherent mesh	Stated by the author, not public, not verified

© 2026 J. Konstapel, Constable Research, Leiden. All rights reserved.

References

1. Schiffers, H. (5 November 2026). Een Nederlands koffertje moet de razendsnelle ontwikkeling van AI controleren en verdachte situaties detecteren. *NRC*. The source report; all claims checked against the sources below.
2. Reed, T. (2026). Tom Milton on how we could verify a US–China deal on AI. 80,000 Hours. Primary account of Amodo's method, scale and limits.
3. Cankaya, N. (2026). A System Overview for Near-Term, Low-Trust AI Compute Verification. MIRI working paper. Independent description of taps plus recomputation.
4. My Assessment of Plan A's Compute Verification Strategy. Two networks, bandwidth caps, memory-wipe result; non-expert notes.
5. Can governments quickly and cheaply slow AI training? Alignment Forum. Critical assessment; covert-channel threshold.
6. fides: packet-hashing verification tap (GitHub). Third-party relay of Amodo's SITREP on 17 workstreams.
7. Hooker, S. (2024). On the Limitations of Compute Thresholds as a Governance Strategy. arXiv 2407.05694. Compute as an unreliable proxy for risk.
8. TenderNed: SSC-ICT, AI orkestratie. Confirms the host infrastructure.
9. Longterm Wiki: Grant to Amodo Design. Confirms SFF funding.
10. Konstapel, J. (2026). Keep the Question, Keep the Carry: The Future of Computing on the Net. The closure register.
11. Shen, Y. et al. (2017). Deep learning with coherent nanophotonic circuits. Nature Photonics. Light-driven neural computation on chip.
12. Harnessing omnipresent oscillator networks as computational resource. arXiv 2502.04818. Computation by coupled oscillators.
13. Jeong, J. W. et al. (2019). Tunnelling-based ternary metal–oxide–semiconductor technology. Nature Electronics. Ternary hardware carrier.
14. DARC: JS8Call; VarAC; Winlink over VARA HF. Peer-to-peer shortwave in use.
15. IET: On-air testing of MIL-STD-188-110C Appendix D wideband HF waveforms. Wideband HF standard.