Newsletter issue 1 - 17 Mar 2026

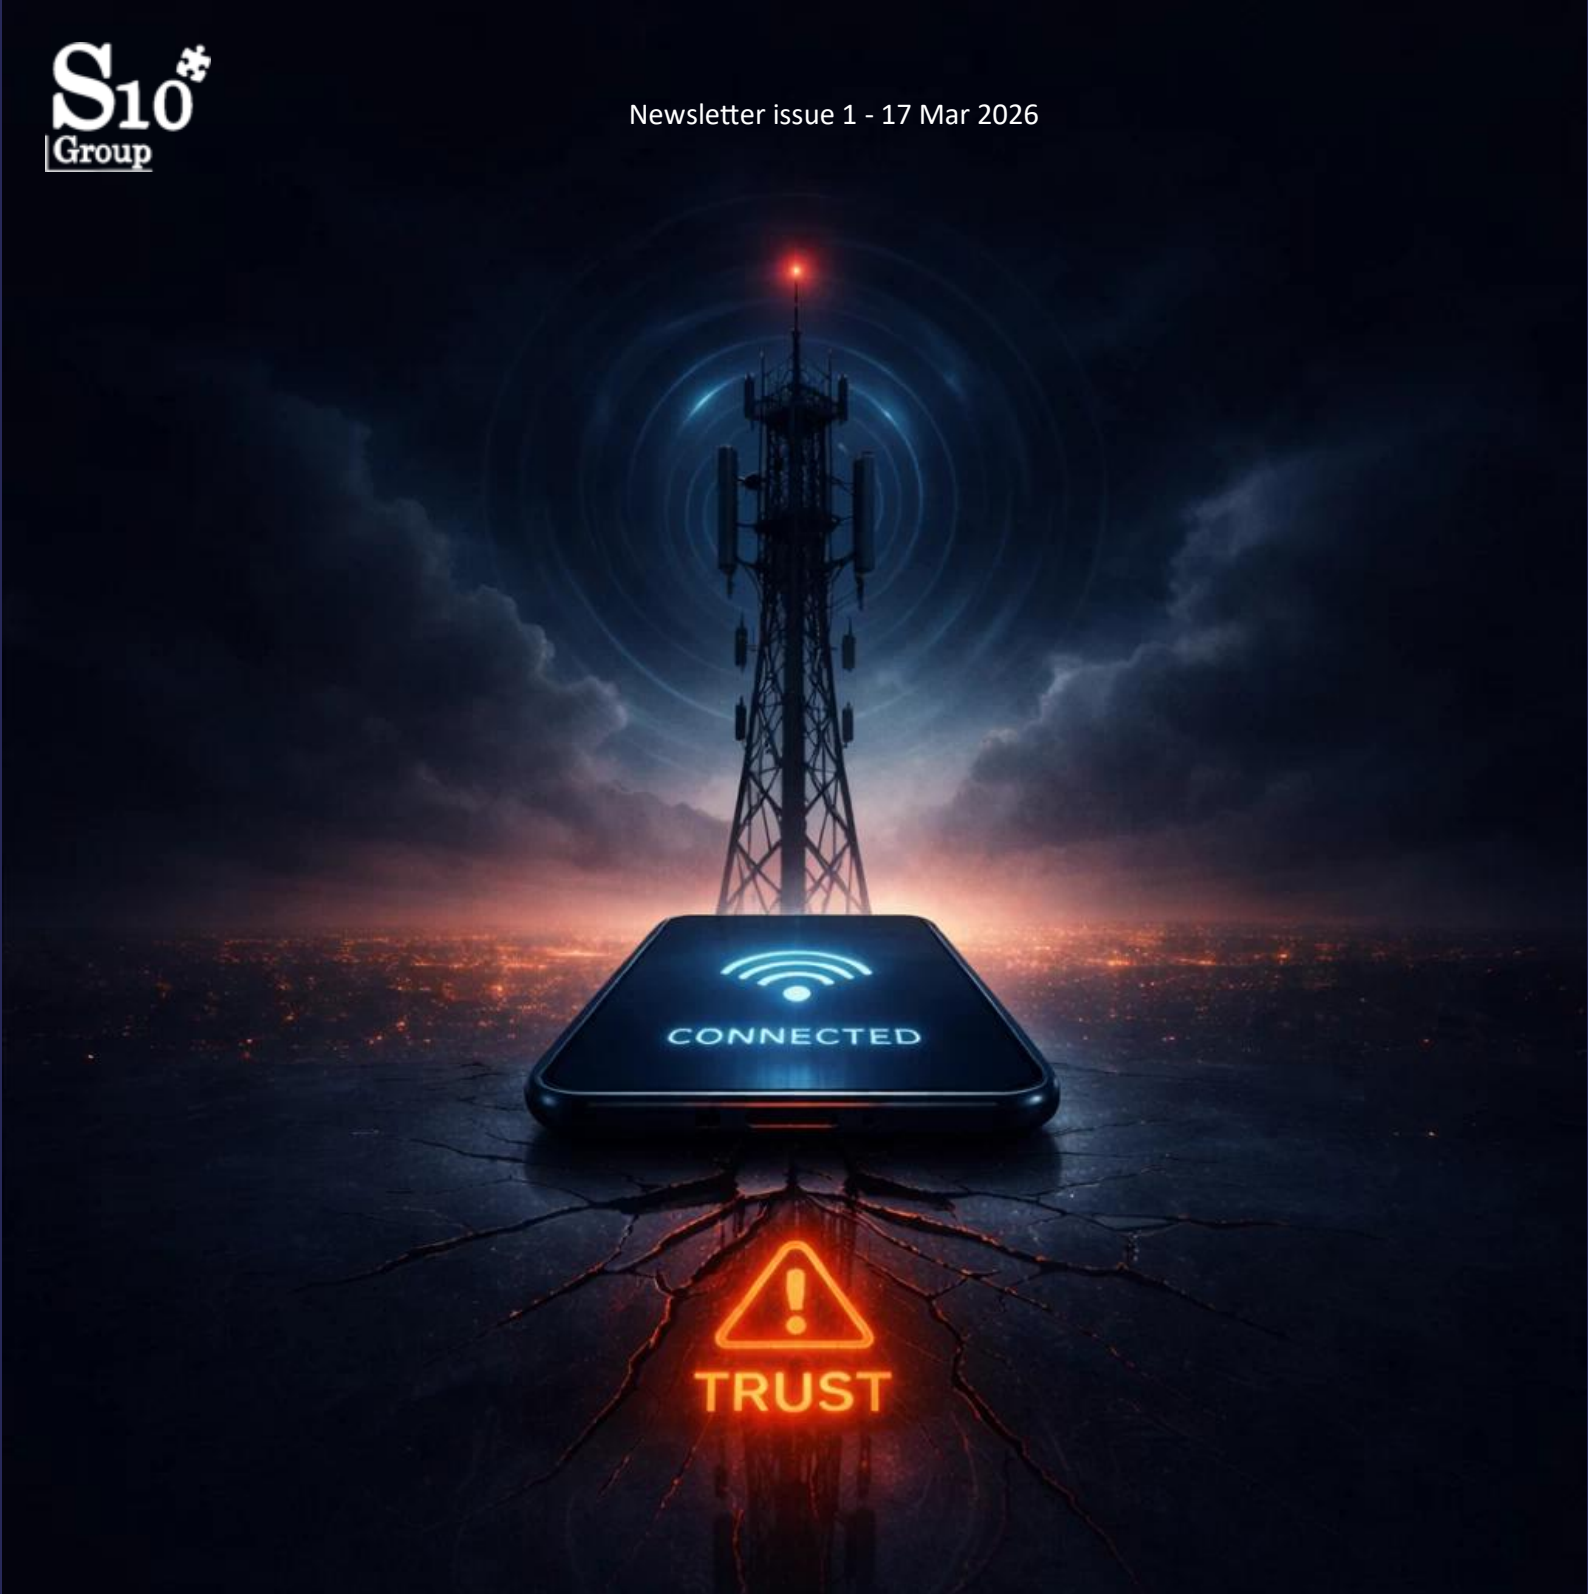

The Odido lesson:
when “access” still works, but trust doesn’t

Why identity trust is the real fault line in modern incidents.

In February 2026, Dutch telecom provider **Odido** disclosed a cyber incident affecting personal data tied to **more than six million accounts**. *(Link to Reuters).* Odido is one of the country's major telecom providers, serving millions. This incident caught my attention for the scale and because the moment data is already stolen, you're no longer preventing impact, you're managing consequences…

The Odido story isn't interesting because attackers were "clever." It's interesting because it shows how quickly a familiar pattern becomes large when warning signals exist, but decision pathways lag. *(Link to IO+)*

There's a moment in many incidents where nothing looks "down"… and yet control is already slipping.
Not because systems failed.
Because **trust** did.

What makes the reporting uncomfortable is precisely that it isn't exotic. Not a movie-plot zero-day.
More like a chain of small, plausible weaknesses that combined into something large: social manipulation of an employee, access into a customer contact system, and permissions that enabled access to a broad set of records. *(Link to IO+)*

What happened (the facts)

Odido stated the incident involved personal data from a customer contact system, and that **passwords, call details, and billing data were not involved**. *(Link to Odido)*

Public reporting describes attackers using social engineering — posing as internal IT/helpdesk to persuade staff to approve access — rather than exploiting a novel software vulnerability. *(Link to IO+)*

Several write-ups also highlight an uncomfortable detail: this general method was not hypothetical. Warnings about this kind of approach had been circulating, including in relation to widely used SaaS environments and identity/access patterns. *(Link to IO+) (Link to NOS)*

That last point matters, because it changes the leadership question from:
"How could anyone predict this?"
…to something more practical:

When warnings exist, the gap is rarely "more tools." It's whether the organisation can reduce the time between knowing and acting — turning a credible signal into fast, governed containment that limits further compromise and restores control.

A practical point: this is also where capability can genuinely change outcomes — **not by creating more alerts, but by creating control**.

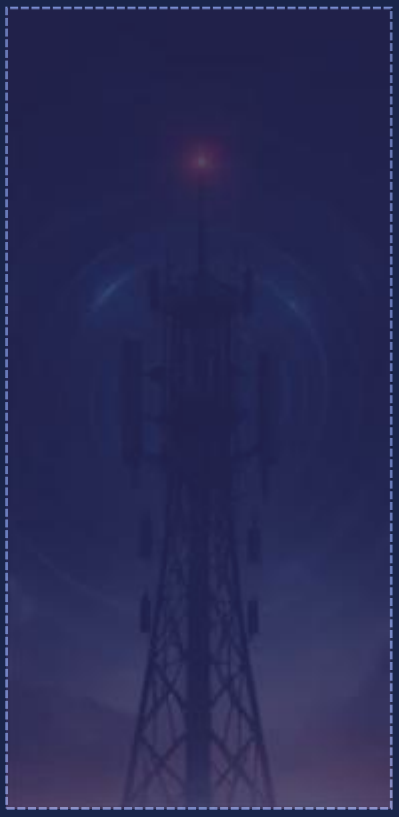

-ACCESS IS NOT LEGITIMACY

The moment services still run but trust is broken, the problem is already governance — not only recovery.

When you can reliably **see lateral movement** and **detect data-exfiltration behaviour** early enough, you can trigger containment actions that limit further spread and additional exposure.

The difference isn't "visibility for reporting." It's **visibility that enables a governed move**: reduce privileged pathways, isolate risky segments, and keep essential operations running while trust is rebuilt.

The decision boundary that mattered

Most organisations will read this and think: "So… train people better."

Training matters. But notice what that framing does.
It quietly places the centre of gravity on the individual who took the call.

A more useful lesson is structural: a secure environment shouldn't rely on a single person being impossible to manipulate — and the way access is designed and governed matters. *(Link to IO+)*

This is the boundary leaders tend to underestimate:

We treat identity controls as "security".
But in practice, identity controls are **governance** — because they define who can act, what they can reach, and how quickly we can reduce risk without freezing operations.

In other words: the incident doesn't start when data leaves.
It starts when leaders realise, they no longer know what "legitimate access" means.

And once legitimacy is uncertain, every response step becomes governance:

- How broadly can we restrict access without breaking the business?
- Who has the authority to do it fast?
- What do we keep running while we regain confidence?

This is the part that doesn't show up in many incident decks:

Detection buys awareness. Only pre-decided containment buys time.

-BOARD QUESTION

Do we know who decides when "still operational" is no longer "safe enough to trust"?

The move from system failure to failure

A system failure is visible: outage, latency, broken process.

A trust failure is quieter: the system is still running, but you can't prove who is driving.

That is why "we detected it" is not the same as "we controlled it".
And why "we restored services" is not the same as "we regained governability".

Odido customers could still use services, according to the company's own communication. *(Link to Odido)*
But the incident still carried weight because identity and access pathways touched sensitive data at scale. *(Link to Reuters)*

Three board questions to ask next time

These aren't technical questions. They force clarity **before** the next uncomfortable hour arrives.

1. **When identity trust is uncertain, what is our "minimum operational mode"?**
   What must keep running, and what can be intentionally degraded — by design — to protect control?

2. **Who can trigger rapid access restriction, and what evidence is "enough" to act?**
   Not "who should be consulted", but who has authority to pull the first lever when the picture is incomplete.

3. **Where do we have concentration risk in access — and do we know it?**
   If one compromised pathway can touch "too much," is that an accepted architectural choice… or an accidental one?

If you want to make this practical, the next step isn't more policy, it's a simple pressure-test of the first-hour decisions.

One practical pressure-test suggestion

Don't start by writing a new policy.

Run a 60-minute identity-trust exercise with three timed phases:

T+0 to T+10: credible signal, incomplete evidence
What do we do immediately that is reversible? Who authorises it?

T+10 to T+30: suspicion of persistence
How do we invalidate active trust (sessions/tokens), and what breaks operationally?

T+30 to T+60: stabilisation
How do we progressively restore controlled access without re-opening the same pathways?

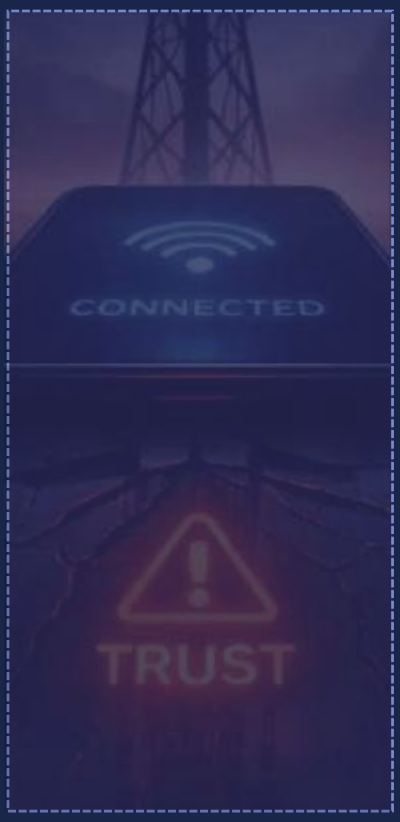

-PRESSURE POINT

The most dangerous phase is often not outage.
It is the moment normal access continues while leadership can no longer be sure what that access means.

The value isn't the tabletop discussion.
The value is discovering whether your organisation can execute the first move without debate, confusion, or operational shock.

A closing thought

So, I'll leave you with a question I find more useful than "How do we prevent every breach?":

If identity trust broke tomorrow, would your first hour be governed… or negotiated?

Resources for CIOs/CISOs to make this practical

- *Identity under pressure — 7 board questions (Trust vs Authority)*
- Three pressure-test briefings: *(Pressure test for leadership control / First-hour containment decisions / Why pressure testing matters)*

I can also share how I pressure-test this in practice through a **resilience assessment** in your own sandbox environment, with your existing security stack enabled.

If any of these would be useful, feel free to contact me.

sgemert@s10group.com

-WHAT CONTAINMENT CHANGES

Containment gives leadership a way to reduce exposure before uncertainty turns into wider consequence.