


The request that looked legitimate

Trusted channels, and what happens when a process is followed correctly and the request itself is fraudulent.


The control verified the channel. Nobody verified the request.



The gap, seen through a trusted channel

Part 1 examined time inside a system, and Part 2 examined data held outside one. Part 3 examines something subtler: an institution releasing data voluntarily, through a legitimate process, to someone who should never have received it.

No system is breached in such a case, and no payment is made, so both measured controls remain silent by design.



Revolut, and requests from a government domain


In September 2026, Revolut confirmed that an unauthorised third party had used a legitimate government agency's email domain to submit fraudulent requests for information, and that customer data had been released in response.¹ The company described it as a sophisticated external impersonation scam, blocked the address on detection, and alerted the agency, law enforcement and regulators.

According to customer notifications shared publicly, the exposed information included the identity document and verification selfie each customer submitted at registration, together with account statements and transaction histories.¹ Revolut stated that its systems and customer funds were unaffected, and that a limited number of individuals were impacted.

Revolut has not said how the government domain came to be used, and nothing here should be read as suggesting how it was. Parties claiming responsibility have made ransom demands; those are claims.¹

Every step was the right step

A LAWFUL-ACCESS REQUEST, STEP BY STEP



Each control verified what it was designed to verify; the only question that mattered sat outside all four. The pattern is not new: in 2022, forged emergency data requests sent from compromised law-enforcement accounts were reported to have obtained customer data from major technology companies.²

Controls of this kind verify the channel through which a request arrives, on the reasonable assumption that a trusted channel implies a trusted sender. Once that assumption fails, each subsequent step performs correctly and the outcome is still a disclosure.

Nothing to reverse, nothing to count

A disclosure of this kind involves no intrusion to detect and no transaction to decline. The funds stay where they are, which is precisely what the institution is able to report, and the loss, such as it is, is borne later and elsewhere, by customers whose identity documents now circulate.

WHERE AI ENTERS

Automation has not invented impersonation. What it has done is reduce the cost of producing requests that are convincing in tone, format and timing, and of producing many of them. A control that relies on a request looking right becomes weaker as looking right becomes cheaper.

That is the sense in which AI accelerates this gap rather than creating it: the underlying weakness is trust placed in a channel, and the economics of exploiting it have simply improved.

Questions worth asking this quarter

- 1 How are requests from authorities, regulators and law enforcement verified independently of the channel through which they arrive?
- 2 Who is authorised to release client data in response, and is any release above a defined scope subject to a second approval?
- 3 Would we notice an unusual pattern across such requests (volume, timing, the profile of clients named) before rather than after the data had left?
- 4 Which of our other trusted channels (partner APIs, service accounts, internal requests) rest on the same assumption that a legitimate source implies a legitimate request?

NEXT IN THE SERIES

Part 4 — Identity at the source. What happens to onboarding checks once the data they rely on is itself exposed.

Sources, and how each is used

Company, regulator and industry-body publications are cited directly wherever they exist. Press reports are used where they are the only public record, chiefly for client notices not otherwise published and for claims made by threat actors, and are marked as such.

- 1 **PRESS** Revolut's confirmation as given to TechCrunch, 12 September 2026, and Bloomberg, 12 September 2026; The Register, "Revolut falls for fake government requests, hands over customer data", 14 September 2026. Revolut has not published a standalone statement. Details of exposed data are drawn from customer notifications shared publicly and reported by The Register. Ransom demands are unverified claims.
- 2 **PRESS** KrebsOnSecurity, "Hackers Gaining Power of Subpoena Via Fake 'Emergency Data Requests'", March 2022, and related reporting at the time.