They Were Already Inside. Nothing They Did Looked Like an Attack.

A national power grid operator, responsible for electricity distribution across a vast region and operating under stringent government regulation, faced sustained interest from advanced persistent threat groups focused on its SCADA and industrial control systems. The objective of those groups was not theft. It was **the ability to interrupt power distribution at a moment of someone else's choosing.**

We were engaged, alongside a NERC-CIP-certified partner, to build the capability to find them. Not to buy a product and declare the problem solved — to establish intrusion detection that understands industrial protocols, an architecture in which a compromised office network cannot reach a protection relay, and a threat hunting practice that goes looking rather than waiting.

A US bank reading this should not skip to the next document. In its Spring 2026 `Semiannual Risk Perspective`, the OCC named foreign state-sponsored actors as a continuing threat to the federal banking system. CISA, NSA and the FBI assess that Chinese state-linked actors are positioning inside IT networks specifically to enable lateral movement into the operational systems that run critical infrastructure — and that the financial sector sits downstream of the energy and communications sectors they have already reached. **The adversary defending against here is the adversary your regulator is warning you about.**

The operator is not named. Critical infrastructure engagements are not publicised, and the details that would identify the client are the same details that would assist an attacker. The method transfers. The method is what follows.

WHAT WE BUILT

01

ICS-aware intrusion detection

Detection that parses industrial protocols rather than counting packets. A tool that only understands IP headers is blind to the thing that matters in an OT network.

02

Segmented network architecture

Critical operational systems isolated from everything else, so that a compromise of the enterprise network is an incident rather than a catastrophe.

03

Threat hunting playbook

A repeatable practice of going to look for an adversary who has not yet done anything detectable, on the assumption that they are already present.

04

Behavioural baselining

An understanding of what normal looks like on this specific grid, because "unusual" has no meaning without it and every grid is unusual in its own way.

05

Deployment without disruption

Every change made against a system that cannot be taken offline, cannot be rebooted casually, and cannot be restored from a snapshot.

+

Transfer to the operator's own teams

The playbook and the baselines were built to be operated by the client after we left. Capability that departs with the consultant is not capability.

WHY THIS ADVERSARY IS DIFFERENT

Ordinary security assumes an attacker who wants something and will act to get it. Ransomware encrypts. Fraud moves money. Both are noisy, because both require the adversary to do something you would notice. **A pre-positioning actor does none of that. Their objective is to be present, capable, and unnoticed for as long as possible.**

This inverts almost every assumption in a conventional security programme. There is no malware to find, because they bring no tools — they use yours. They log in with valid credentials through approved remote access and run native system binaries that your own administrators run every day. This is what the industry calls living off the land, and it is not a clever trick at the margin. It is the entire tradecraft.

The time horizon compounds it. These actors have been found resident in critical infrastructure networks for months and, in documented cases, for years — not because the defenders were negligent, but because there was nothing to alert on. Every individual action was authorised. Every account was legitimate. The only signal available was that the pattern of authorised activity did not match what that network had ever done before, and no product ships knowing what your network has ever done before.

THE ASSUMPTION THAT FAILS

That an intrusion produces something recognisable. Signature detection has nothing to match when the adversary's only tool is your own administrator's console.

WHAT REPLACES IT

Behavioural baselining plus deliberate hunting. You cannot detect the unknown. You can detect the deviation from known — if you know.

THE ARCHITECTURE

Assume the enterprise network is compromised and build so that the assumption is survivable. Segmentation is what makes the assumption bearable.

WORKSTREAM 01

Intrusion Detection That Understands the Grid

An intrusion detection system that does not parse industrial protocols is, in an OT network, an expensive way to count packets. The entire meaning of the traffic is in the protocol payload, and that payload is where the grid's physical reality lives.

WHAT WE ACTUALLY DID

We deployed detection that decodes industrial control protocols to the level of the command — not merely that a connection occurred between two addresses, but what instruction was issued, to which device, from which engineering station, at what time, and whether that combination had ever occurred before. In an OT network the packet is not the unit of meaning. The command is.

All collection was passive. Detection was fed by network taps and span ports that cannot transmit, so the monitoring apparatus was incapable of interfering with the process it monitored. This is not a preference. In an environment where a protection relay operates on a timescale measured in milliseconds and a mistimed instruction can damage physical plant, an active scanner is itself a threat to availability.

THE CHALLENGES WE ENCOUNTERED

- **Configuring IDS tools to accurately analyse ICS traffic.** This was the hardest technical problem in the engagement. Industrial protocols are implemented with vendor-specific variation, undocumented extensions and behaviour that diverges from the published specification. A parser that works on the standard fails on the deployment.
- **The baseline did not exist.** Anomaly detection presupposes a model of normal. Nobody had a model of normal for this grid, and the people who understood the process best had it as intuition rather than as data.
- **Legitimate operations look exactly like attacks.** Maintenance windows, seasonal load patterns, emergency manual interventions and vendor support sessions all produce traffic that a naive model flags immediately — and each false positive spent credibility we could not replenish.
- **Encrypted and proprietary traffic.** Portions of the network carried traffic that could not be decoded at all, creating blind spots that no amount of tuning would resolve.

HOW WE RESOLVED THEM

- Parsers were validated against the client's actual traffic rather than against the protocol specification, and every divergence between the two was documented as an artefact of that deployment. What the standard says is a starting hypothesis, not a description of the network.
- The baseline was built from observation across a period long enough to capture the operational cycle, and it was built `with` the control engineers rather than presented to them. They knew which of the anomalies were the grid being the grid; no model discovers that unaided.
- Known operational patterns were modelled explicitly rather than suppressed. A maintenance window is not an exception to be silenced — it is expected behaviour with a defined shape, and activity that resembles a maintenance window outside a maintenance window is precisely the signal.
- Blind spots were documented as blind spots and treated as a known limitation with compensating controls, rather than papered over. An honest gap that is written down can be defended around. A gap nobody has recorded is simply an unguarded path.

HOW WE PROVED IT WORKED

Detection was validated against replayed traffic containing known adversary behaviour patterns, and against operational traffic to establish the false positive rate under real conditions. The test that mattered was not whether the system caught obvious malicious activity — it was whether it distinguished an authorised engineer performing an unusual action from an adversary performing that same action with the same credentials. **Where it could not, we said so, and compensated architecturally rather than claiming a detection capability we did not have.**

> **RISK AVOIDED**
>
> A detection layer that produces confident green status while an adversary operates freely inside it — the worst possible outcome, because it converts an unknown risk into a falsely known one and stops anyone from looking further.

> **WHAT A US BANK SHOULD TAKE FROM THIS**
>
> Your payment switch, your card processing environment and your privileged administration hosts are operational technology in every sense that matters. Ask whether your monitoring understands what a transaction instruction means, or only that a connection occurred.

WORKSTREAM 02

Segmentation as Blast Radius Engineering

Segmentation is usually sold as prevention. It is not. It is damage limitation, and it is the only control on this list that still works after the adversary is already inside — which, against this actor, is the condition you should assume you are in.

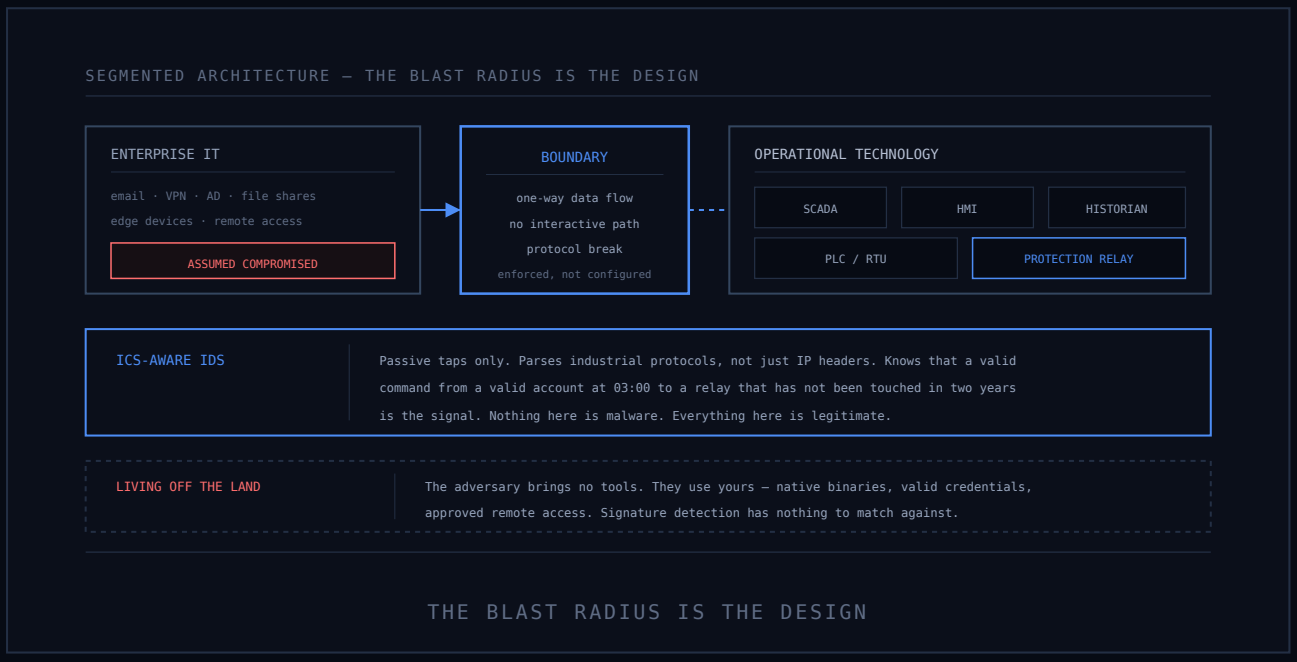

WHAT WE ACTUALLY DID

We established an architecture in which critical operational systems are isolated from the enterprise network, with data crossing the boundary in one direction and no interactive path crossing it at all. An operator can see what the grid is doing from the corporate side. Nobody can reach in from the corporate side and change what the grid is doing.

The design premise was explicit and uncomfortable: **assume the enterprise network is already compromised.** Not as a rhetorical device — as a design input. Email, VPN concentrators, the directory, edge appliances and remote access are exactly where this actor is documented to establish itself, and the architecture must remain sound with all of them in hostile hands.

THE CHALLENGES WE ENCOUNTERED

- **Resistance to change within the operational teams.** The grid's existing infrastructure was complex, and the people running it had earned their caution. Every proposed change was a change to a system where an error does not produce a stack trace — it produces an outage, and possibly a damaged asset.
- **The real topology was not the documented topology.** As with every long-lived operational estate, connections existed that no diagram recorded, placed there over years by people solving real problems under time pressure.
- **Legitimate cross-boundary needs.** Vendor support, historian replication, and reporting all required data to move from OT to IT. Every one of those flows was a potential path back.
- **Minimal disruption was an absolute constraint.** The grid cannot be taken offline for security work. There is no maintenance window in which the region agrees to go without electricity while a firewall rule is tested.

HOW WE RESOLVED THEM

- The operational teams' resistance was treated as expertise rather than obstruction, because that is what it was. They knew which systems would not tolerate what, and the design was better for having been argued with. A security programme that overrides operations in a physical-consequence environment is not a security programme; it is a new hazard.
- The true topology was discovered through passive observation rather than trusted from documentation, and the undocumented connections we found were, predictably, the ones that mattered most.
- Cross-boundary flows were reduced to the minimum that operations genuinely required, and each survivor was made one-directional with a protocol break, so that the data crosses but a session cannot.
- Changes were staged incrementally with a tested rollback at every step, and each was validated in a representative environment before it touched production. Nothing was deployed on the assumption that it would work.

HOW WE PROVED IT WORKED

The boundary was tested by attempting to cross it. Reachability from enterprise segments to operational systems was probed systematically, including through the legitimate flows, on the assumption that a determined adversary would use exactly those paths rather than looking for an exotic one. Rollback procedures were exercised rather than documented, because an untested rollback is a hope. **The measure of the architecture is not that the rules are correct in the configuration. It is that the path does not exist when someone actively looks for it.**

RISK AVOIDED

A phished credential on the corporate network becoming a hand on the grid. This is the documented path: compromise IT, move laterally, reach OT, wait. Segmentation is the control that breaks the chain in the middle, after every preventive control has already failed.

WHAT A US BANK SHOULD TAKE FROM THIS

Ask what an attacker with a valid credential on your corporate network can reach in your payments environment. If the honest answer requires a caveat about flat networks and legacy exceptions, you have the same architecture the grid had before this engagement.

WORKSTREAM 03

Threat Hunting: Going to Look

Detection waits. Against an adversary whose entire strategy is to do nothing detectable until it matters, waiting is indistinguishable from not looking.

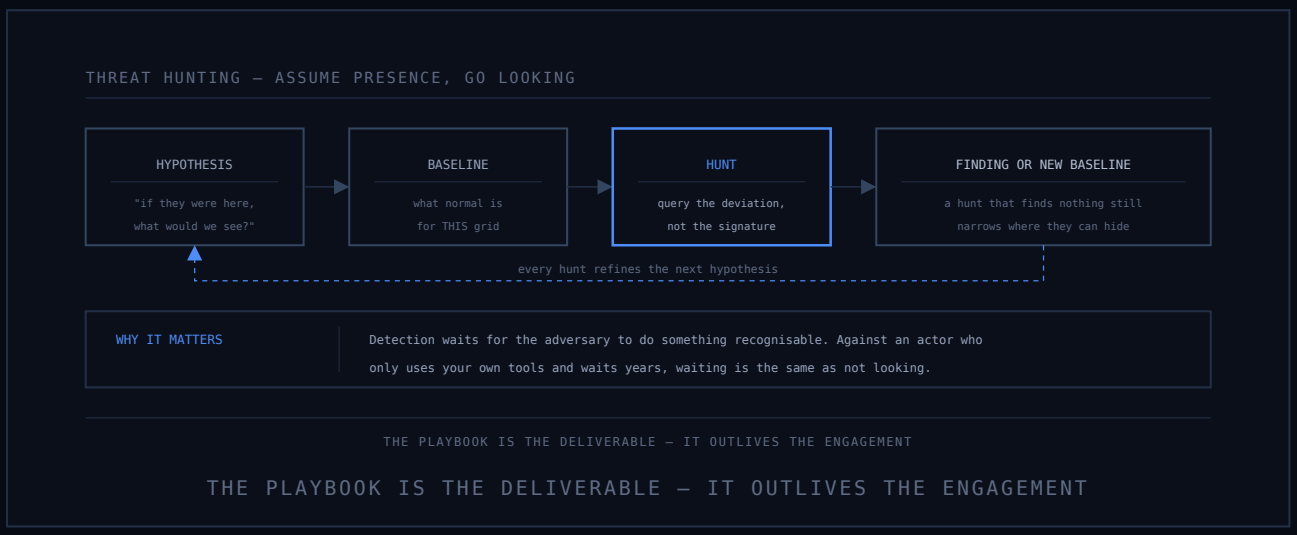

THE PLAYBOOK IS THE DELIVERABLE — IT OUTLIVES THE ENGAGEMENT

WHAT WE ACTUALLY DID

We developed a threat hunting playbook aligned to industry practice and to the specific tradecraft of the actors targeting this sector. A hunt begins with a hypothesis — if this adversary were resident in this network, what would be observably different? — and proceeds by querying for that difference against the baseline. It does not wait for an alert. There will not be one.

The playbook was the deliverable, deliberately. A hunt we ran is a point-in-time result with a short shelf life. A playbook the operator's own team can run every month, refining the hypotheses as the threat evolves, is a capability. We were engaged to leave the second thing behind.

THE CHALLENGES WE ENCOUNTERED

- **Hunting requires knowing normal, and normal was undocumented.** The same gap that undermined anomaly detection undermined hunting, for the same reason and with the same cost.
- **A hunt that finds nothing feels like a wasted week.** This is the single most common reason hunting programmes are quietly defunded, and it rests on a misunderstanding of what a negative result means.
- **Hypotheses go stale.** Tradecraft evolves. A playbook written against last year's techniques tests for behaviour the adversary has already abandoned, while providing the comfortable feeling of having looked.
- **The scarce resource is people, not tooling.** Hunting requires someone who understands both the adversary and this specific grid, and that combination cannot be bought as a product.

HOW WE RESOLVED THEM

- Baselining was made a prerequisite of hunting rather than a parallel effort, because a hunt without a baseline is not a hunt — it is a search for something suspicious-looking, which finds whatever the searcher already expected.
- Negative results were recorded as findings in their own right. A completed hunt that finds nothing has narrowed the space in which an adversary could be hiding, and that narrowing is the product. The playbook records it as such so that the value is visible to whoever funds the next quarter.
- The playbook was structured for revision, with hypotheses tied to current threat intelligence and an explicit expectation that they would be retired. A playbook that has not changed in a year is not stable; it is abandoned.
- Hunts were run jointly with the client's own engineers rather than delivered to them, so that the capability transferred rather than the report. This is slower and it is the only version that survives our departure.

HOW WE PROVED IT WORKED

The playbook was validated by having the client's own team execute it without us. If a hunt could only be run by the people who wrote it, it was not a playbook — it was a consultant's notes with a cover page. Hypotheses were tested against replayed traffic containing known adversary tradecraft to confirm that a hunt would surface the behaviour it claimed to surface, rather than reassuring everyone by finding nothing in a network where there was nothing to find.

> **RISK AVOIDED**
>
> A hunting programme that exists on paper, runs quarterly, finds nothing, and is understood by everyone involved to be a compliance exercise — while the adversary it was designed to find sits in the network it never actually examined.

WHAT A US BANK SHOULD TAKE FROM THIS

Ask when your team last hunted for an adversary that had triggered no alerts, and what hypothesis they tested. If hunting only happens after an alert, it is incident response wearing a more fashionable name.

WHY A US BANK SHOULD CARE ABOUT A POWER GRID

Banks are designated critical infrastructure. The actors pre-positioning in the energy sector are the actors your regulator has named in writing, using the same tradecraft, for the same strategic purpose. **The grid is simply where they have been found most often, because the grid is where people have been looking.**



The mapping is close to exact. A bank's payment switch is a real-time system whose availability is the product. Its privileged administration hosts are engineering workstations by another name. Its core ledger is the historian — the authoritative record everything else reconciles against. And the objective translates directly: an adversary who can stop settlement at a chosen moment has achieved, against a bank, precisely what an adversary who can open a breaker has achieved against a grid.

YOUR CONSTRAINT	HOW THIS ENGAGEMENT ADDRESSES IT	MAPS TO
State-sponsored actors	Detection and hunting designed against an adversary who brings no malware, uses valid credentials, and waits — rather than against commodity threats.	OCC SRP SPRING 2026
Operational resilience	Segmentation engineered so that compromise of the enterprise network does not reach the systems that must keep running.	OCC OP RESILIENCE
Ransomware recovery	The same architecture that contains a nation-state actor contains an encryptor. Blast radius is indifferent to motive.	FFIEC
Third-party access	Vendor support paths reduced to one-directional flows with a protocol break, rather than trusted on the strength of a contract.	INTERAGENCY TPRM 2023
Dwell time	Hunting as a standing practice against an adversary who generates no alerts, rather than detection that waits for one.	CISA GUIDANCE
Regulated environment	Delivered alongside a NERC-CIP-certified partner into a stringently regulated operator, without an outage.	CRITICAL INFRA

WHAT THE CLIENT DID THAT MADE IT WORK

It would be dishonest to present this as a result we delivered to a passive recipient. The operator gave us access to critical systems, supplied detailed documentation of the existing infrastructure, and put their control engineers in the room throughout. Regular review sessions kept both teams aligned as the design met the reality of the plant.

This matters as a transferable lesson rather than as an acknowledgement. **An OT security programme delivered to an operations team fails. One delivered with them works.** The engineers who resisted the early proposals were the reason the final architecture survived contact with the grid — they knew which systems would not tolerate what, and no external team discovers that in time to matter. A bank running this programme against its payments environment should expect the same dynamic and should plan for it as a design input rather than treating it as a stakeholder problem to be managed.

Assume presence

Every control in this engagement follows from one premise: the adversary is already inside, they look exactly like your own administrators, and nothing they have done so far was detectable. Design from there and the programme holds. Design from prevention and you will build something that reports green until the day it matters. The operator is not named — critical infrastructure engagements are not publicised, and the identifying details are the ones that would help an attacker. The method is ours, and the method is what we are offering to walk your team through.

THE QUESTIONS WORTH ASKING YOUR OWN TEAM

- What can an attacker holding a valid corporate credential reach in our payments environment? Not what should they reach — what can they, today?
- Does our monitoring understand what our transaction traffic means, or only that connections occurred between addresses?
- Do we have a documented baseline of normal for our critical systems? If not, what exactly is our anomaly detection comparing against?
- When did we last hunt for an adversary who had triggered no alerts, and what hypothesis did we test?
- Our vendor support paths — are they one-directional with a protocol break, or interactive sessions we trust because of a contract?
- If we assumed our corporate network were already compromised, which of our current controls would still be doing anything?

WHERE THIS GOES NEXT

The uncomfortable part of this work is that success is invisible. There is no incident to point at, no breach averted that anyone can measure, and the programme's value is entirely in the outage that did not happen on a date nobody knows.

Which is why the artefacts matter: the baselines, the playbook, the tested boundary. They are what you have when someone asks what you actually did, and they are what a regulator will ask to see.

TECH4BIZ SOLUTIONS | CODENUK | AUDITTRAX LABS

tech4bizsolutions.com codenuk.com auditraxlabs.com

contact@tech4biz.io