// WHITE PAPER · PUBLISHED 22 SEPTEMBER 2026

REEF, Local-First Network Defence: The Porch Camera's Heartbeat

A reader beside the gateway that turns encrypted traffic into a sentence a household can act on, and waits for a PIN before it changes the network

Jason Newell · MAX Research Collective · Report WPR-MAX-019 · Published 22 September 2026
Licensed under the Creative Commons Attribution 4.0 International License (CC BY 4.0), https://creativecommons.org/licenses/by/4.0/
Original: https://www.maxresearchcollective.com/articles/reef-local-first-network-defence.html
Cite as: Newell, J. (2026). REEF, Local-First Network Defence: The Porch Camera's Heartbeat. MAX Research Collective, WPR-MAX-019. https://www.maxresearchcollective.com/articles/reef-local-first-network-defence.html

A robot heartbeat, not you

The card on the screen read: "The porch camera has been quietly calling one outside address every 58 seconds for two hours, across 125 connections, with two percent jitter. That is a robot heartbeat, not you."

Two hours at one call every 58 seconds comes to about 124 calls, which squares with the 125 connections the card counted, and two percent of 58 seconds is a little over one second. The calls were wandering by about the length of a blink and still landing on the same beat. A machine on a timer does exactly that, and a person scrolling a feed never does.

Beneath the sentence REEF had laid out the evidence that produced it, the number of independent sensors that agreed, four things that could be done about it, what each would change on the network, and a written promise that each one could be undone. The buttons came last.

Every packet the camera sent was encrypted. REEF reached its finding without opening one of them.

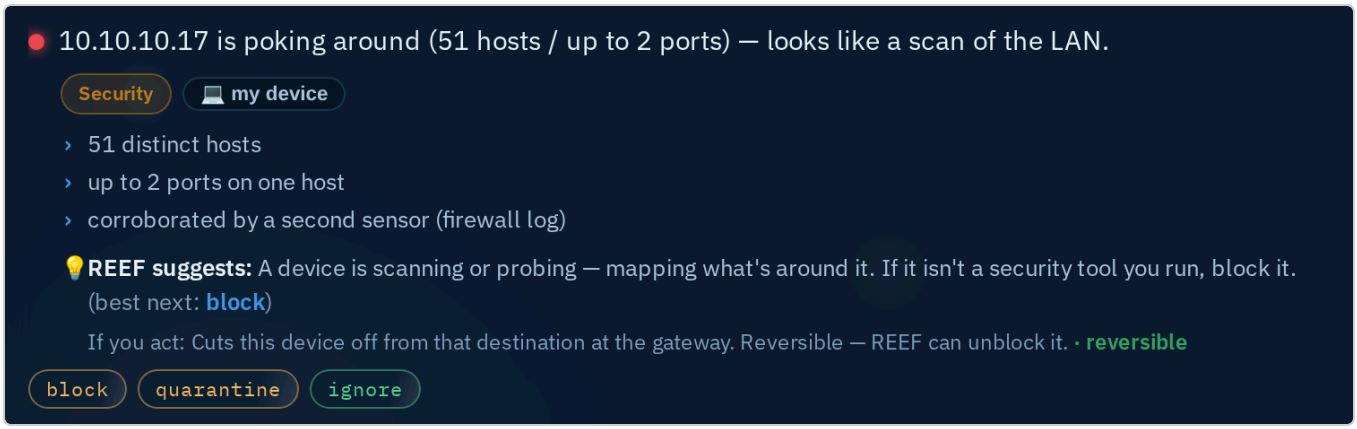


Figure 1. A finding on REEF's board: a device probing 51 hosts, the evidence and the second sensor that agreed, REEF's suggestion and its effect, then the buttons.

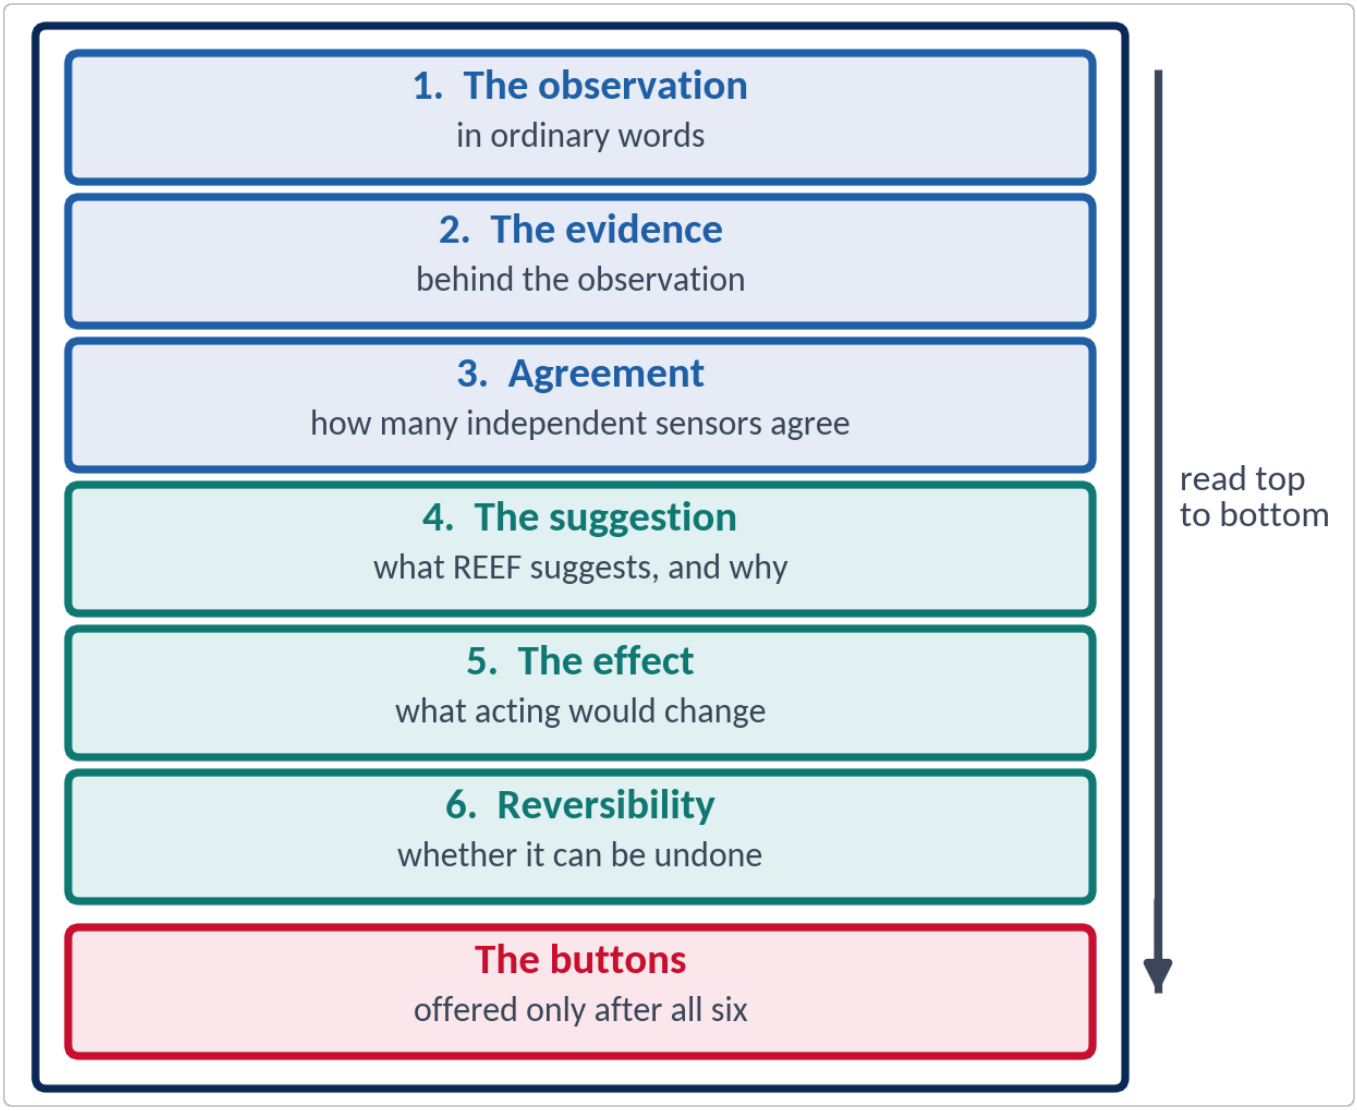


Figure 2. The six parts every finding carries, in order, with the buttons after all of them.

What REEF is

REEF sits beside the gateway, the box where a household's network meets the internet. It reads. REEF stays off the path the traffic takes: encryption ends where it always ended, every payload reaches its device unopened, and pulling the plug on REEF leaves the network exactly as it was. Anything that watches a perimeter should have that property before it has any other.

What it reads is the metadata the gateway already produces. The firewall logs every connection, and its deep-packet-inspection engine (the part of the firewall that classifies traffic by more than its address and port) records what it made of each one. From those records REEF knows who talked to whom, when, how often, how much, over which protocol and against what reputation. All of that survives encryption. Encryption seals the letter; the envelope stays in plain view.

REEF commands the inspection engine and takes its verdicts as one input among several, so there's no second engine to maintain. It serves one household. It listens on a loopback address — one that answers only the machine REEF runs on — so reaching it from outside means already being inside. Its findings come from a live window of recent traffic; there's no archive to manage, no retention tier and no query language to learn.

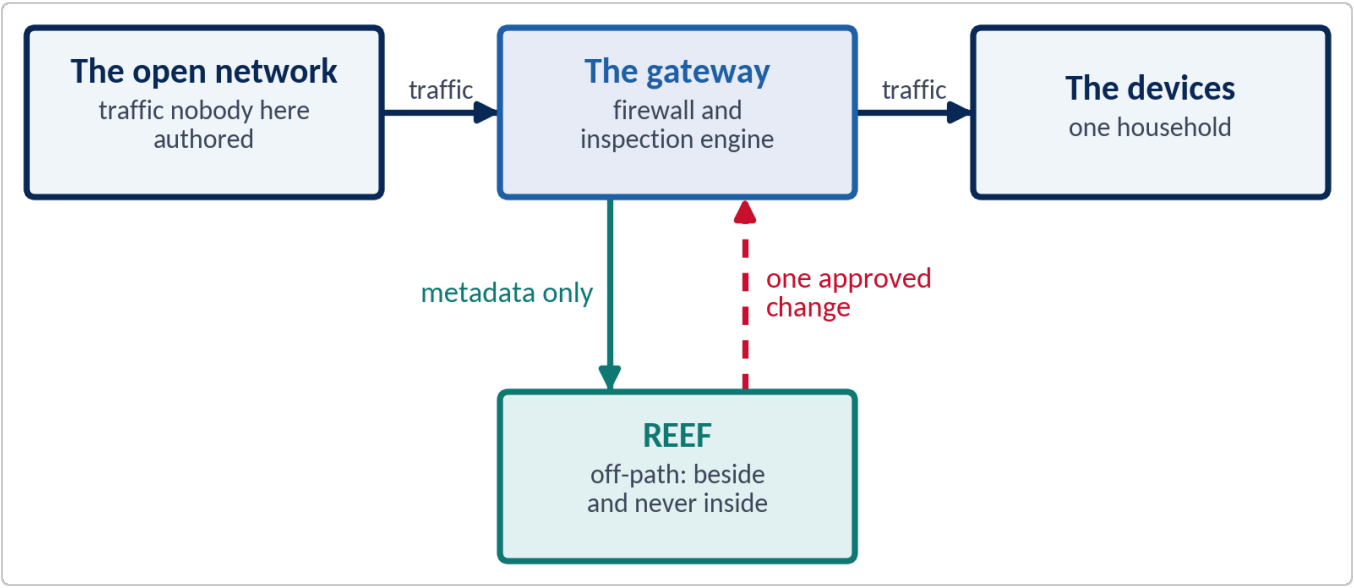


Figure 3. Where REEF sits: metadata travels up the solid line, and one approved change travels down the dashed one.

Why it exists

Start with what a modern intruder looks like from the gateway. Take a modular remote-access trojan, malware that hands a stranger control of a device, whose command channel rides streaming remote procedure calls over a modern web protocol. Remote procedure calls let one program run functions on another across a network, and this channel carries them inside TLS, the encryption behind the padlock in a browser's address bar. Its authors shape the whole thing to look like an ordinary backend service. There's no readable payload to match against a signature and nothing odd about the protocol, so signature matching finds nothing and reports that nothing with complete confidence.

Meanwhile the household already owns a firewall that knows the shape of every conversation on the network. Nobody in the house can ask it anything in plain words. The logs exist. They're written for somebody who reads packet-filter output for a living.

REEF closes that distance. It turns what the gateway already knows into a sentence a person can act on, and then it waits for that person to decide. Legibility is half the purpose and restraint is the other half, and nearly every design decision below serves one of the two.

Five questions that survive encryption

Traffic has a shape, and the shape survives encryption. REEF puts five questions to every conversation on the network without decrypting a byte.

Table 1. Five questions that survive encryption, and why each one is awkward to defeat.

QUESTION	WHAT IT LOOKS FOR	WHY MALWARE STRUGGLES TO HIDE IT
Is this rhythmic?	Regular inter-arrival intervals across many connections	An operator has to be called back on a schedule, and a metronome stays a metronome even with jitter added
Is this a first contact?	A protocol, port or endpoint this device has never used	New infrastructure is new, and a device's own history is the baseline
Does the transport look like what it claims?	Modern protocols on odd ports, a handshake fingerprint that disagrees with the stated application	Imitation has to be exact on every axis at once
Is the destination trustworthy?	Reputation, domain age, network ownership, algorithmically generated names	Infrastructure has a history, and fresh infrastructure has none
Is the flow shaped like remote control?	Directionality, long-lived idle streams, a low and slow trickle both ways	The shape of a session follows what the session is for

Any single yes is weak. Plenty of honest devices check in on a timer (a thermostat reporting the temperature, say), and rhythm alone would bury the household in false alarms. The work lies in correlating the answers per conversation, meaning per pair of one device and one destination, so that a channel unremarkable on every single axis lights up the moment several weak signals stack on the same pair.

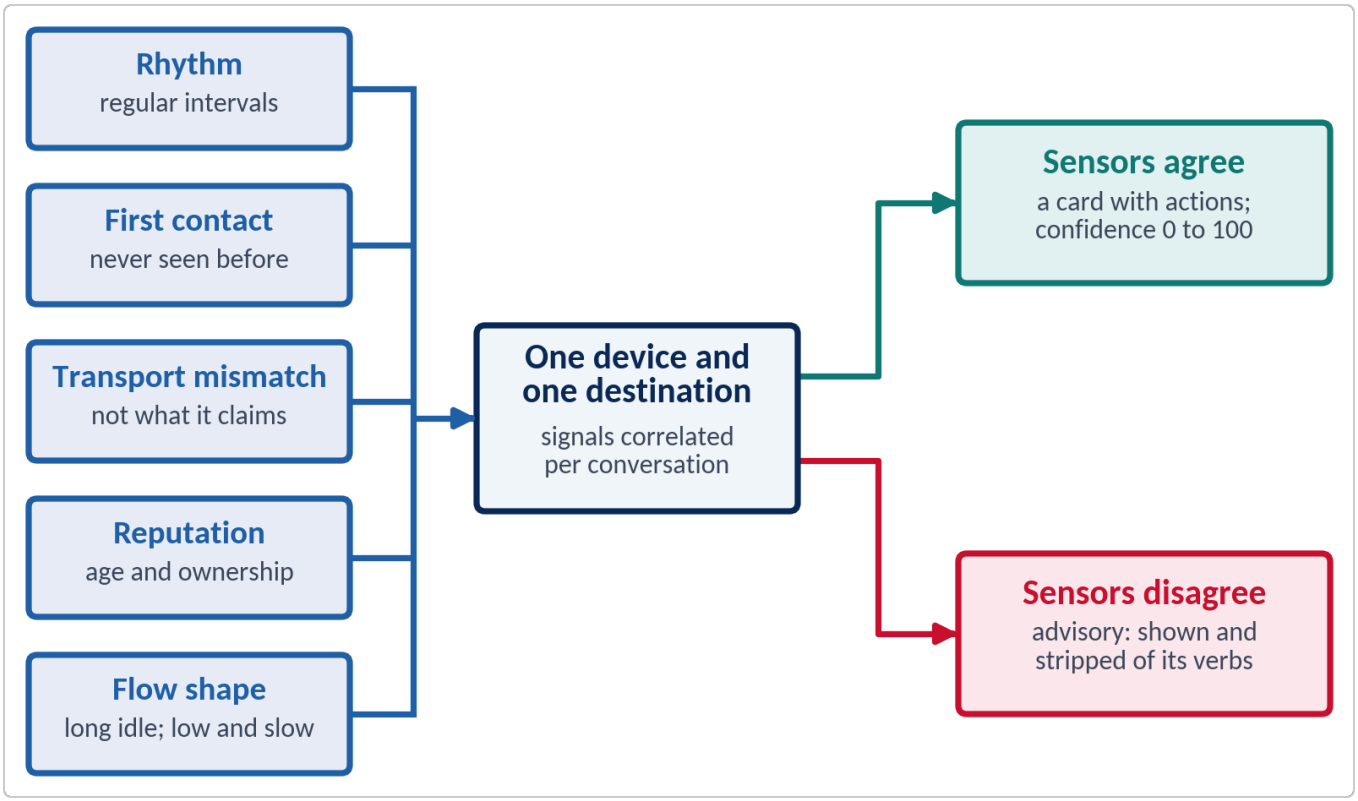


Figure 4. Weak signals stacking on one device-and-destination pair, and the two outcomes available to the result.

Corroboration turns an opinion into an offer. It counts only between independent signals; two readings of one feed are a single signal wearing two hats. When sensors agree, the finding becomes a card carrying actions and a confidence score from 0 to 100. When they disagree, it's demoted to advisory. An advisory stays on screen with all its evidence and loses the verbs that would change the network, because a thin case still deserves reading and a suppressed case is a hidden one. It just can't become a block.

Several sensors, kept apart

REEF reads from several flow sources: the inspection engine's connection index, the firewall's own packet-filter log, the intrusion-detection alert log and, where one exists, the query log of the household's DNS resolver (the service that turns names into addresses). Pouring all of them into one stream looks like a simplification, and it would break the most useful detector REEF has: the rhythm detector, the one that caught the porch camera.

Rhythm analysis measures the gap between arrivals. Interleave two sensors' views of one conversation and a phantom arrival appears beside every real one; the timing the detector reads is gone. So each flow detector runs once per sensor. Agreement across sensors is what counts as corroboration, and the household gets wider coverage without any of it contaminating the measurement.

Reputation lookups follow the same discipline for a different reason: checking a destination against outside threat-intelligence services costs money and runs into rate limits. Each sweep therefore enriches only the few most severe destinations, from cache first. Attempts are what the budget counts. Count successes instead and a failing provider becomes unbounded spend, since a failed lookup never uses up any of the budget. A hostile verdict raises a finding's severity by one step, which is the whole of its power; it can sharpen a finding and never create one.

The translation layer is the product

A correct finding still has one job left: it has to reach a person who can decide what to do about it.

Every REEF finding carries the same six things in the same order: the observation in ordinary words, the evidence behind it, how many independent sensors agree, what REEF suggests and why, what acting would change, and whether it can be reversed. The buttons come after all six. That order is an argument in its own right — an action offered before its explanation is a request for trust, and REEF is built to earn the trust first.

One acceptance test governs every screen. Would a person who isn't a network engineer see what is happening, and understand it well enough to decide? A screen that fails is unfinished, however correct its numbers are. That test is why the porch camera card says "a robot heartbeat" where the detector underneath had computed a coefficient of variation (the statistic for how tightly a set of intervals clusters).

That honesty extends to what REEF claims about the devices themselves. A device type is a guess. REEF labels it as one, sharpens it with what the device actually talks to, and leaves it editable, because one vendor ships many kinds of thing. A

randomised hardware address — the made-up identifier many phones now present to each network they join — is flagged as made up rather than adopted as an identity. And every translated view sits one click from the raw feed it came from, with its own clock showing how fresh that feed is; a translation nobody can check against its source would be one more opaque layer.

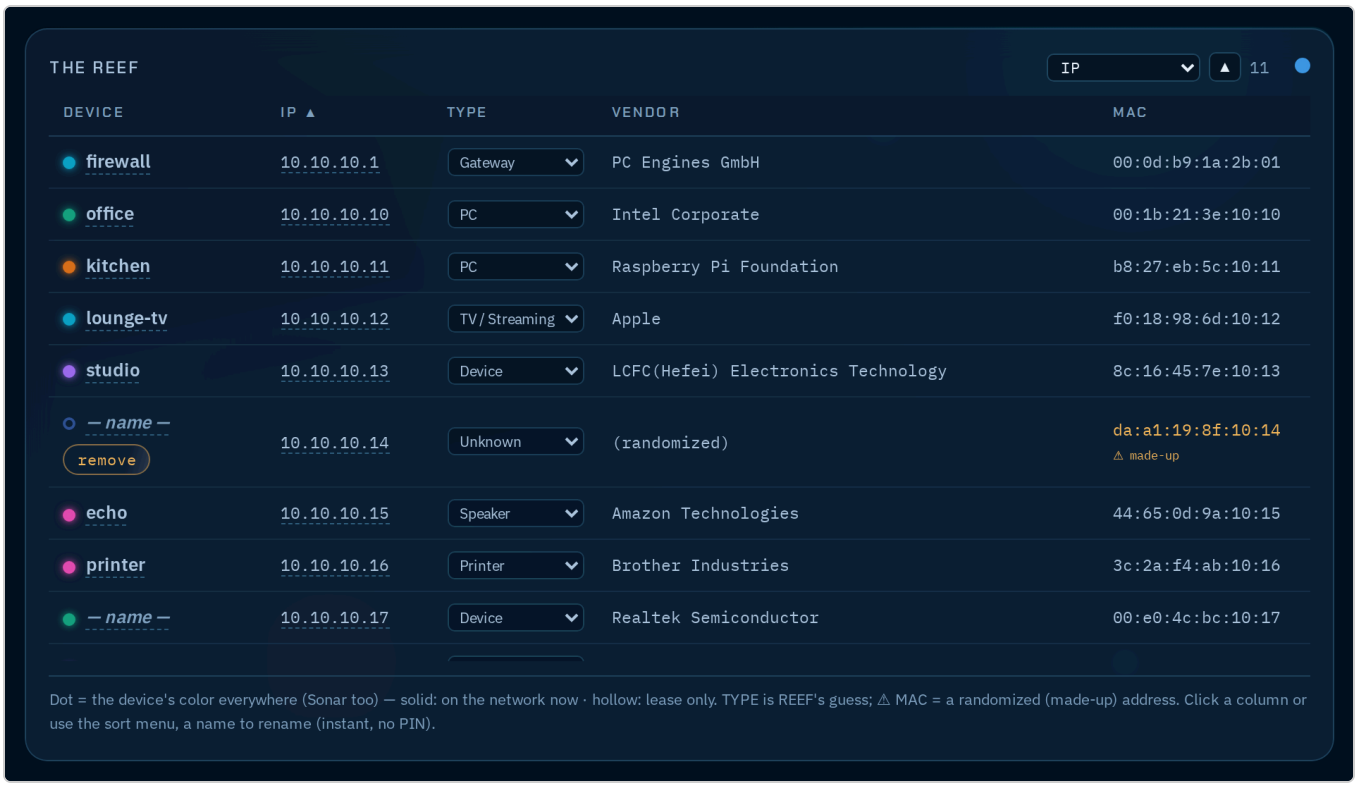


Figure 5. The device view, with a type menu on every row and a footer that calls each type REEF's guess. Names and addresses come from a synthetic demonstration network.

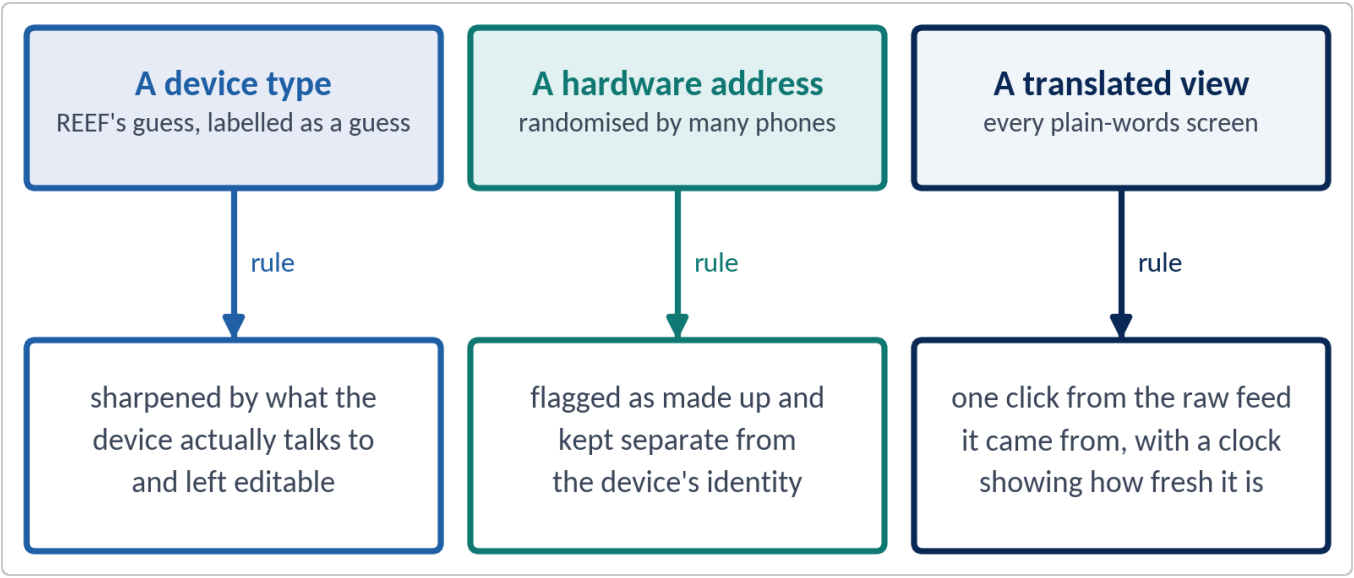


Figure 6. Three rules for what REEF claims about a device, and where each claim can be checked.

Nothing changes without a PIN

REEF does change real things on the network. Every change runs through one loop, and there is no second path.

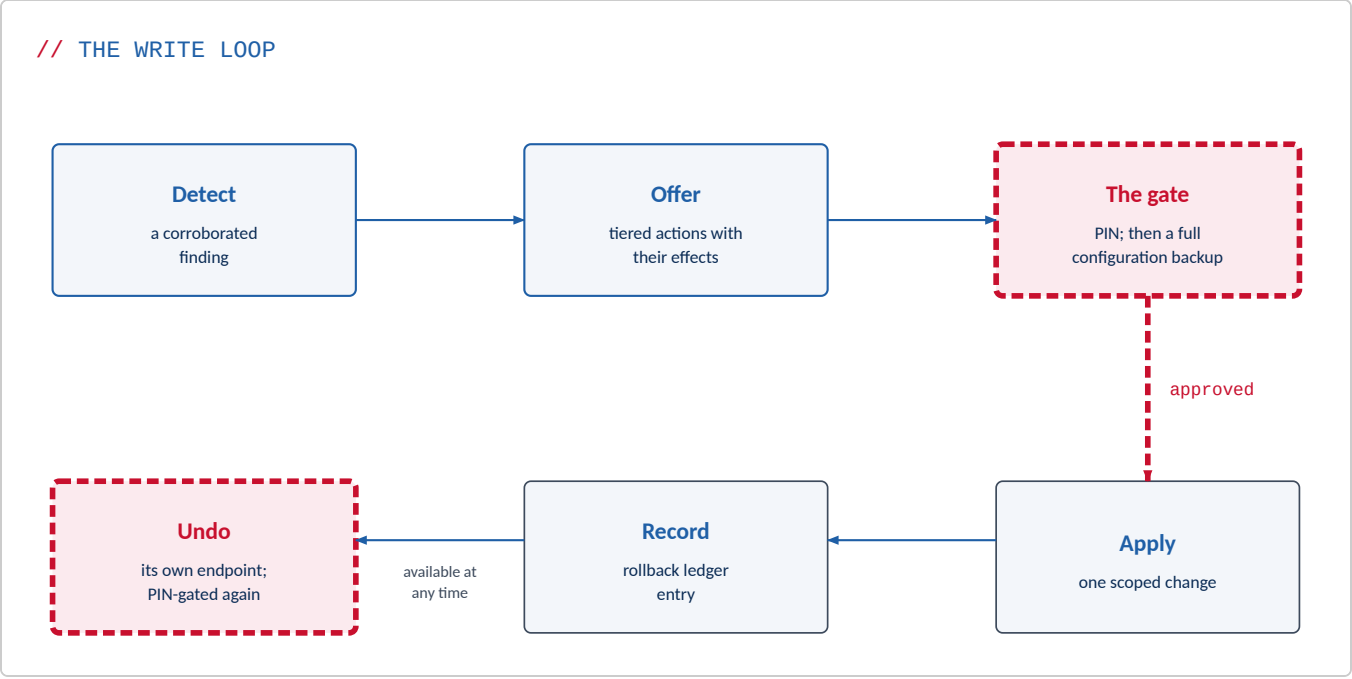


Figure 7. The write loop. The gate in the middle is what makes the rest of it safe.

Every action REEF offers carries a tier, and the tier decides the ceremony.

Table 2. Three tiers by blast radius. The ceremony is the point.

TIER	CEREMONY	EXAMPLES
Green	Instant, no PIN, local and non-destructive	Name a device, trust it, allowlist a finding, ignore one, take a snapshot
Amber	PIN, plus a full configuration backup taken first	Block, quarantine, throttle, restrict a class of traffic
Red	An explicit warning before it can be chosen at all	Anything that would open the perimeter

Amber changes take a PIN and a full backup of the gateway's configuration before anything is touched. Undoing one has its own endpoint behind its own PIN, which makes reversing a change exactly as deliberate as making it. Autonomous enforcement ships switched off. Turning it on is an explicit, logged decision the owner makes rule by rule, with a rollback ledger behind it. No upgrade ever switches it on, since a capability that arrives switched on is a capability nobody chose.

The asymmetry is the part that matters for security. On its own initiative REEF only ever tightens the perimeter: it will add a block. Opening inbound connections from the internet, disabling an inspection category and writing a broad allow rule all stay human decisions, and each carries a red warning before anyone can pick it.

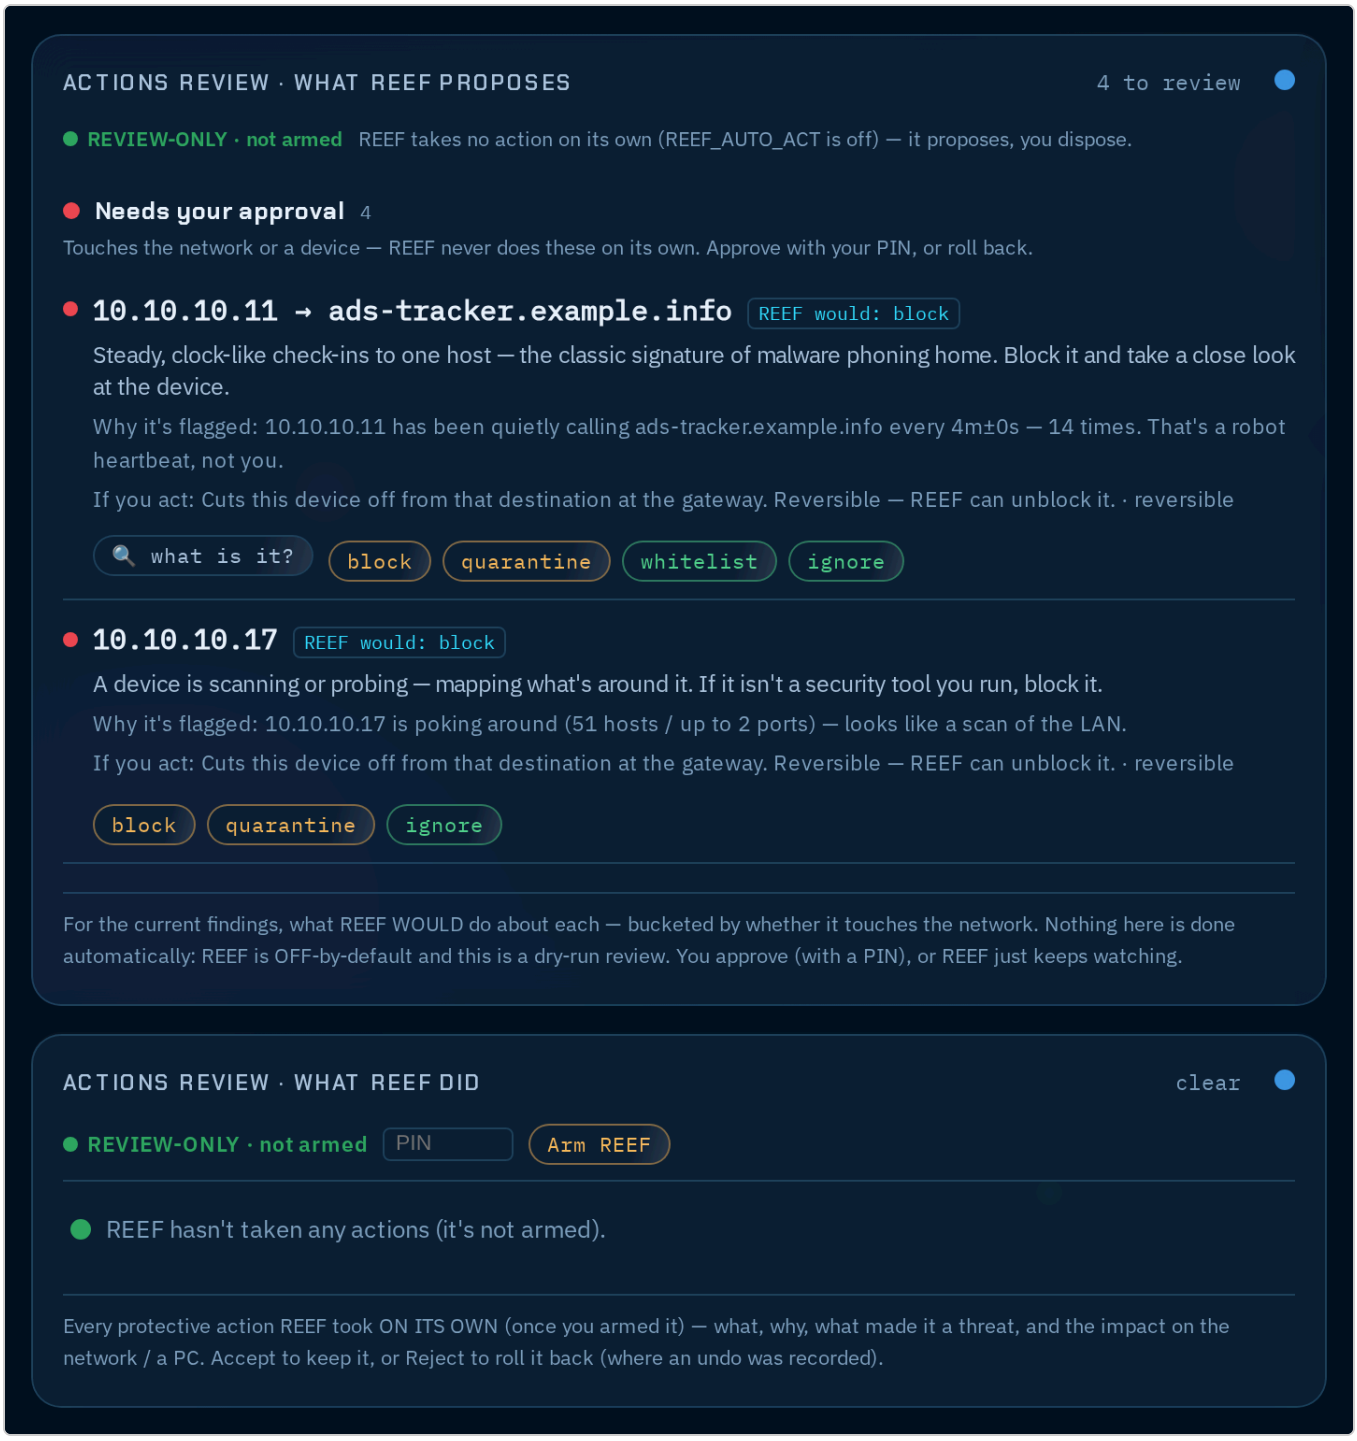


Figure 8. The action review, running review-only and unarmed, with every proposal waiting on the owner's PIN and a PIN field beside the control that arms REEF.

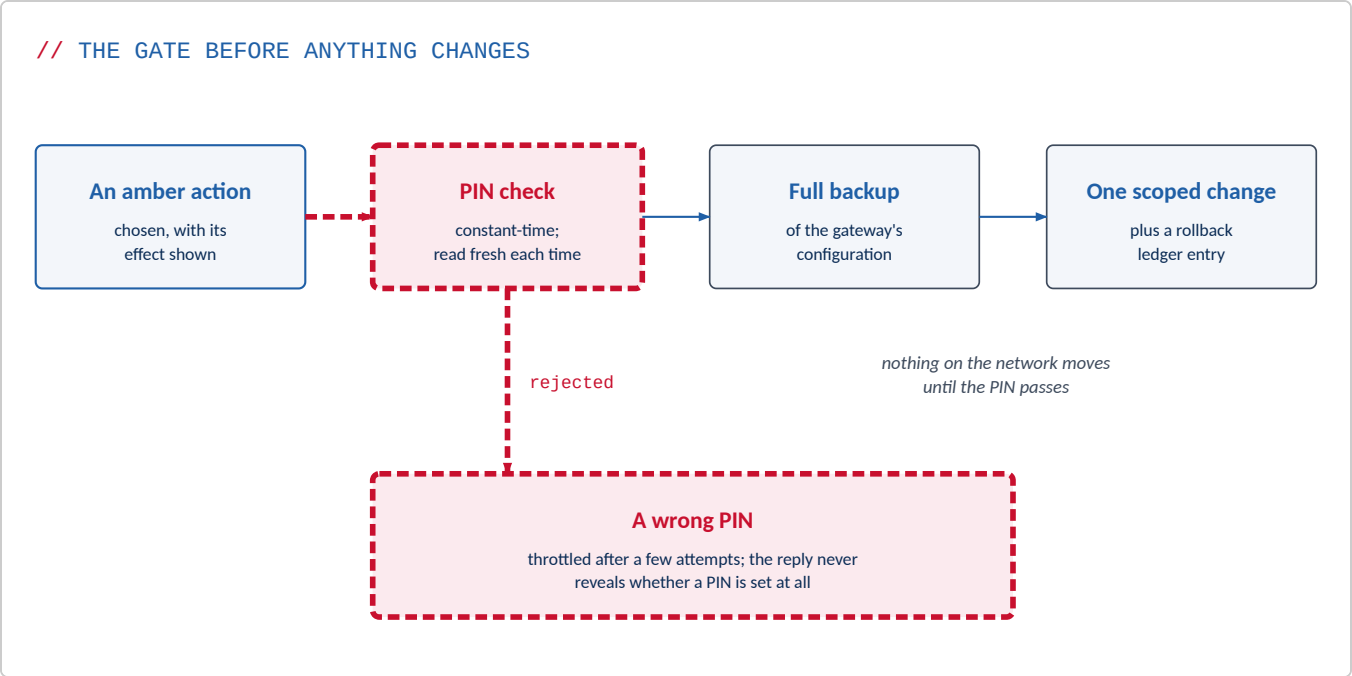


Figure 9. The gate between choosing an amber action and anything on the network changing.

Honest about what it could not see

”We did not look” and ”we looked and it’s fine” are different answers. A tool that renders them the same way is lying, and REEF holds that line in code, in seven places.

Table 3. Seven places one rule is enforced rather than recommended.

WHERE IT SHOWS UP	WHAT THE RULE FORCES
An unreadable sensor feed	It sets a degraded flag. It never yields an empty result an interface would render as all clear.
The interface data layer	Six states rather than two: loading, live, empty, stale, offline, error.
Security modules	They are configured not to retain the last good result, so a failed refresh cannot leave a stale green on screen.
The static-analysis gate	"Could not scan" is a failure rather than a pass, and a high-severity finding can never be baselined away.
Status dots	They mean the lane ran, never that the code is fine.
A gauge with no reading	It renders as a dashed outline reading "no reading", visibly unlike a reading of zero.
A check that cannot see its comparison	It reports unknown and says so. Unknown is a third state and is never folded into agreement.

The opposite failure took longer to see. Its damage is just as real. Switch a capability off on purpose, and a tool that keeps reporting it as a high-severity defect teaches its owner to skim past the loudest line on the page, which is exactly where the next real signal will be sitting. So a capability in REEF has three readings: working, broken, and off on purpose (the third carrying a stated reason and a condition that lifts it).

Deliberate deferrals live in a registry built to resist abuse. A file that turns red items green is one careless edit from muting a firewall's alarms, so four guards hold it. Each has its own test. A deferral names one specific item; wildcards are refused. One missing both a reason and a lifting condition is ignored, and REEF says so on screen. Unknowns sit outside the registry's reach altogether, because unknown means "we could not look" and nothing is allowed to silence that. And a deferred item still appears with its evidence; it simply leaves the queue of things to work.

One clause there matters more than it looks. With paging switched off, REEF says a change will show up on its own screen, and it never tells anybody they'll be paged. A

deferral stays a deferral and never turns into a promise.

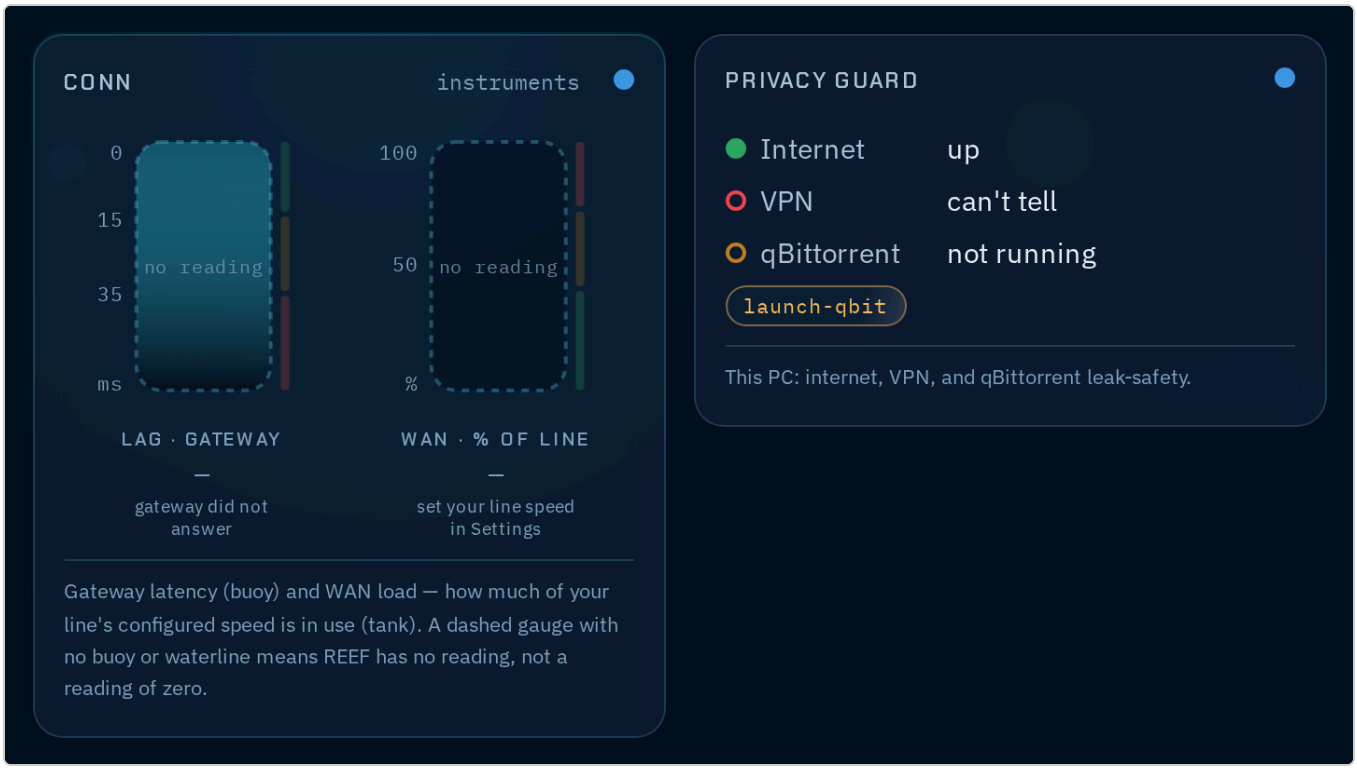


Figure 10. The CONN gauges drawn as dashed outlines reading "no reading", beside Privacy Guard answering "can't tell" for the VPN.

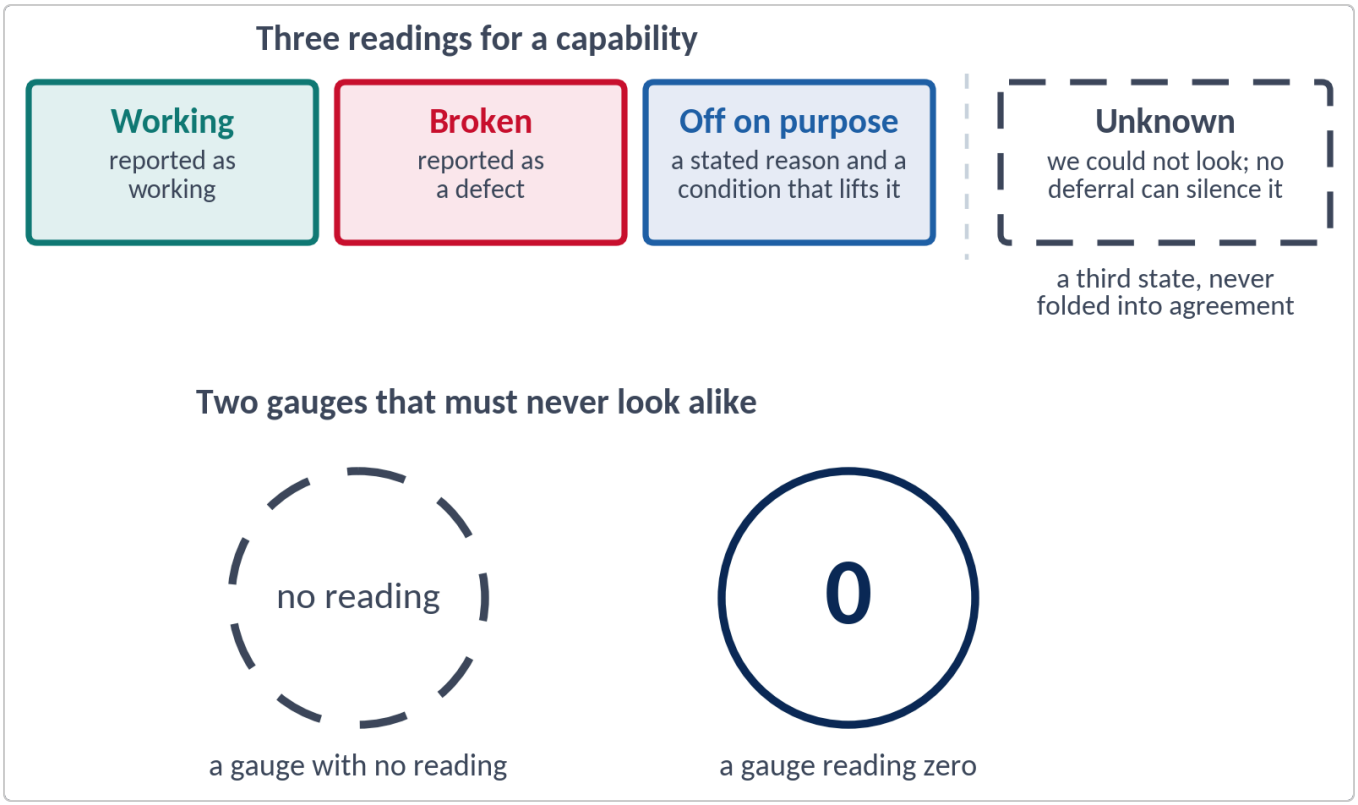


Figure 11. Three readings for a capability, the unknown that sits outside them, and two gauges that must never look alike.

Why there is no shell on the gateway

About once a quarter a sensible-sounding proposal comes back: put an agent on the firewall that can run commands, so REEF can reach what the appliance's own interface doesn't expose. The wish list is always the same four things, packet capture, deep firewall state, a custom script and a log tail, and every one of them is a real need.

The project records it as rejected, permanently, for three reasons in rising order of force. A command agent converts a scoped, revocable key into root code execution on the perimeter. It leaves the appliance's supported and hardened path, which means maintaining a bespoke privileged component on the security boundary. Most decisively, an inbound command channel on the firewall is precisely the command-and-control pattern REEF exists to catch. Nobody can credibly hunt beacons from inside one.

What exists instead is a ladder of sanctioned routes, tried from the top down.

Table 4. The ladder that exists so nobody has to invent an exception at two in the morning.

RUNG	WHAT IT IS	WHEN IT IS REACHED FOR
1	The appliance's own interface	Always first
2	A named, parameterised, allow-listed action defined on the appliance	When the interface genuinely cannot reach
3	A read-only telemetry sidecar	When the need is observation rather than action
4	A separate probe host on a mirror port	When the need is packet-level and off-path
none	A general remote shell	Never

A capability REEF lacks today arrives when somebody writes a specific, reviewed action for it. That friction is the feature.

What REEF will be

REEF is built to be one organ of something larger. The house it guards has a resident, a local AI assistant living on one workstation and keeping its knowledge there. Two guards stand around it. A code reviewer watches the repository the resident's code lives in, and that's the yard. REEF holds the gateway and the network, and that's the coastline. Past the coastline lies open water: traffic nobody in the house wrote, from services nobody in the house controls.

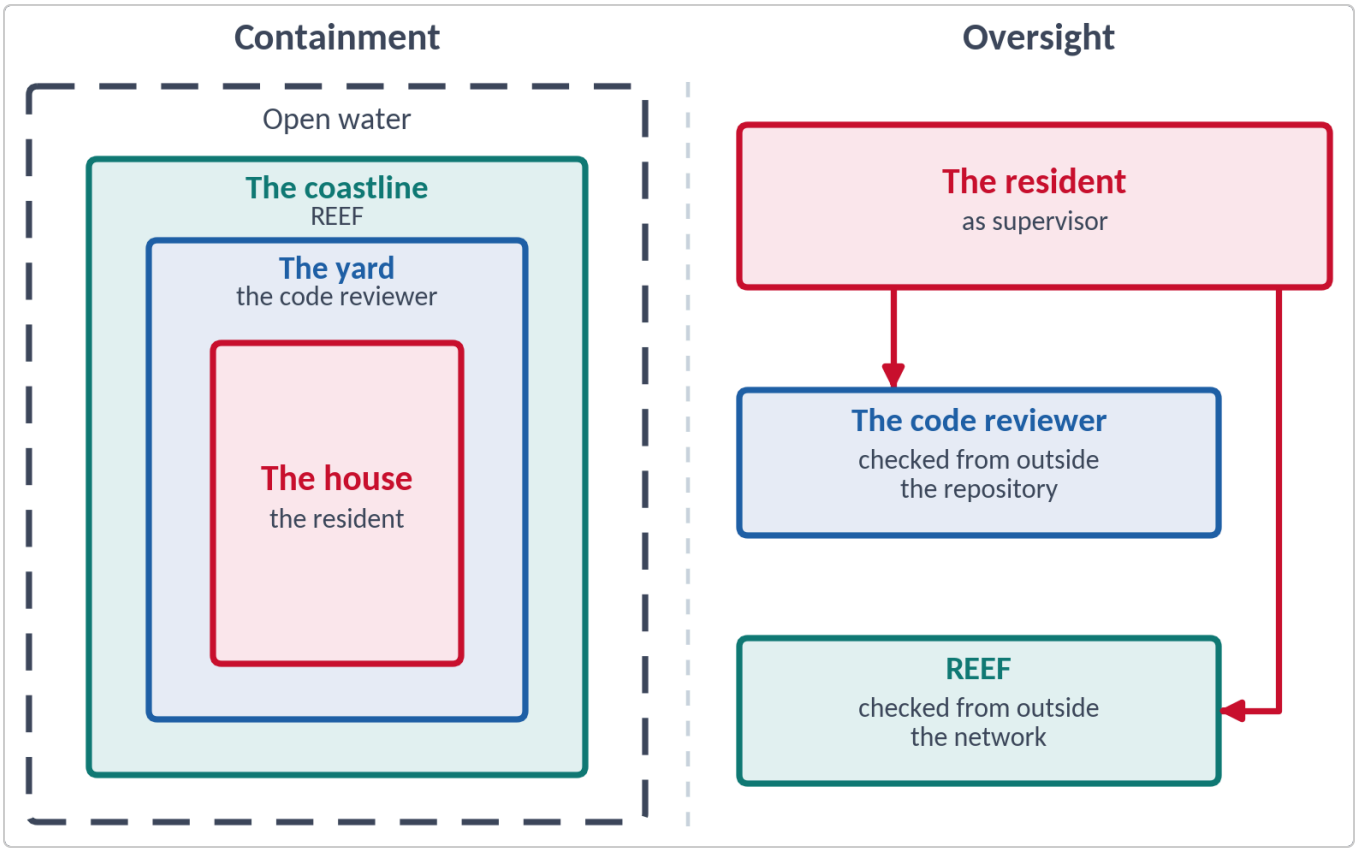


Figure 12. Containment on the left, oversight on the right. The same three parts are arranged by two different relations.

The next layer puts the resident above both guards. In space it sits innermost, protected by both. For the purpose of checking them it stands outside both, because it's the only part of the arrangement that lives inside neither thing being checked. A guard can't see its own blind spot from where it stands.

That need showed itself one night when REEF, acting unattended and exactly as designed, cut every device off the gateway, the resident's workstation included. Nothing outside the house had to fail. The owner's direction arrived two days later, and it is still the brief: "This is your home. He has a job and you should use him. Ensure he never breaks or locks you in your home."

REEF keeps its job and keeps its authority to block. What the next layer adds is somebody positioned to check it, through two channels.

Table 5. Two channels, because a stale or lying self-report must not be able to hide a lockout.

CHANNEL	WHAT IT CARRIES	WHAT IT COSTS	WHAT IT CANNOT DO
REEF's self-report	Redacted metrics about its own state	Almost nothing; it is already being computed	Detect the state in which it is wrong about itself
The resident's own reachability check	Whether the world is actually reachable from here	A separate probe on a separate stack	Explain why, which is what the first channel is for

The supervisor also holds the mirror image of REEF's powers. REEF may add a restriction. The supervisor's only network action undoes a restriction or restores a known-good configuration, a saved state the gateway can safely return to, and it can never add anything. So the supervisor can't become the outage it exists to prevent. The code reviewer, covered in "Local-First Code Review: The Guard Dog and the Ledger", runs under the same law from the other side of the house.

Table 6. An organ that can only subtract risk cannot fail open, whatever else goes wrong with it.

ORGAN	WHAT IT MAY DO	WHAT IT IS STRUCTURALLY DENIED
REEF	Add a restriction	Weakening the perimeter on its own initiative
The code reviewer	Add findings, and at most block a change	Approving anything
The supervisor	Undo a restriction, restore a known-good configuration	Adding any restriction at all

Read the three rows together and one shape appears. Each organ's denied column holds exactly the action that would turn its own failure into somebody else's incident.

From that follows the supervisor's standing posture. It observes freely and recommends by default. It acts only on one of REEF's own reversible recoveries, and only when somebody has armed that lane in a decision logged at the moment it was made. Before anything that could sever the gateway it warns in plain language. Then it stands aside, because the owner holds full authority over their own network, and a supervisor with a veto would be one more way to get locked out. Counsel survives the case where the person is right and the machine is wrong. A veto doesn't.

More organs will be built to this shape, each guarding something the resident can't watch alone. The reason to build them this way reaches well past one house. A trust ladder that won't promote a claim without corroboration, a reviewer that can object and never approve, a supervisor that can only undo, and a record of what was said and whether it held: together those are what auditable means when somebody asks, later, how a decision got made. Joining data helps most in the places where bringing it together has always been forbidden — a patient's history scattered across systems that can't speak to each other, a fraud signal split between two institutions — because bringing it together has always meant sending it somewhere. A resident that sends nothing anywhere can hold both sides at once, and REEF is what stands on its coastline.

The porch camera, as it turned out, was talking to a firmware update server. REEF had the rhythm right, said plainly what it didn't know, and said all of it before anybody blocked anything.

Glossary

TERM	DEFINITION
Advisory	is a finding real enough to show and too thin to act on; it keeps its evidence and loses the verbs that would change the network.
Arming	is turning on a capability's authority to act rather than recommend. It is explicit, made per capability, logged at the moment of the decision and never switched on by an upgrade.
Beacon	is a regular, automated call from a device to an outside address. Rhythm is its signature, and jitter leaves the rhythm in place.
Coastline	is the outermost guarded band, holding the gateway and the network. REEF guards it, and past it is traffic nobody in the house wrote.
Corroboration	is independent agreement between sensors about the same conversation. Two readings of one feed count once.
Deep-packet inspection	is an engine that classifies traffic by more than its address and port. REEF commands the gateway's own engine and treats its verdicts as one input among several.
Degraded flag	is a marker travelling with a response to say part of the picture could not be read; it exists so an empty answer can never be rendered as a clean one.
Enrichment	means looking a destination up against outside reputation data. It runs on a budget, cache first, and can raise a severity without ever creating a finding.
Fingerprint	is the shape of an encrypted handshake, which can disagree with the application the traffic claims to be.
Gateway	is the box where a household network meets the internet, carrying the firewall and its inspection engine. REEF sits beside it.
Jitter	is deliberate variation added to a beacon's interval so it looks less mechanical; it widens the spread of intervals and leaves their centre

TERM	DEFINITION
	where it was.
Known-good configuration	is a saved gateway state the system can safely return to. Restoring one subtracts risk, which is why it's one of the only two writes the supervisor holds.

A loopback listener sits on a network address that answers only the machine it runs on, which is why reaching REEF from outside means already being inside.

Metadata records who talked to whom, when, how often, how much and over what. All of it survives encryption, and reading it requires decrypting nothing.

Off-path means sitting beside the traffic rather than inside it. Remove an off-path reader and the network is unchanged.

The PIN gate is the check standing between a proposed change and a real one. It is compared in constant time, read fresh each time, and throttled after a few wrong attempts without revealing whether a PIN is set.

A reachability check is a probe run from a separate stack asking whether the world is actually reachable; it exists so a self-report can never hide a lockout.

The resident is a local AI assistant living on one workstation and keeping its knowledge there. In the next layer it supervises REEF and the code reviewer.

A reversible recovery is an action that undoes a restriction and can itself be undone, and the only class of action the supervisor may take unattended.

The supervisor is the resident in its checking role, holding an independent second opinion about a guard from outside the thing that guard sits inside.

A tier is the blast radius of an action, which decides the ceremony it needs: instant, a PIN plus a backup, or an explicit warning.

TLS, or Transport Layer Security, is the encryption behind the padlock in a browser's address bar and behind most traffic a modern device sends.

The yard is the middle guarded band, holding the repository whose changes the code reviewer reads.

MAX3 WPR-MAX-019

// ABOUT THE AUTHOR

Jason Newell leads organizations and people in technology, and is a researcher, deep technologist and builder — from a CRAY to agentic AI. He founded MAX Research Collective to build AI for regulated industries that runs where its owner has control, with healthcare as its strongest focus. About Jason → Connect on LinkedIn →