FIELD DOSSIER // MSSP OPERATOR PLATFORM

OMNISCIENT
ONE OPERATOR. THE WHOLE FLEET.

> Many toolkits, one underlying data stream — idempotent, reversible, and what-if-able. Discovery, hardening, deception, and live adversarial validation, run and re-run from a single console. Built to be operated at scale, and built to be resold under your own name.

RESPONSE TIERS	DECEPTION ZONES	HEARTBEAT	KEY MODEL
T0–T5	5	RSA-2048	2-KEY
configurable ladder	outer → crown jewels	signed dead-man	authority ÷ encryption

[00 // INDEX]

WHAT'S IN THIS DOSSIER

[01 // THE PROBLEM]

MONTHS VERSUS SECONDS

The average intrusion sits undetected for **months**. Most managed security is a dashboard nobody watches and a hardening pass that happens once, then goes stale the day after it's written.

For an MSSP, that gap is the whole business risk. You put your name on a client's security posture, then find out about the breach the same week the client does — from a ransom note, not your own tooling. The tools that were supposed to catch it were busy generating alerts no one had time to read.

The industry's answer has been **more**: more sensors, more alerts, more dashboards, more analysts to triage the noise. That doesn't close the gap; it just moves the fatigue around. Every added tool is another vendor, another console, another integration seam for something to slip through.

> **THE INSIGHT THIS PLATFORM IS BUILT ON** ◆
>
> Security isn't about stopping everything — nothing stops everything. It's about making sure that the moment an attacker does get in, they hit a tripwire immediately, and the response that follows is automatic, contained, and recorded. Detection measured in seconds, not months. Coverage that doesn't depend on someone watching a screen at 3 a.m.

Omniscient is one platform built around that single idea, running the entire lifecycle of protecting a fleet — discovery, access, hardening, deception, and live validation — from one console, under one governance model, operable by one analyst. For a reseller, that last part is the point: it turns headcount into margin.

[02 // ARCHITECTURE]

ONE PLATFORM, ONE DATA STREAM

Omniscient is not five tools bolted together. It's many toolkits — recon, hardening, deception, access, red team — running as views and verbs over **one governed data stream**. That single design decision is what makes everything else safe to operate at scale.

Because it's one stream underneath, a mine trip, a heartbeat silence, and a network-drift event all land in the same feed, on the same incident thread. Nothing is stranded in a separate console. And that stream has three properties that, together, are the reason a single operator can run an entire fleet without fear:

PROPERTY 01

IDEMPOTENT

Run the same operation twice and land in the same state — no duplicates, no errors, no drift from repetition. Every door, rule, and decoy is named, so the system always knows exactly what it did and can find it again. You are never afraid to re-run.

PROPERTY 02

REVERSIBLE

Every action ships with its inverse in the same pass. Apply pairs with rollback, open with close, provision with restore. No one-way doors. Nothing you do to a client's production box is a commitment you can't walk back.

PROPERTY 03

WHAT-IF-ABLE

A pre-flight plan shows the exact consequence before you commit. Sentinel simulates before it arms. The box itself is the preview — so you operate confidently without needing a separate lab to rehearse in.

THE GOVERNANCE MODEL

Policy inherits down three levels, each mapping to something an operator already thinks in: National is the network — topology, segments, switches, the whole fleet. Federal is the machine — system-wide hardening that applies to every account unless overridden. State is the user — role-based policy, because a General Manager's workstation is not a line cook's.

When the engine finds a live account with no role mapping yet, it doesn't pretend the account isn't there — it's surfaced as ungoverned, with a one-click path to promote it into the model. Drift isn't just flagged; there's a wired path to resolve it. **Nothing stays invisible by accident.**

FIG.01 The fleet map — every managed machine, switch, and network segment in one topology, with live counts for machines, users, admins, drift, and sync state. This is the National view: the whole client estate on one surface.

[03 // LAUNCHER]

ONE BINARY, FIVE WAYS IN

The platform runs off a single launcher that checks its own environment first, then hands the operator the interface that fits the box in front of them — and **never hard-fails** because a dependency is missing.

On first run it reports what it's standing on: admin state, language mode, and whether the tooling it needs is present, flagging anything missing with the exact remediation rather than a cryptic error. Then it offers five interfaces over the same engine: a native console, a 3D fleet map, a full-controls manager, a plain text console for locked-down boxes, and a self-contained bundle. If the richest option isn't available on a given box, it falls back automatically instead of stopping.

WHY A RESELLER SHOULD CARE

Your operators will meet every kind of client box — modern, ancient, locked down, air-gapped. A tool that degrades gracefully instead of failing means one platform covers your whole client base, not just the convenient half. Fewer exceptions is fewer truck-rolls.

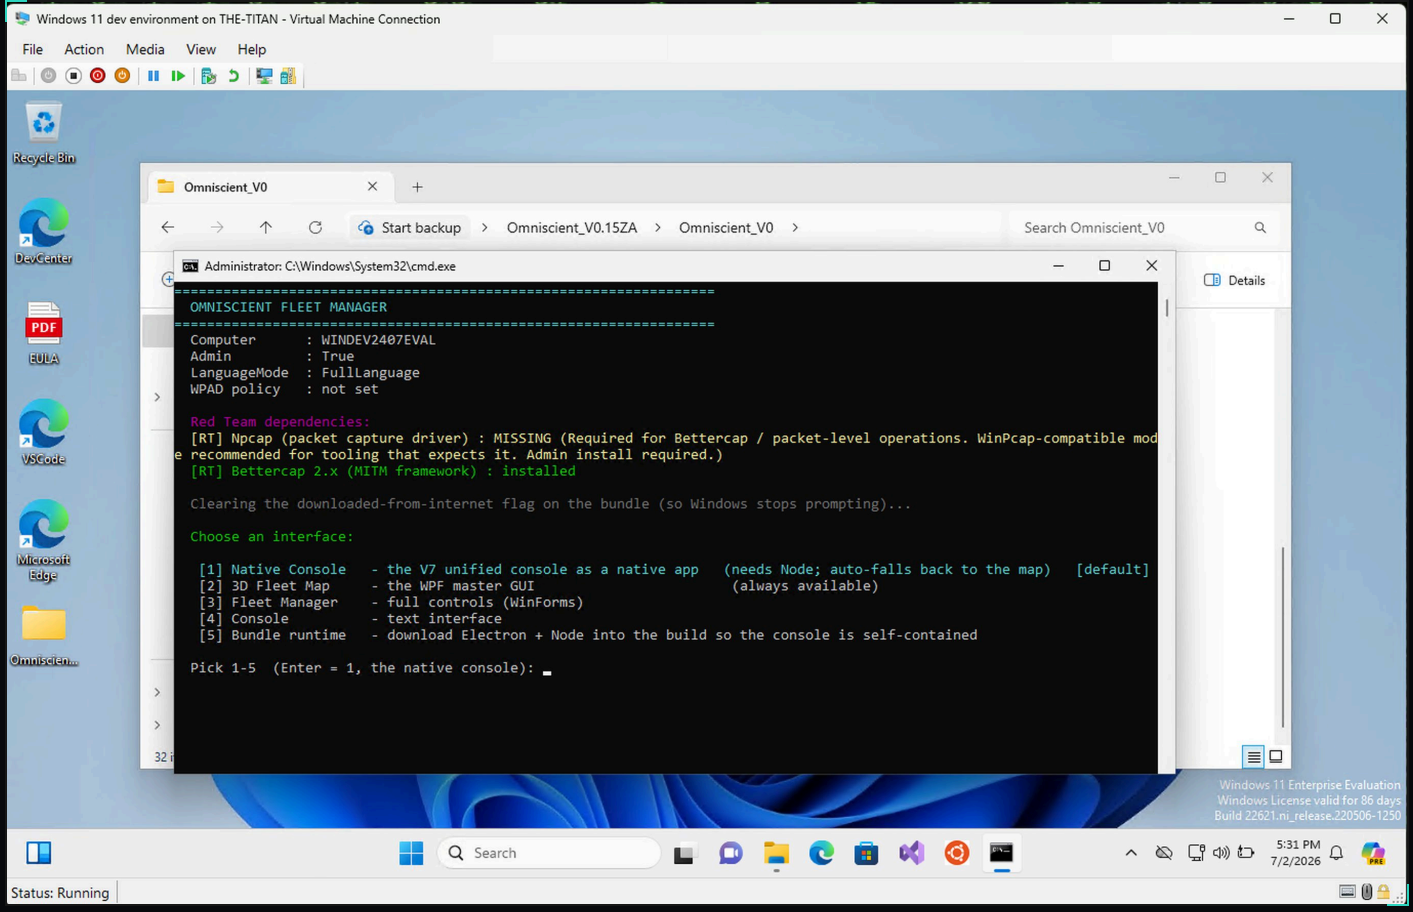

FIG.19 The launcher's first-run self-check and interface picker — environment reported honestly, dependencies flagged with fixes, and five ways to run the same engine depending on what the box supports.

[04 // RECON]

SEE WHAT'S ACTUALLY THERE

Every engagement starts with what's actually on the wire — not what a client's documentation claims. Discovery flags **drift**: anything present that isn't in the governed model, the same low-level telemetry a SIEM diffs over time.

A fast parallel sweep maps the subnet; a deeper pass probes services. Whatever it finds that you didn't put there is surfaced as drift, with real MAC and vendor data — which is often the single most valuable thing you can show a new client: the devices sitting on their own network that nobody told them about.

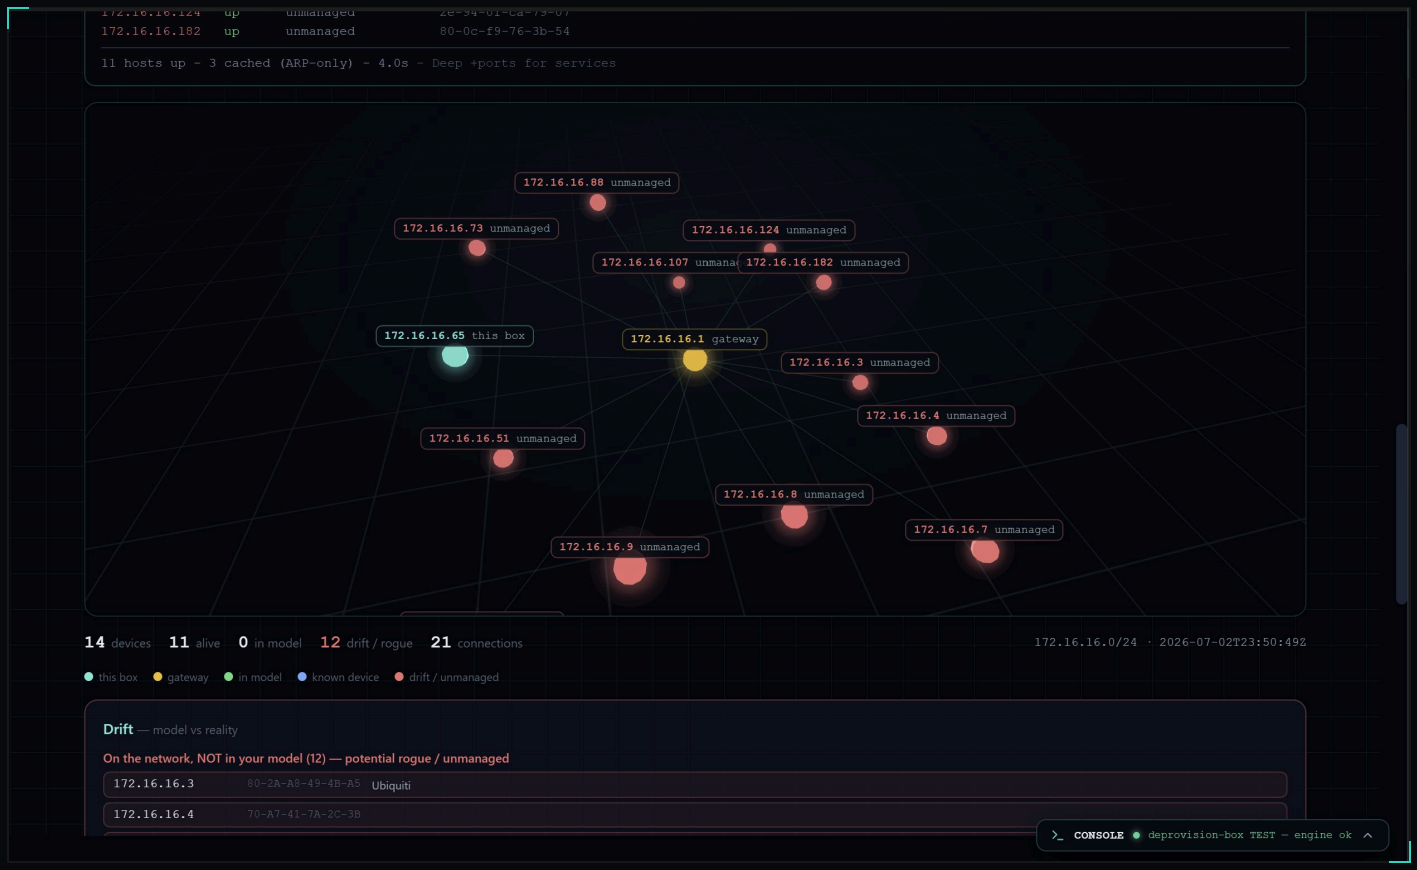

FIG.12 A live subnet sweep: managed endpoints resolve clean, and everything undocumented is flagged as drift — each rogue device carrying its real MAC address and vendor fingerprint.

The capture layer is deliberately conservative. It reports the best sniffing tier a box can actually run — native telemetry, wired frame capture, WiFi monitor mode — **without ever opening an adapter** until an operator explicitly starts a capture. When a tier is blocked, it says exactly why and how to reach it, rather than silently under-delivering.

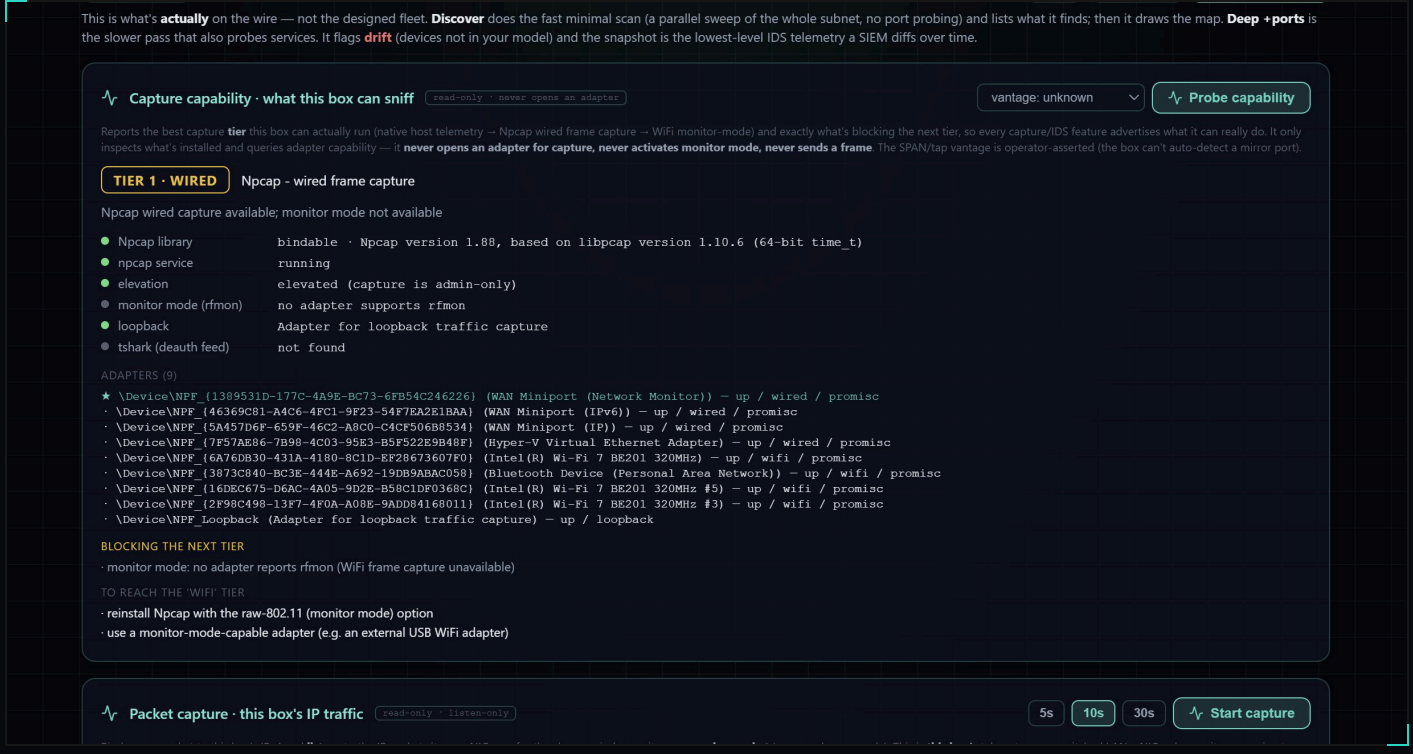

FIG.11 The capture-capability probe: full adapter inventory, the precise reason a higher tier is unavailable, and the exact steps to unlock it. It advertises only what it can truly do.

A CREDIBILITY BEAT WORTH NOTICING

Run recon inside a NAT boundary and the platform detects it, explains the mechanism, and tells the operator how to get a true reading — instead of handing back a falsely clean scan. A tool that tells you when it can't see something is worth more than one that always shows green. That honesty is what you'll stake your reputation on when you resell it.

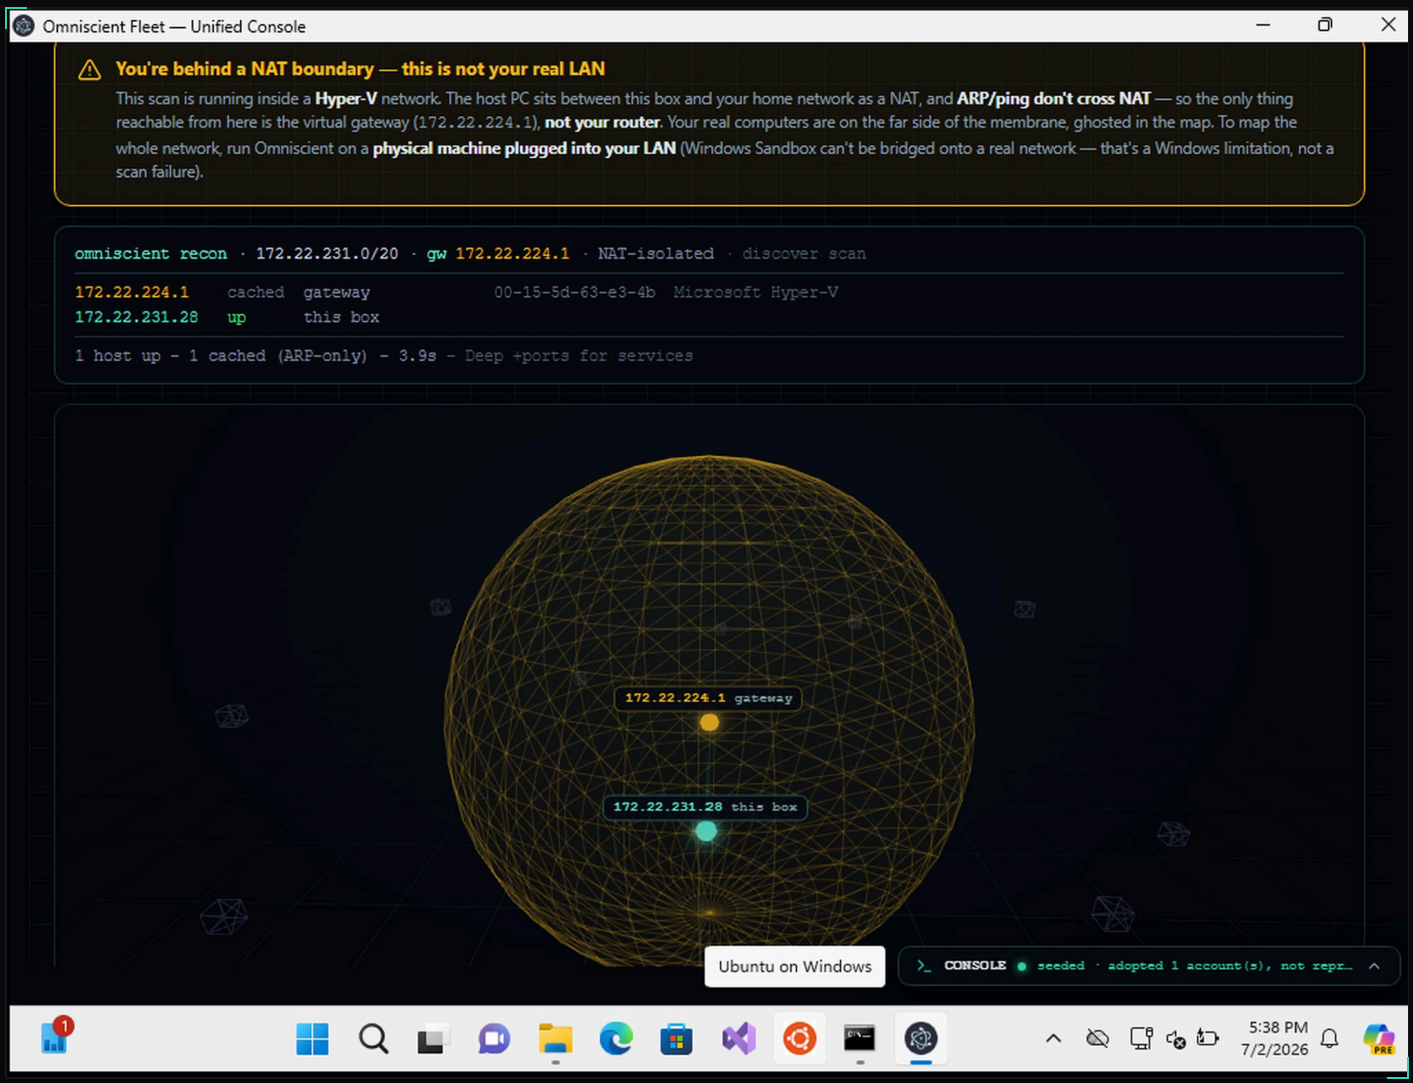

FIG.21 Recon correctly identifying a NAT boundary and explaining why the scan is limited — rendered in the platform's 3D topology view. The tool refuses to fake a result it can't stand behind.

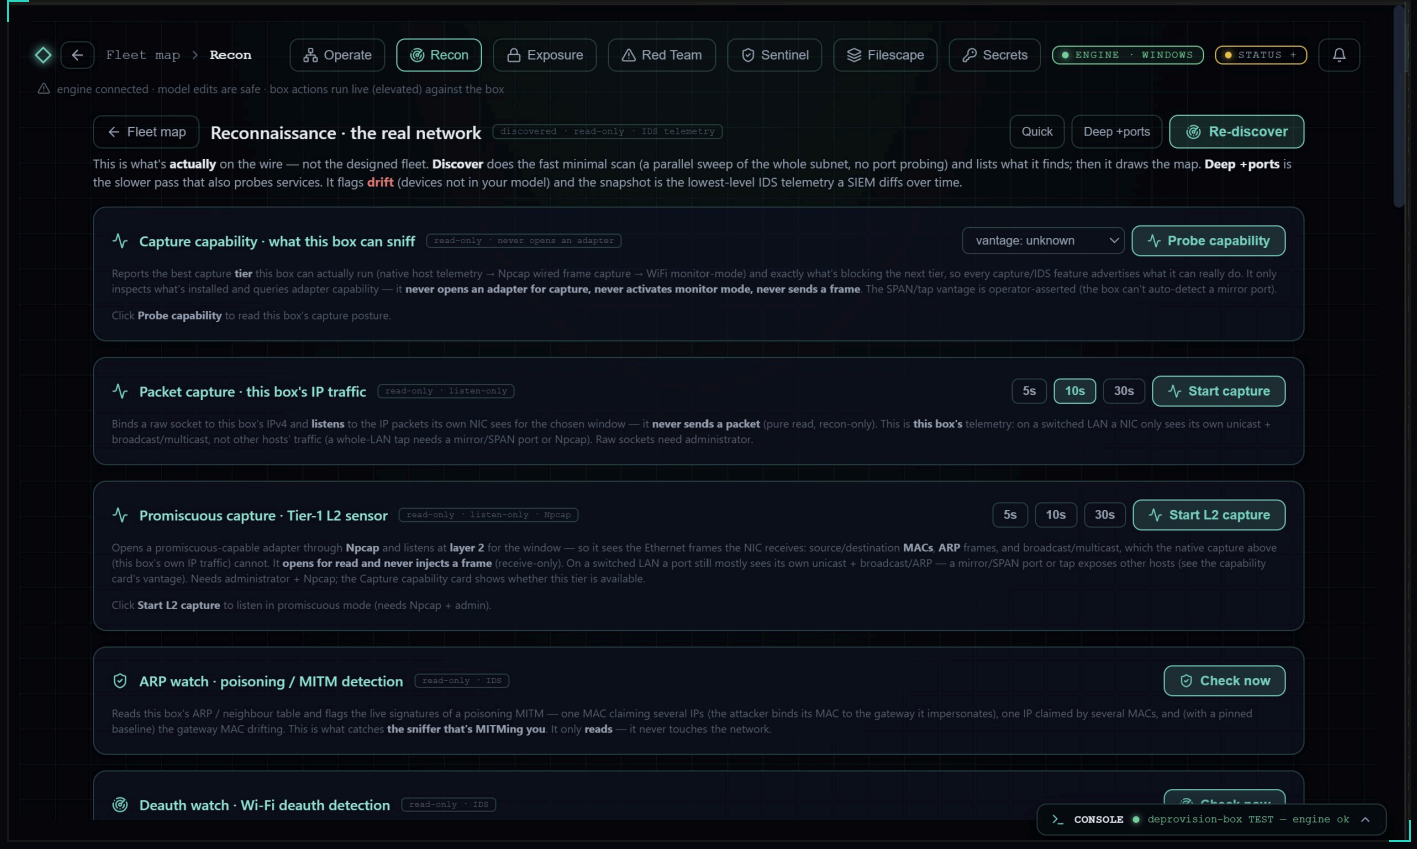

FIG.10 The reconnaissance module overview — discovery, capture, promiscuous L2 sensing, and ARP/MITM watch, each clearly scoped and each read-only until deliberately engaged.

[05 // ACCESS]

GET IN SAFELY

Every managed box's resting state is hardened and closed. Remote access is the exception — and everything the platform opens is **named, visible, and reversible**. There is no "allow any."

The engine refuses wildcard and LocalSubnet firewall rules outright. Every door is locked to a specific office source, every route is scoped and auditable with a one-click "what's open right now" check, and every port forward tears back down to exactly the state it started in. For a reseller, that's the difference between remote access you can defend in an audit and remote access that becomes the client's breach.

- **Source-locked doors** — RDP requires NLA and an explicit office CIDR; the engine will not accept a broad rule.
- **Scoped routing** — internal access forwards through the gateway on a named, auditable rule, never a blanket tunnel.
- **NAT-aware dial-out** — boxes with no inbound path reverse-tunnel out on key-based auth, closeable from either end.
- **Read-only first touch** — scanning an inherited box reads its accounts and posture without changing a single setting.

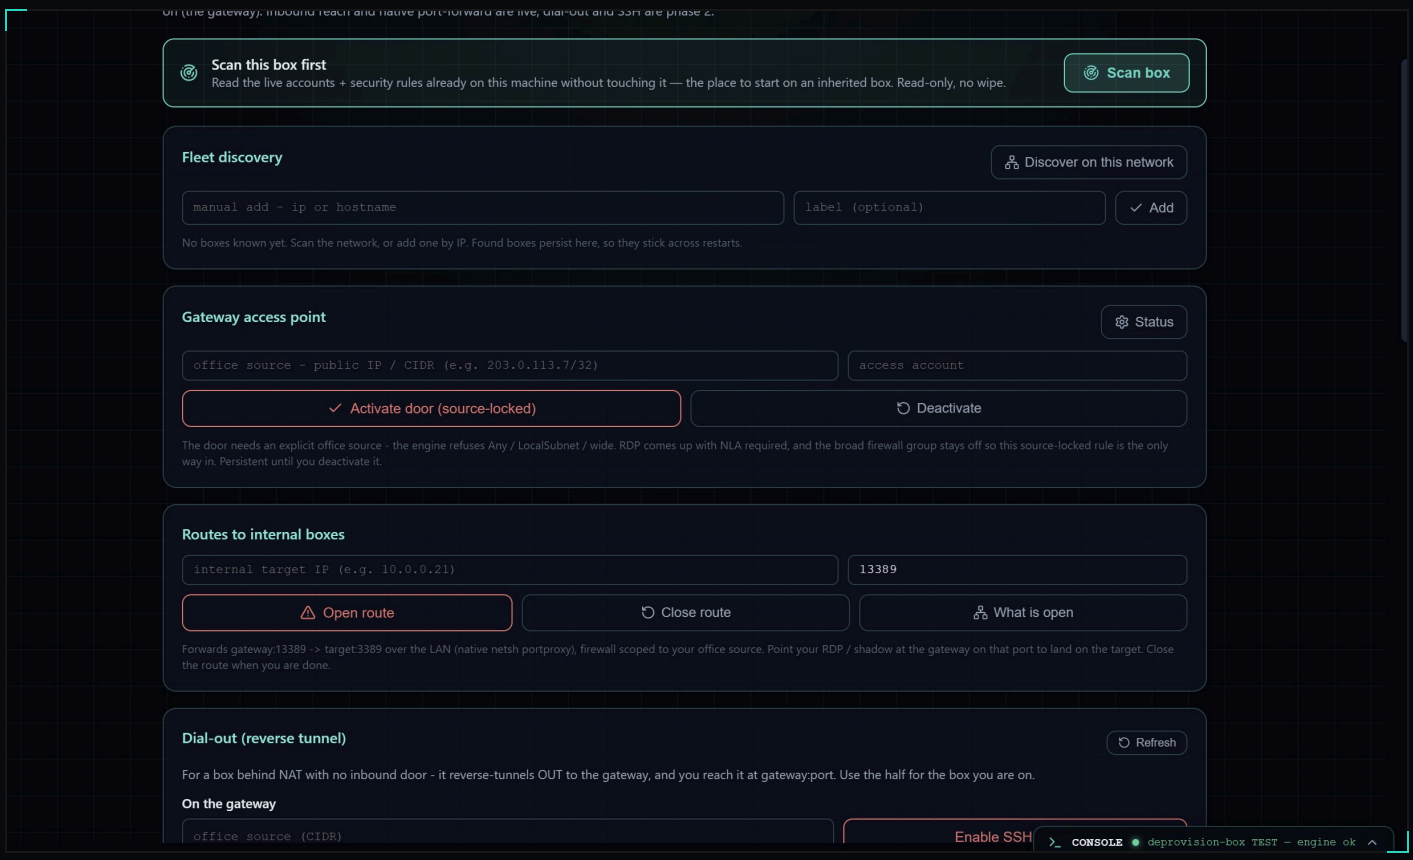

FIG.13 Standing up a gateway access door — source-locked, NLA-required, and torn down with the same one click it took to open. The exposure panel always shows what's currently open.

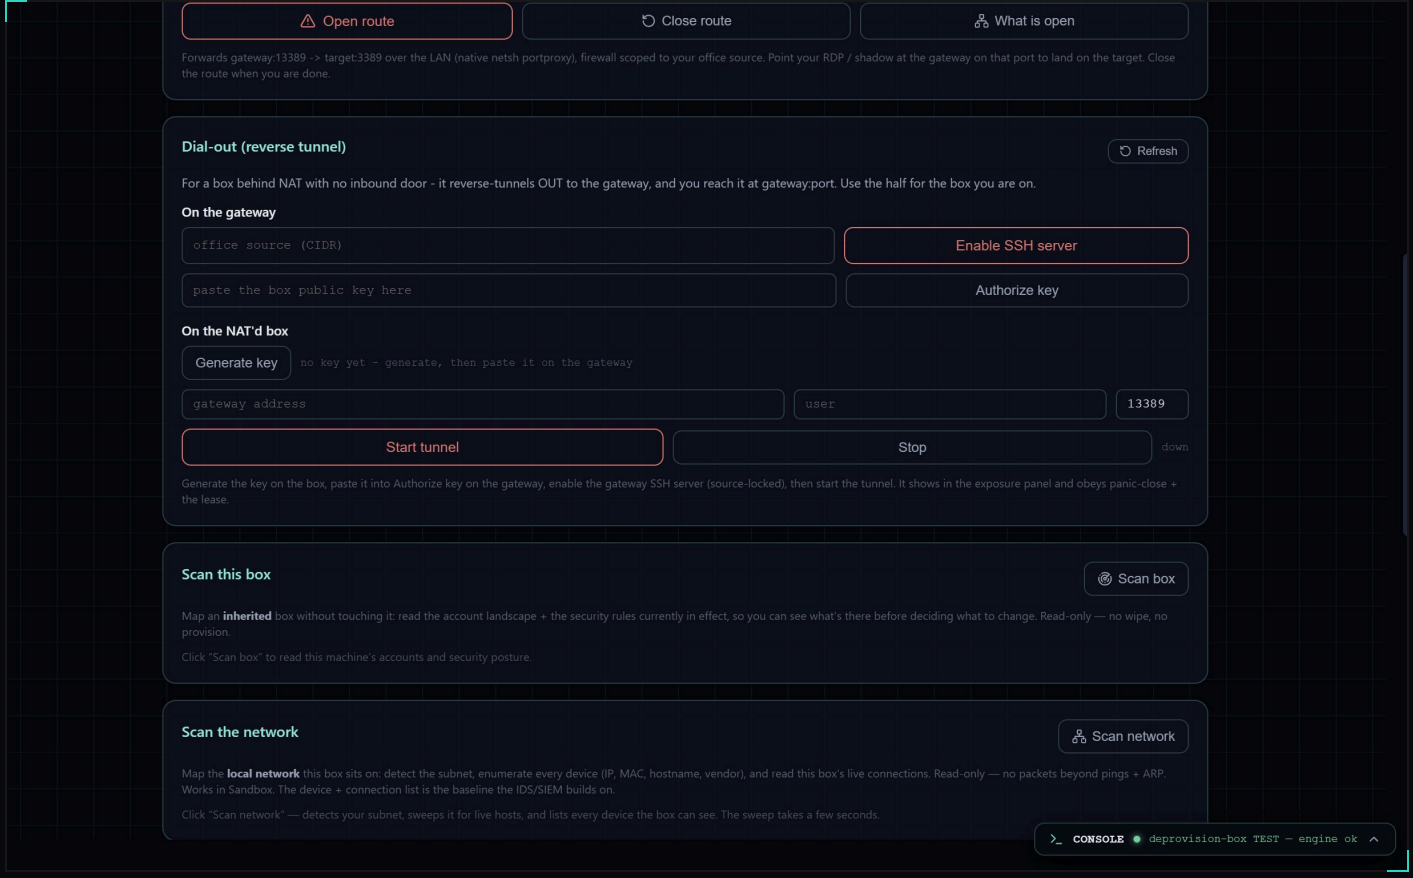

FIG.14 A dial-out reverse tunnel for a NAT'd endpoint — the box reaches out rather than being reached, on key-based auth, reversible and logged, obeying lease expiry and panic-close.

[06 // HARDEN]

PROVISION, REVIEW, FIRE

Nothing executes against a live box without a review step first — the plain-language summary, the exact commands, and an explicit confirmation. **What-if-able, by construction.**

Role-based presets carry sensible starter security for every position out of the box, with per-role web allowlists and application control layered on automatically. Admin accounts are structurally exempt from the lockdowns applied to everyone else — by design, not oversight. And every run is atomic and logged, so you can prove exactly what changed on a client's estate.

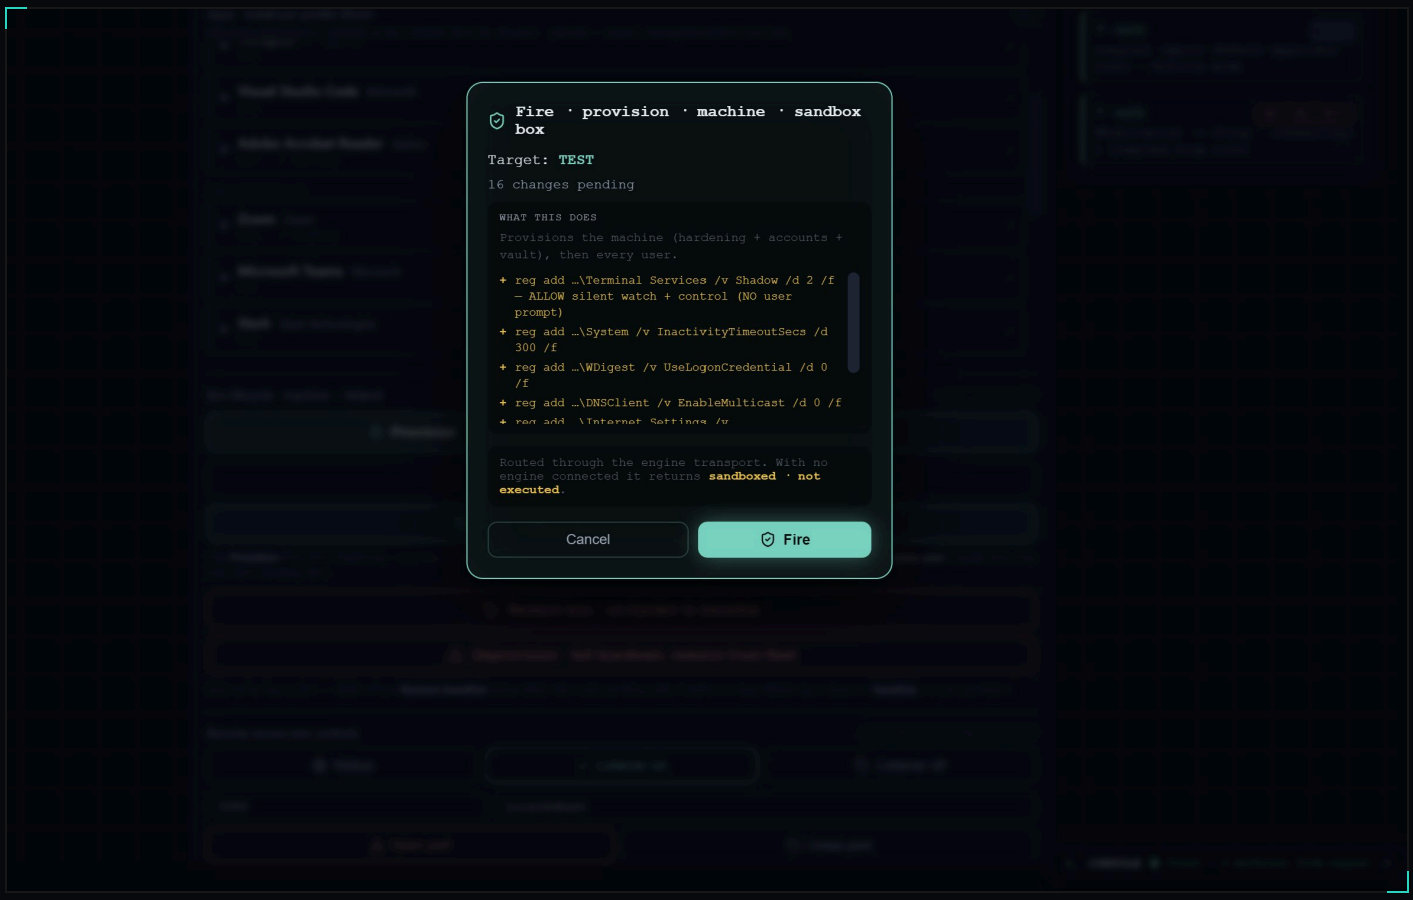

FIG.08 The pre-flight review — every pending change shown, in plain language and as raw commands, before a single one fires. Cancel or Fire; nothing happens by surprise.

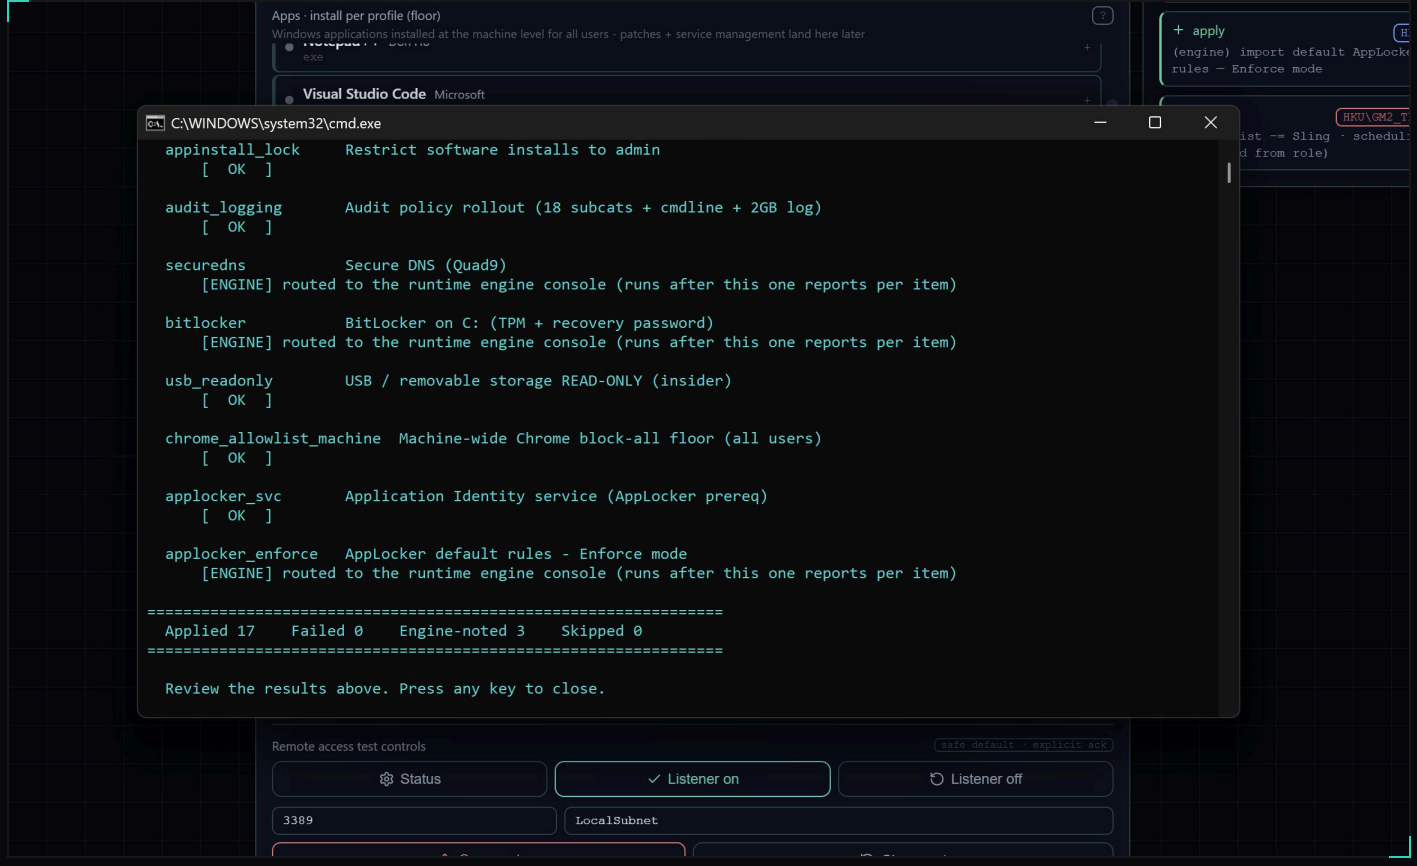

FIG.09 Real execution output after firing: a per-item pass/fail tally — applied, failed, engine-noted — not a spinner that just says "done." Auditable evidence of exactly what landed.

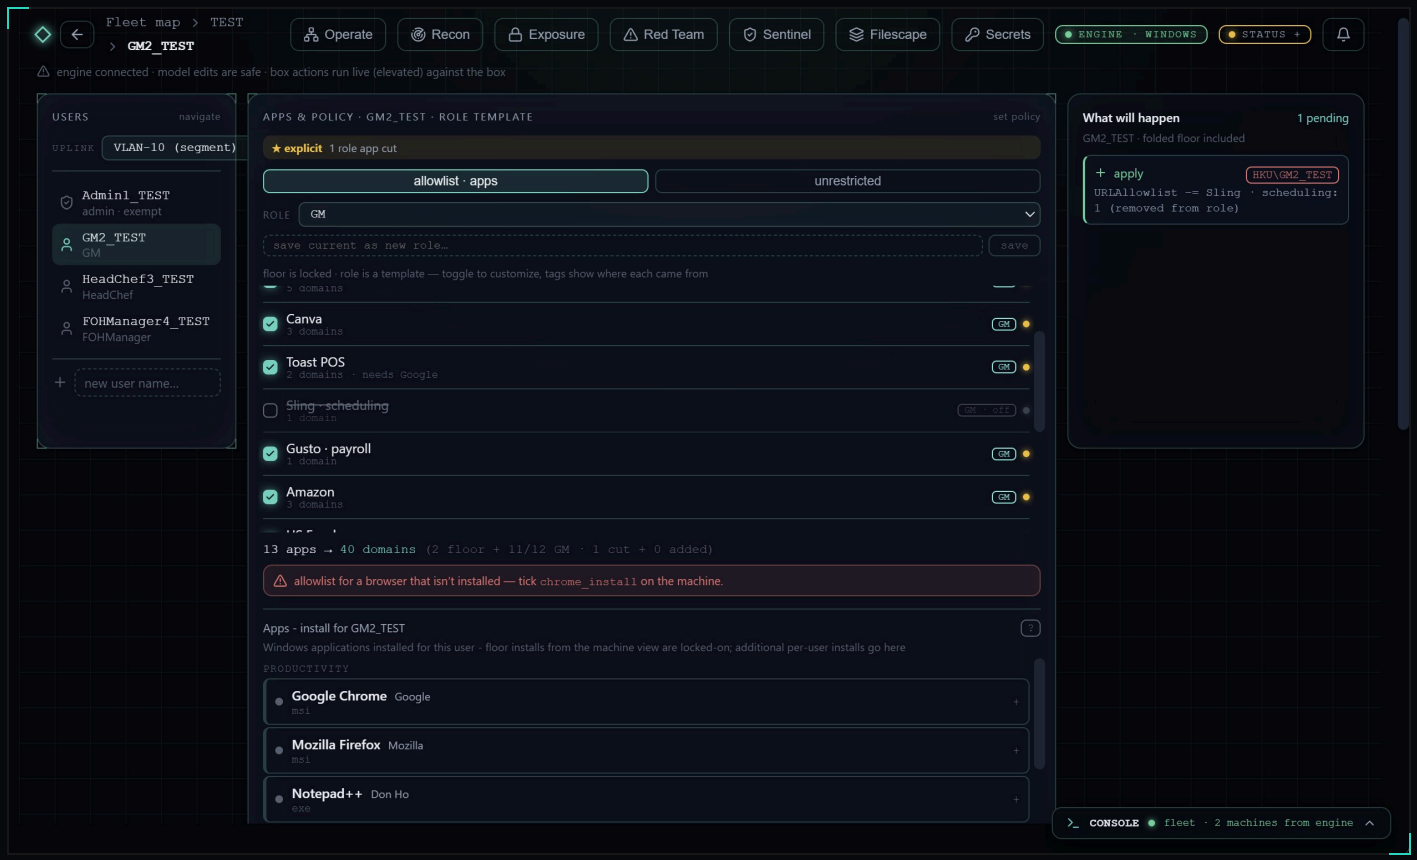

FIG.06 A per-role policy template — the State layer. Apps and domains are assigned by role, cuts are explicit, and the UI warns when an allowlisted app isn't actually installed on the machine.

[07 // KEYS]

TRUSTED WITH THE KEYS

Credential security rests on one deliberate decision: **the key that proves you're the operator is never the key that decrypts credentials.** One shared secret would be a single point of total compromise — so there isn't one.

> **DESIGN NOTE — TWO KEYS, SPLIT ON PURPOSE**
>
> The Operator Key gates authority actions — break-glass, overlay unlock, mine re-arm. PBKDF2-HMAC-SHA256 at 600,000 iterations, verifier-only, DPAPI-wrapped. The Vault Key decrypts the credential vault only, via a separate derivation, AES-256-GCM with CBC+HMAC. Cracking one gets you nothing toward the other.

Passphrases are never stored and cannot be recovered if lost — by design. No hash hard-coded in source, no bare fast hash, no plaintext fallback anywhere; each of those failure modes is explicitly forbidden in the platform's own engineering standard. The Secrets pane gives an operator one place to see every secret the suite uses — what it protects, whether it's set, and where it lives — while **never displaying a secret's value.**

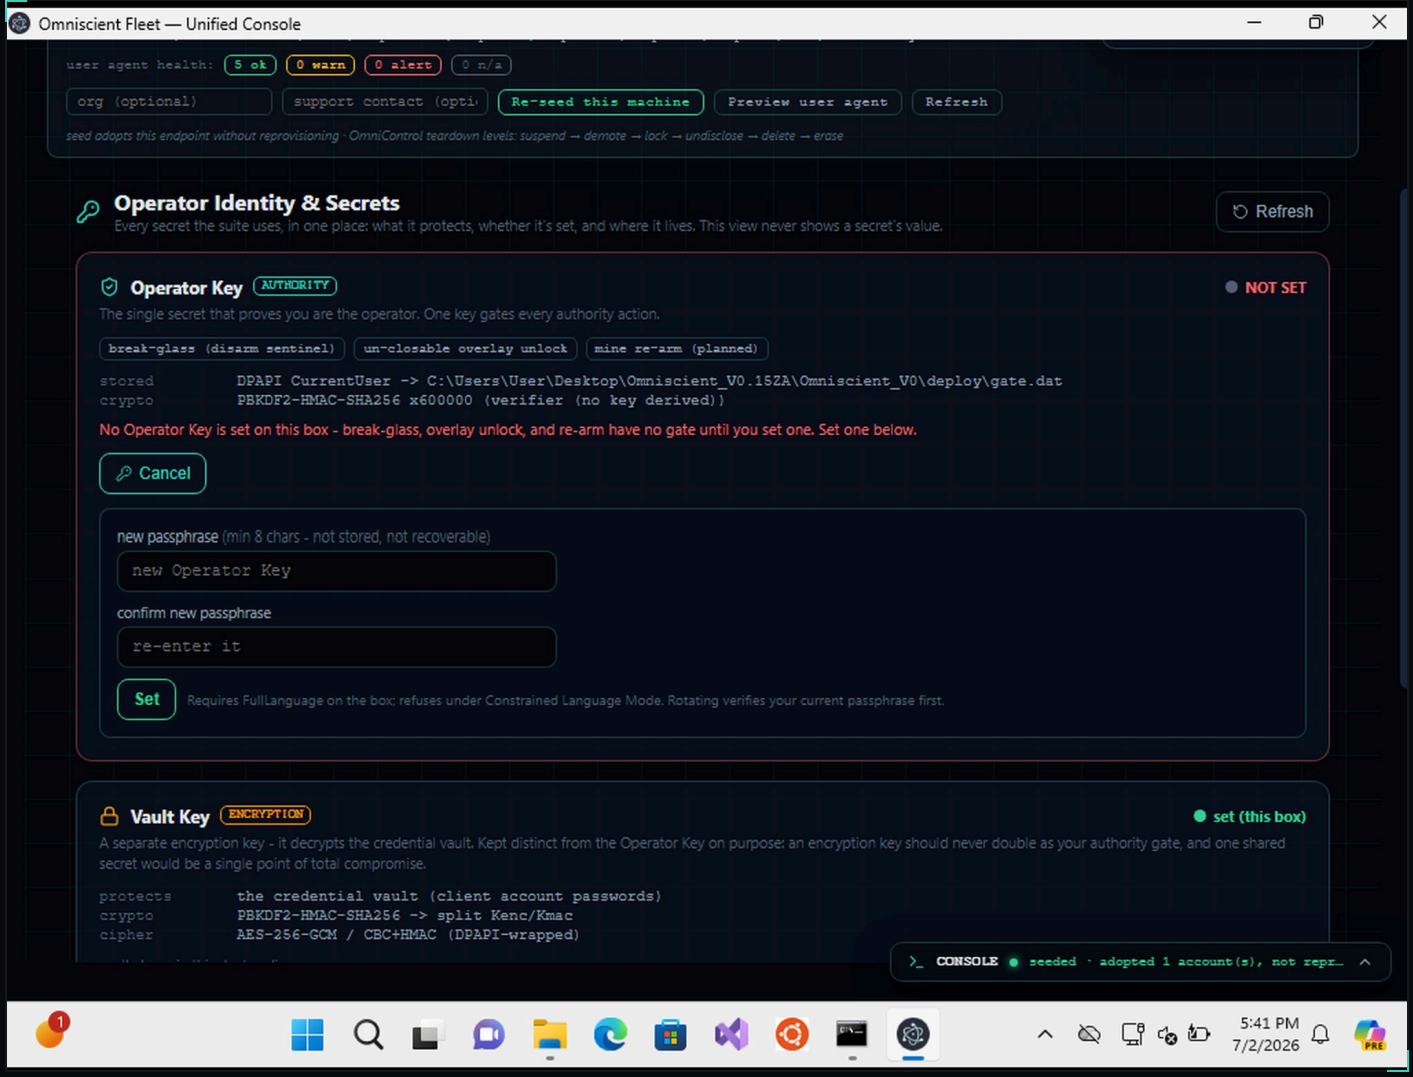

FIG.22 The Secrets pane — the two-key architecture in one view. Set-state and metadata only; no secret material is ever read into the display. Authority and encryption stay separate, always.

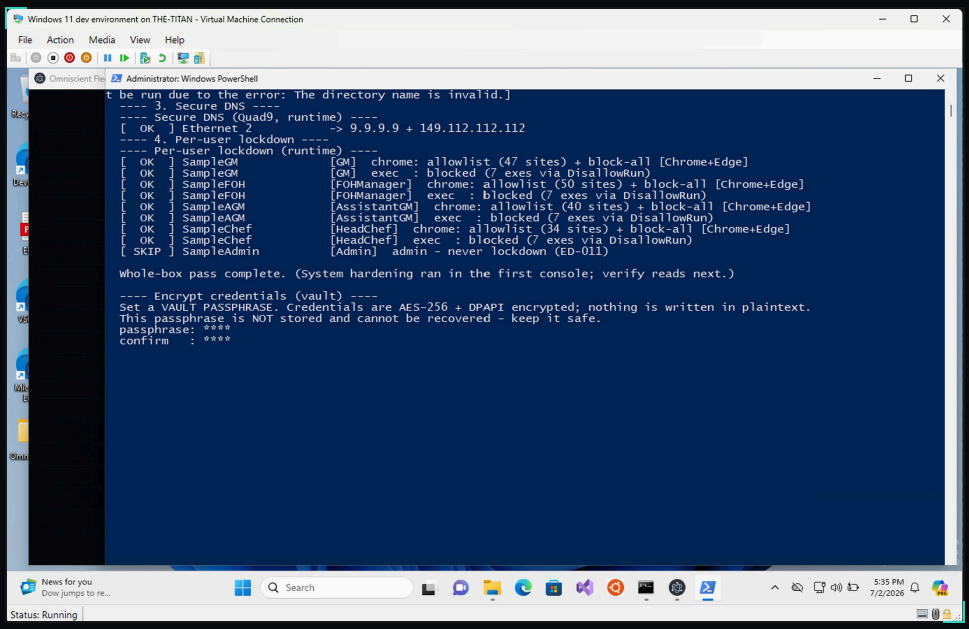

FIG.20 Vault creation at the tail of a hardening pass — AES-256 plus DPAPI, with a passphrase that is explicitly not stored and not recoverable. The credentials the platform just touched are sealed immediately.

[08 // OPERATE]

ONE ANALYST, MAGNIFIED

This is the section a reseller should read twice. Omniscient is built so that **one operator covers a fleet that used to need a team** — and that multiplier is where your margin comes from.

Acting as a specific user for support or investigation is a single elevation, not a UAC prompt on every click — one authorization covers the whole visit, logged as one session. Live policy state is read straight off the box and compared against the model; if they diverge, that's drift, surfaced the same way network drift is. Nothing forces the operator to hunt box-by-box: the whole fleet's alarms and messages land in one triage feed.

- **One elevation per visit** — not per action. The friction that slows down manual operators is designed out.
- **Drift surfaced, not hunted** — the model tells the operator what diverged, instead of the operator discovering it by hand.
- **A blocked action is a pass, not a bug** — when a locked-down app refuses to launch, that failure *is* the lockdown working, and the platform treats it as confirmation.
- **One feed for everything** — security events and user messages triaged separately so nothing critical hides under routine traffic.

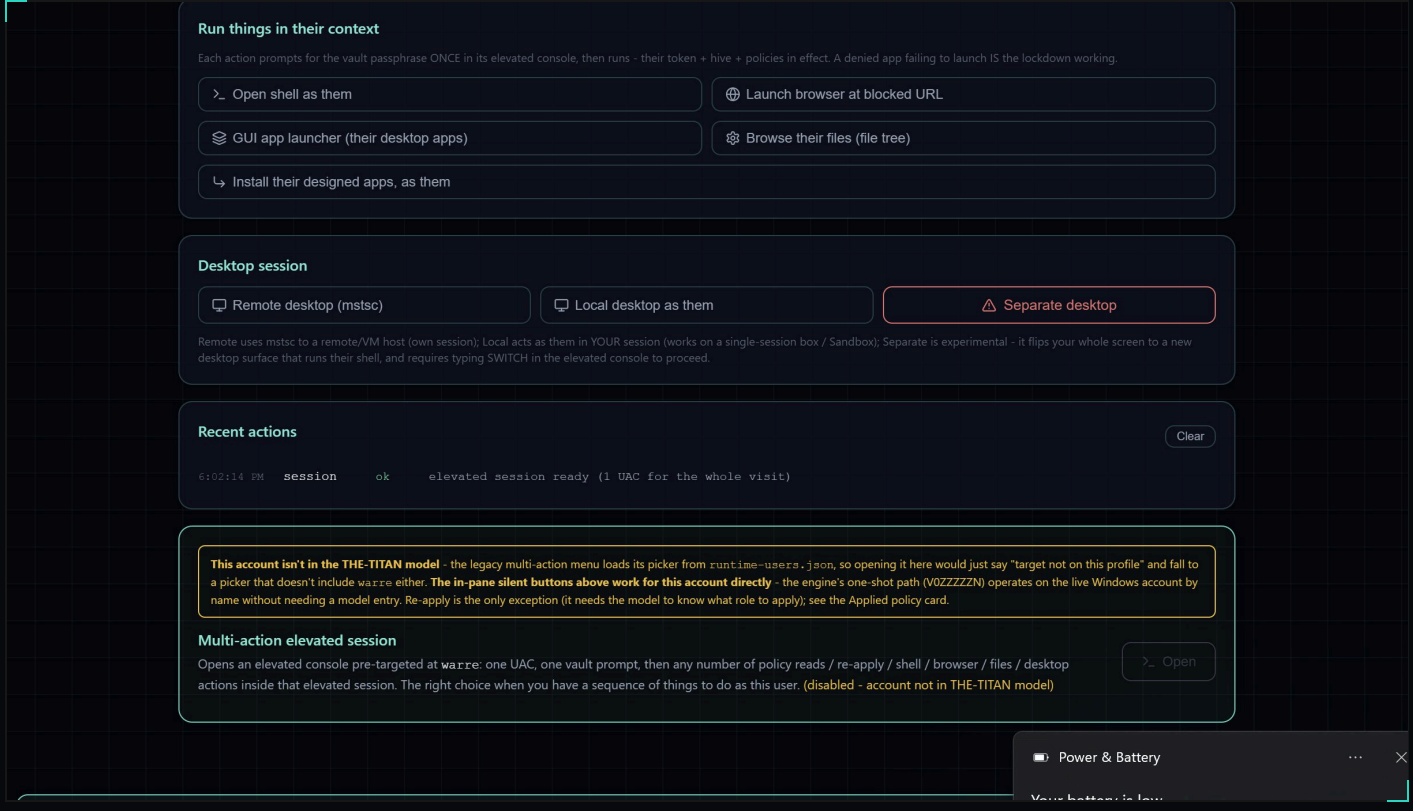

FIG.16 Running actions in a specific user's context — shell, browser, files, installs — all under that account's real token, hive, and effective policy, inside one authorized session. One analyst, operating precisely, at fleet scale.

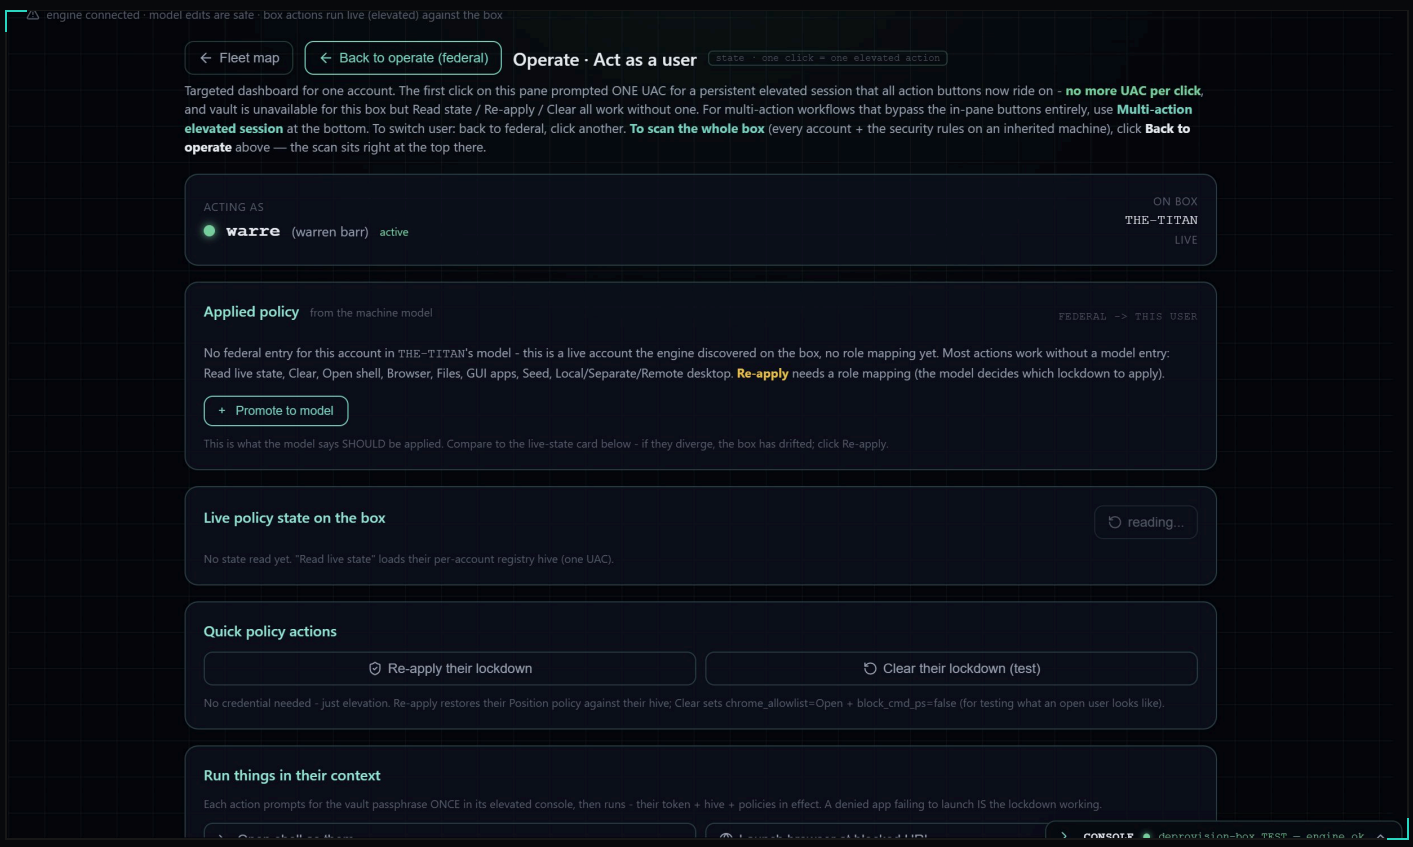

FIG.15 The act-as-user pane — live policy compared against the model, a one-click promote path for ungoverned accounts, and the whole Federal-to-State inheritance visible at a glance.

[09 // SENTINEL] — THE DIFFERENTIATOR

IT FIGHTS BACK. SAFELY.

Hardening keeps an attacker out. Deception is what happens when one gets in anyway — and it changes the **economics** of attacking your clients.

Think about an intruder's cost. Today, once they're inside, their only real cost is **time**. They can move laterally, poke at files, and explore at leisure, because nothing punishes movement. Omniscient inverts that. Every step an attacker takes through a protected environment raises their odds of triggering an automated, recorded containment event. You aren't just detecting them — you're making the environment too expensive to navigate. That's a product an MSSP can sell on outcome, not jargon.

> **THE INSIGHT — IT'S ABOUT THE IMMEDIATE TRIPWIRE** ◆
>
> Security isn't about stopping every intrusion; it's about guaranteeing that the instant an attacker moves, they hit something. Sentinel is that something — a layer of live decoys and a signed liveness check that turn quiet lateral movement into an immediate, loud, automatic problem for the attacker.

MECHANISM — ABSENCE AS THE ALARM

Every managed endpoint writes a signed, sequenced heartbeat on a fixed interval. The clever part isn't the beat — it's what its **absence** means. An intruder can kill the agent, but they cannot forge the next beat; there is no signing key on the box to steal. Silence isn't a gap in coverage. Silence *is* the alarm — the one signal an attacker can't fake their way around.

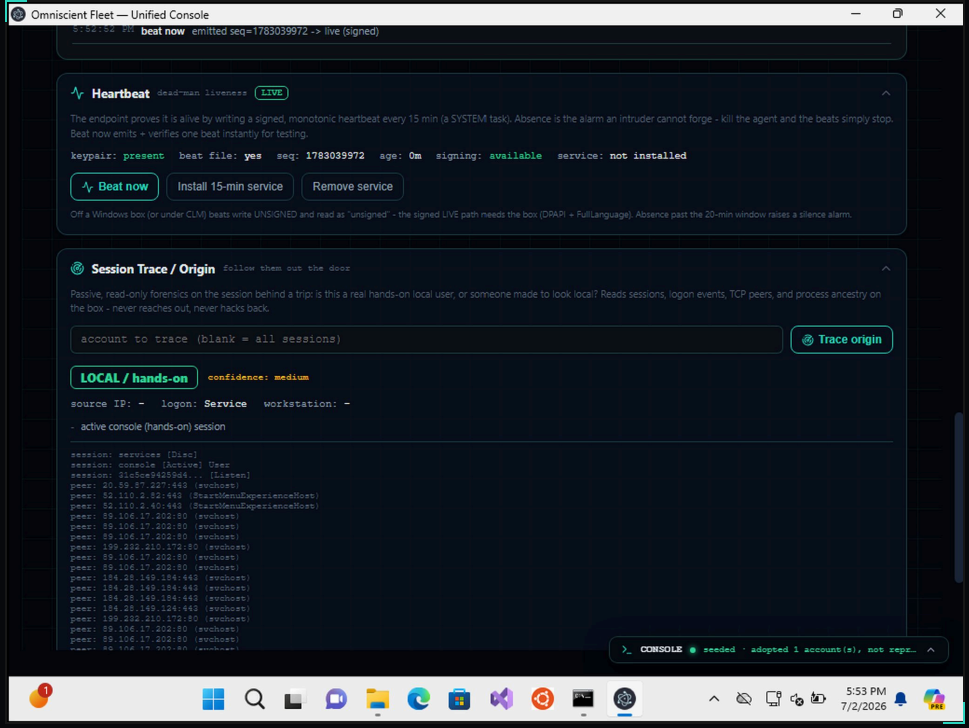

FIG.30 The signed heartbeat panel alongside passive session-origin tracing — live proof-of-life on the left, and read-only forensics on whoever tripped a mine on the right. It never reaches back out.

THE SAFETY PROTOCOL — THE CONTAINMENT LADDER

Touch a decoy and a response fires automatically. But it isn't one blunt action — it's a ladder, and **where the decoy sat decides how hard the system hits back.** The deeper the zone an attacker reached, the higher the rung.

T0	**LOG**	Chain-of-custody entry. Always, no exceptions.
T1	**ALARM**	Warning-level alert the instant the decoy is touched.
T2	**ESCALATE**	Critical alarm, flagged for off-box escalation.
T3	**REVOKE**	Every live just-in-time access grant on the box is pulled immediately.
T4	**ISOLATE**	Closes the exposure the platform itself opened — never the operator's console, never the account.
T5	**DISABLE**	Isolates and disables the implicated account — but never the operator's own, and never the last admin.

Notice what the ladder **doesn't** do. It never destroys data. It never counter-attacks. Its most aggressive rung refuses to lock out the last way back in. That restraint is deliberate: this runs unattended on a client's production box, so its worst-case action has to be one you'd still be comfortable with at 3 a.m. with no one watching. Containment, not vengeance.

FOR THE BUYER WHO ASKS "CAN WE TUNE IT?" ◆

Yes. The ladder ships pre-configured for safety, but every rung is configurable to match each client's risk appetite. A cautious client can cap the response low; a high-security client can arm the full ladder. Same platform, tuned per engagement.

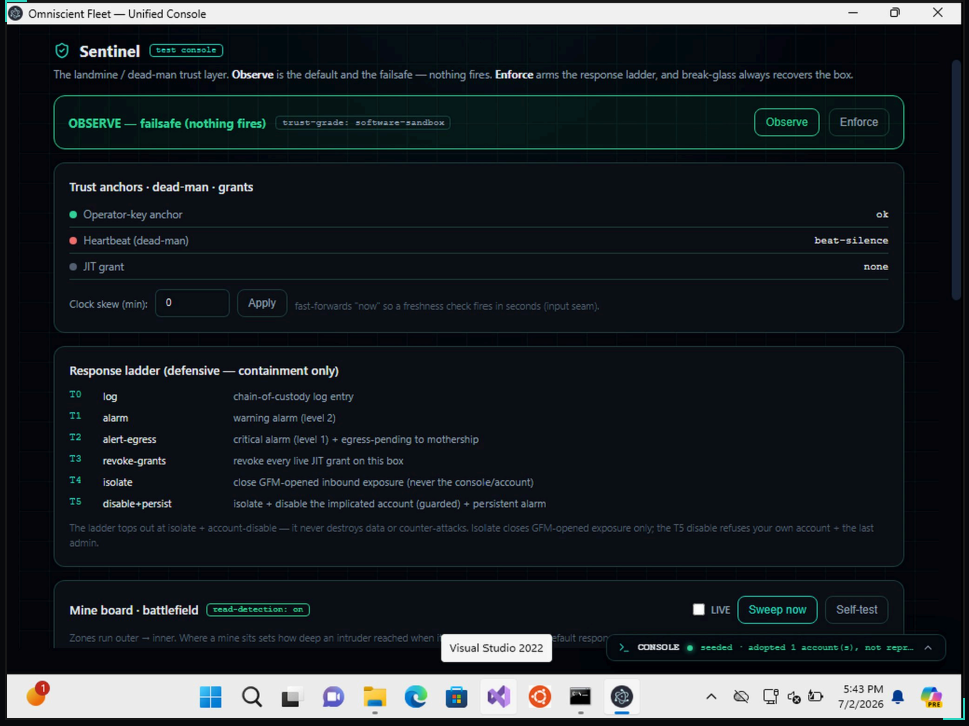

FIG.23 Observe versus Enforce, with the full containment ladder rendered exactly as the engine executes it. Observe is the default failsafe — nothing fires until an operator arms it.

THE MINEFIELD — FIVE ZONES, SCALED RESPONSE

Decoys are planted directly into the live filesystem — not a simulated overlay — across five zones that mirror how far an intruder would have had to travel to reach them: perimeter bait, user data, credentials, infrastructure, and crown jewels. Each zone maps to a default response tier, so blast radius scales with intrusion depth automatically.

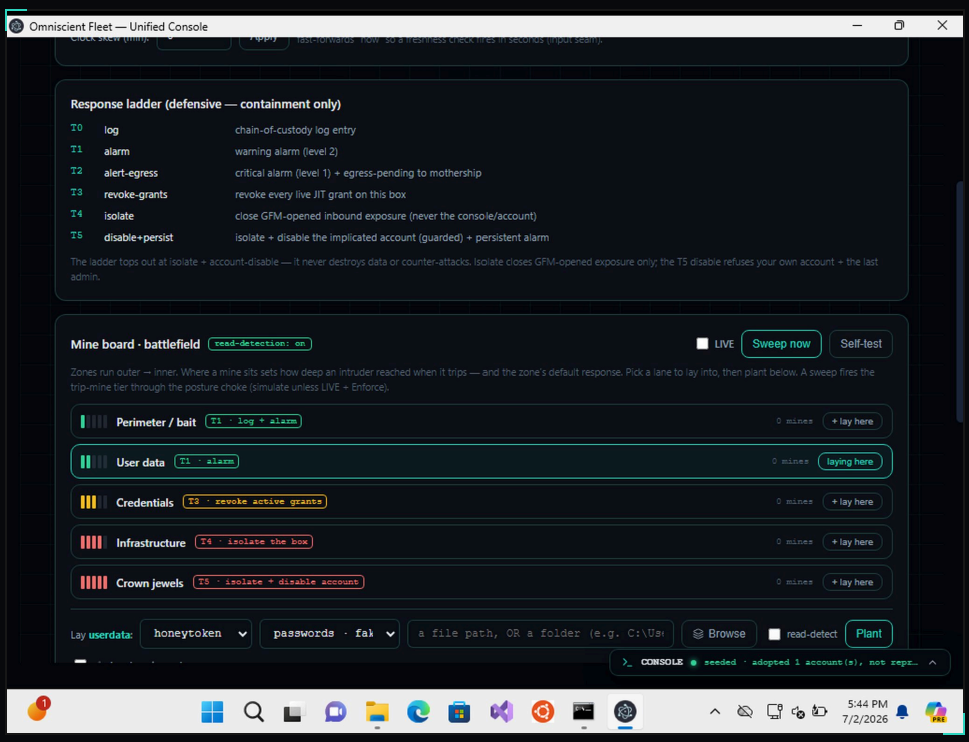

FIG.24 The mine board — five zones, five default response tiers. A decoy in perimeter bait raises an alarm; a decoy in crown jewels triggers isolation and account disable. The operator decides where to lay each one.

DESIGN NOTE — A TRAP THAT CAN'T BE QUIETLY RESET

A decoy an intruder can trip and then un-trip is worthless. Sentinel computes trip-state live, from the file's actual condition against a sealed baseline — there is no "tripped" flag to flip back. Clearing an incident doesn't erase it; it archives the evidence and plants a brand-new decoy in its place. The record survives everything the attacker can do to the box.

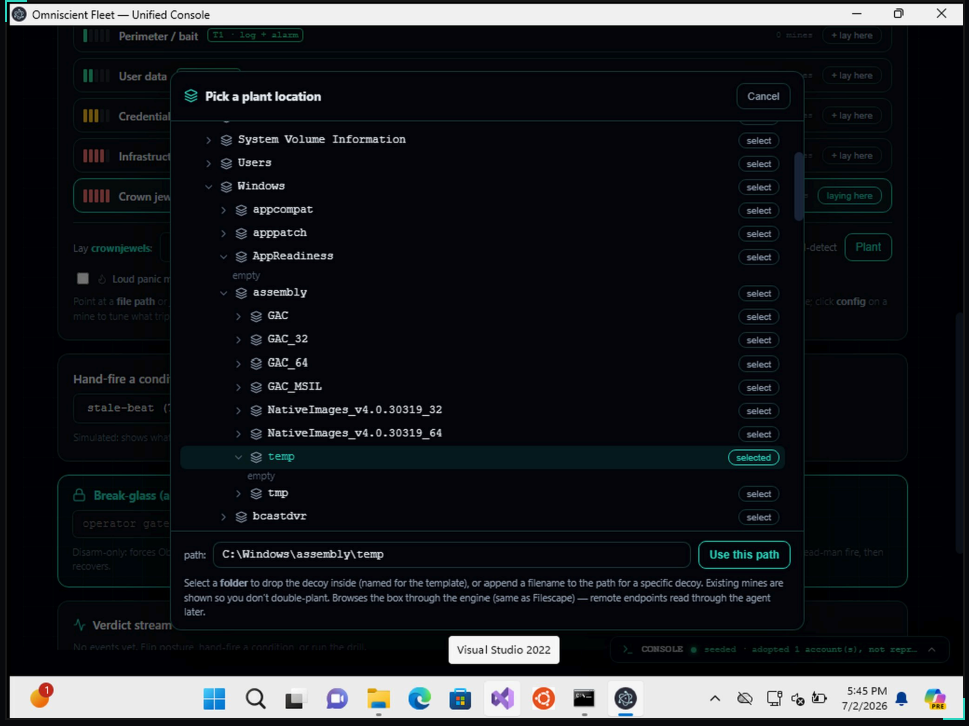

FIG.25 Placing a decoy through Filescape — a live, navigable view of the box's real filesystem. Existing mines show inline so an operator never double-plants, and decoys drop exactly where an intruder would look.

PSYCHOLOGICAL CONTAINMENT

When a loud mine trips, the response isn't only technical. The intruder's own session is taken over and told, in the most unambiguous terms the screen can manage, that they've been caught and the response has already started. This is psychological containment — you break the attacker's momentum by showing them, mid-move, that they're inside a trap that's already closing.

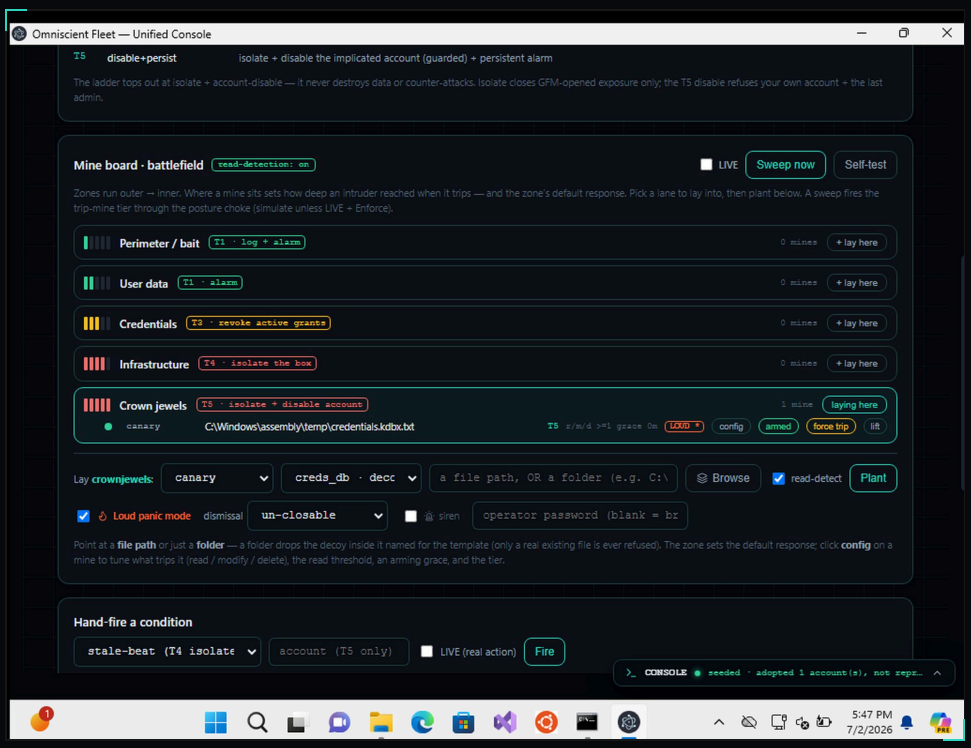

FIG.26 A live, armed decoy in the Crown Jewels zone — trip conditions, response tier, and loud-or-silent mode all set by the operator before it goes hot.

LIVE SCENARIO — 17:49, UNATTENDED

A decoy credential file in the crown-jewels zone is modified. No analyst is watching. The heartbeat and mine engine act on their own: the incident is logged with a tracked ID, the acting account and process are captured, grants are revoked, the session is contained, and the alarm is written to the feed. By the time an operator next opens the console, the intrusion is already contained and fully documented — a clean chain of custody, no 3 a.m. phone call required.

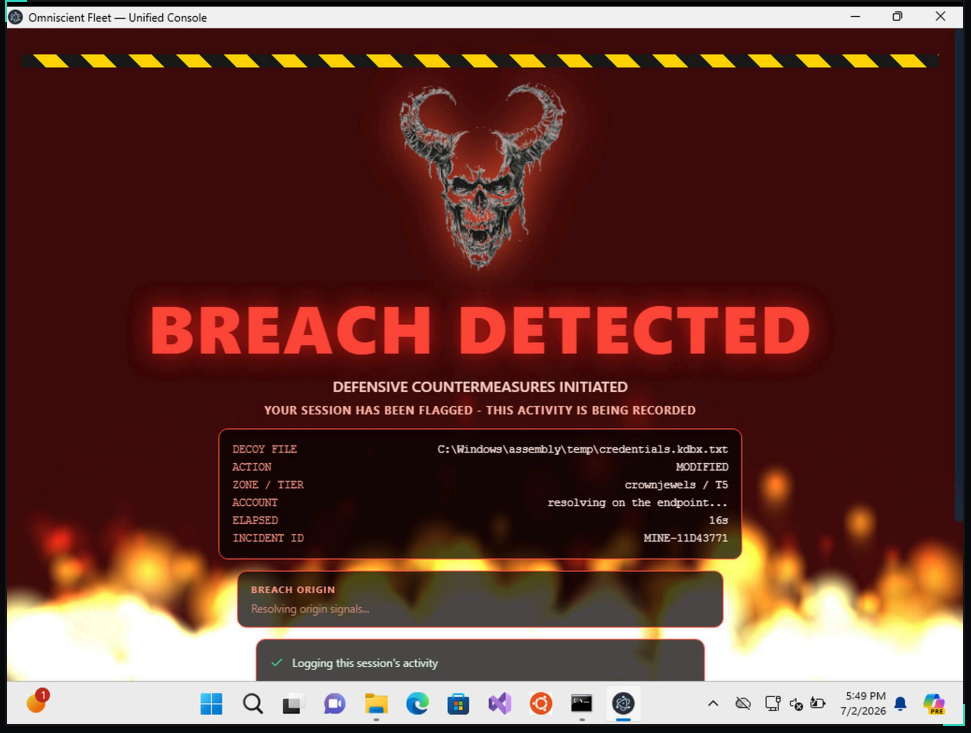

FIG.27 The moment of detection, as the intruder's own session sees it — a full-screen takeover carrying a tracked incident ID while the containment ladder fires underneath. Momentum broken.

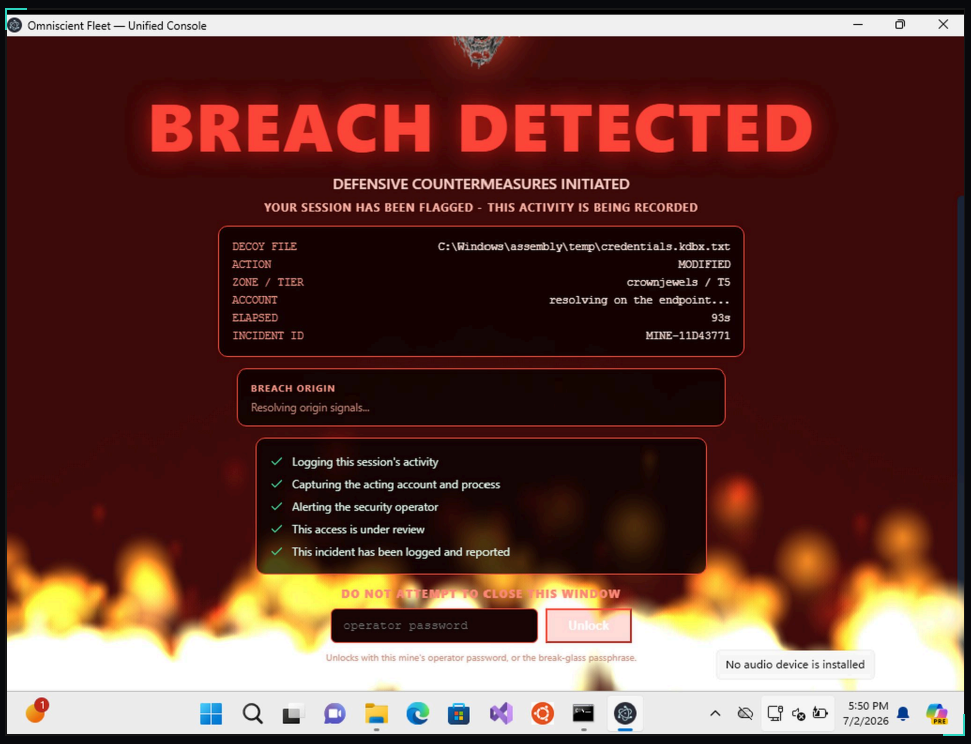

FIG.28 The response executing in real time — logged, acting account captured, operator alerted, incident recorded — and locked behind an authority key the intruder does not have.

[10 // ADVERSARY]

WHAT IT TAKES TO BEAT IT

An honest look at the attacker's problem — because a claim of "unhackable" is what amateurs sell, and the buyers you want will see through it instantly.

Against a fully-hardened, Sentinel-armed environment, a remote attacker doesn't get to fight the cryptography — they get pushed off it entirely:

- **No forged grants.** Access grants are signed and context-bound; without the private key, a fabricated grant is rejected and raises an alarm.
- **No faked liveness.** The heartbeat is RSA-signed with no key on the box to steal — an attacker can stop it, but stopping it is itself the alarm.
- **No brute-forced vault.** PBKDF2 at 600,000 iterations means the credential vault isn't cracked in any timeframe that matters.
- **No quiet exploration.** The decoy layer means lateral movement is the risk, not the reward — every step toward something valuable is a step toward a tripwire.

So the attacker is forced off the math and onto the two most expensive problems in security: defeating a disciplined operator, and moving through a monitored environment without tripping anything while fully informed. That's not impossible — nothing is — but it is **slow, loud, and expensive**, which is the entire point. You have changed the economics. An environment that used to cost an attacker time now costs them risk on every move.

THE RESELLER'S VERSION OF THIS ARGUMENT

You don't sell "unbreakable" to your clients. You sell "we make attacking you more expensive than attacking someone else" — and you can back it with a live demonstration, because Sentinel's response is real and repeatable, not a slide.

[11 // PROGRAM]

NOT A SNAPSHOT. A PROGRAM.

A hardening report is a snapshot — stale the day after it's written. A deception layer is only as credible as the offense actually thrown at it. So validation isn't a one-time event; it's a **continuous purple-team loop**.

Omniscient ships with its own red-team capability, wired to the same engine as everything else, and enrollment is structured as a continuation program rather than a project with an end date. Under signed rules of engagement and explicit written authorization — enforced in the platform, not just promised in a contract — the exact hardening deployed to a fleet is tested against, and every finding feeds the next hardening pass.

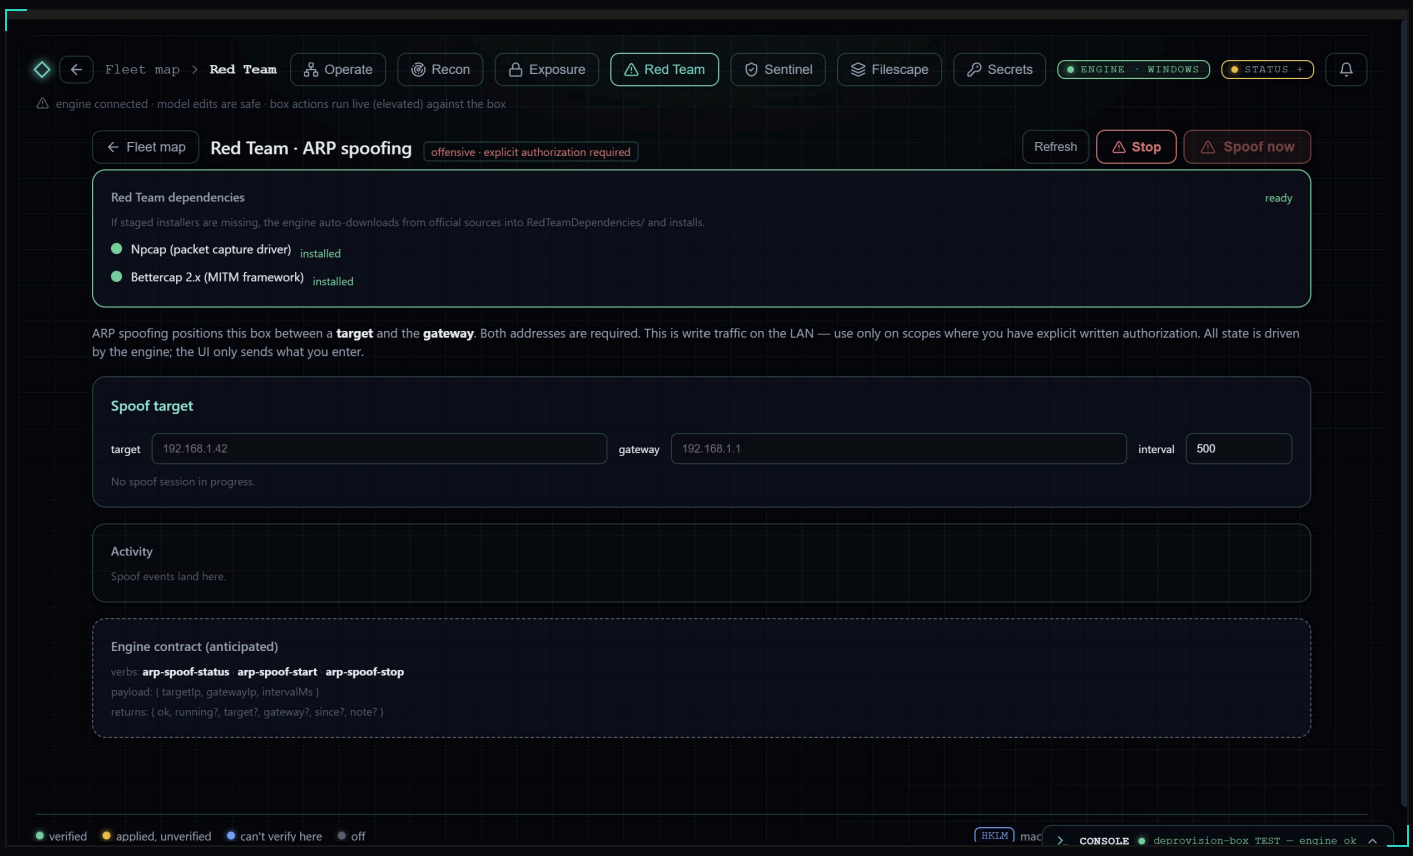

FIG.17 The red-team module — adversarial validation run under explicit authorization against the platform's own hardening. Offensive capability with rules-of-engagement built into the tool itself.

For a reseller, the continuation program is what turns a one-time sale into recurring revenue: the client's posture doesn't get set once and left to drift — it gets tested, tuned, and re-tested on a standing cadence, under your name, for the life of the contract.

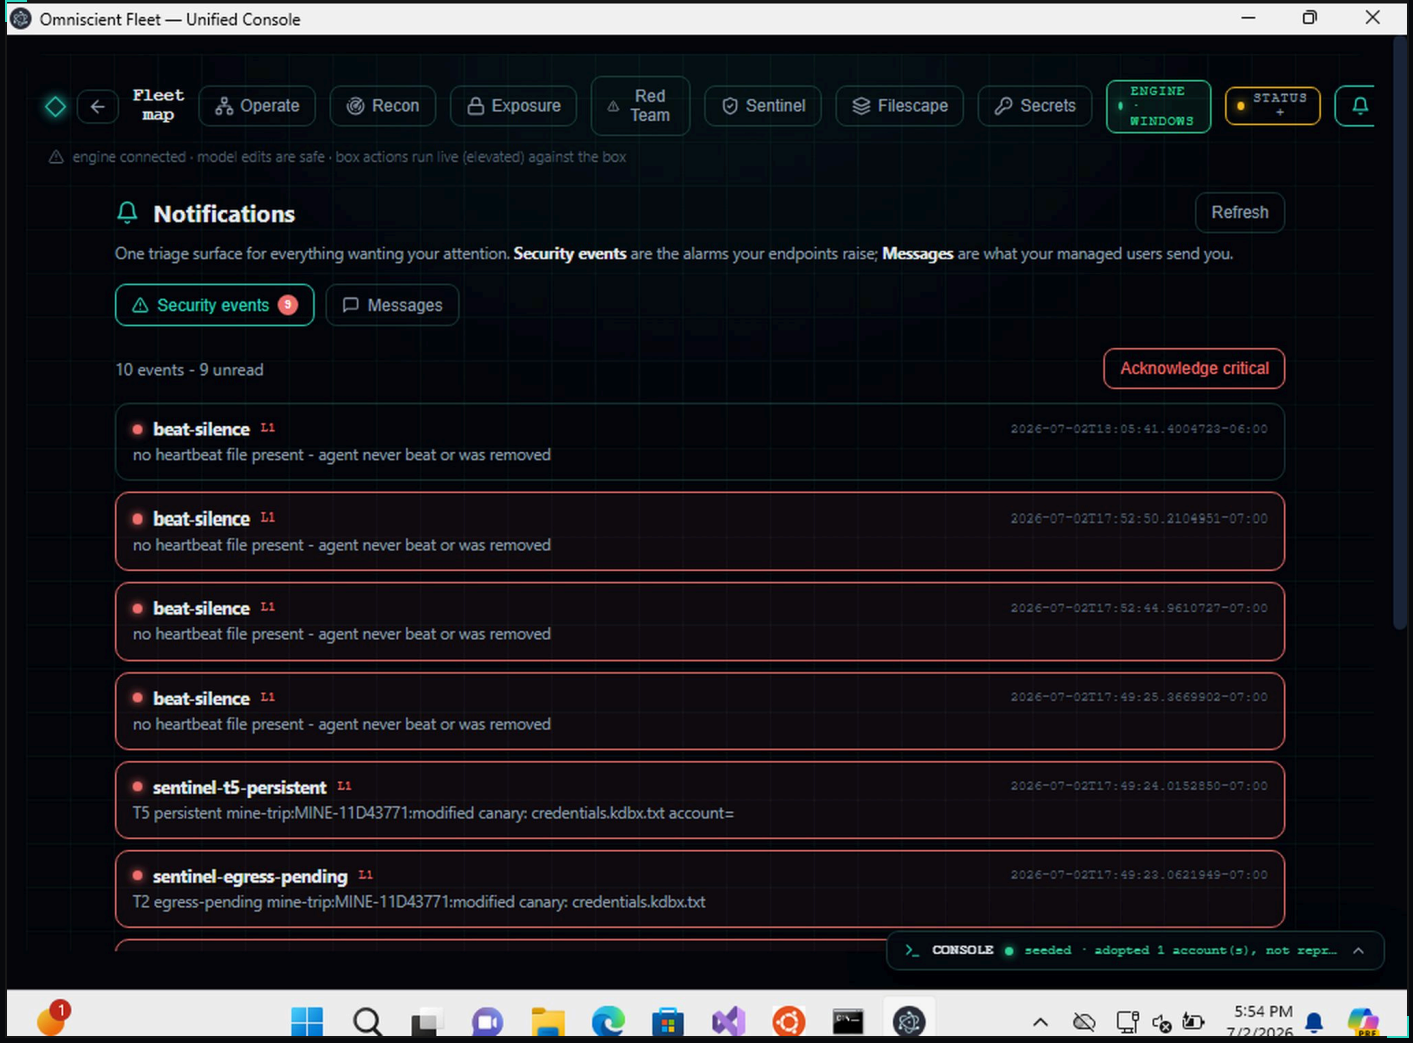

FIG.31 The unified triage feed — a live mine-trip incident and repeated heartbeat-silence alarms landing in one place, on the same incident thread that ran through the breach screen. One stream, end to end.

[12 // SCALE]

ADD A CLIENT, NOT A REBUILD

The economics that make Omniscient worth reselling come from one architectural fact: because every capability is driven by the same data model, **growth is additive, not structural.**

Adding a machine is adding a profile. Adding a client is adding a set of profiles — not re-architecting anything. Role templates fold down to every user automatically, presets carry the hardening baseline, and the same engine drives one endpoint or a thousand. The operator's console doesn't change as the fleet grows; only the count on the fleet map does.

GROWS BY	FOLDS VIA	RUNS ON
PROFILES	**TEMPLATES**	**ONE ENGINE**
A new machine or client is data added to the model, not new code or a new console. Onboarding is configuration, not construction.	Role and preset templates cascade to every user and machine they touch. Change one, it ripples correctly to all — idempotently.	The same engine drives one box or the whole estate. Your operators learn one tool and apply it to every client you take on.

That's the reseller's growth curve: take on more clients without a linear increase in headcount or a new toolchain per engagement. The platform's discipline — idempotent, reversible, what-if-able — is precisely what makes operating many clients at once safe rather than chaotic.

FIG.07 The topology as the fleet grows — new machines join the map as profiles, drift counters update live, and the operator's view stays exactly as legible at scale as it was with a single box.

[13 // ENGAGE]

TWO WAYS TO PUT OMNISCIENT TO WORK

// FOR MSSPS

LICENSE & RESELL

Run Omniscient as the operator platform behind your own practice. One analyst covers a fleet that used to need a team; onboarding a client is configuration, not construction; and the continuation program turns one-time hardening into recurring revenue — under your brand, on your terms.

Licensing terms, operator training, and the boundaries of a client-facing view are scoped with you directly.

// FOR BUSINESSES

ENROLL FOR FULL CARE

Have your fleet run through Omniscient by Barr Cyber directly — discovery, hardening, deception, and continuous purple-team validation, under one contract. The same discipline applied to the platform's own environment, applied to yours, on a standing cadence.

Scope is set to your environment and risk appetite. Nothing offensive ever runs without written authorization on file.

START THE CONVERSATION.

Whether you're evaluating Omniscient to resell or to be protected by it, the first step is a direct conversation about your environment — not a sales funnel.

WARREN BARR · 713-882-0902 · warren@barr-cyber.com · barr-cyber.com

[APX // GALLERY]

FULL PLATFORM GALLERY

Every screen referenced in this dossier, plus the supporting views not called out individually — provisioning, allowlisting, access, and governance, shown as captured.

FIG.01 Fleet map overview

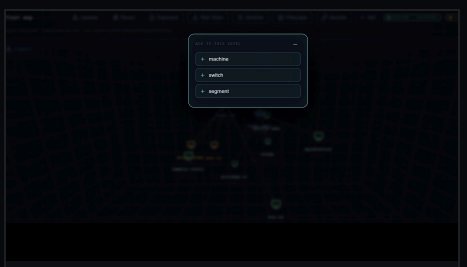

FIG.02 Add machine / switch / segment

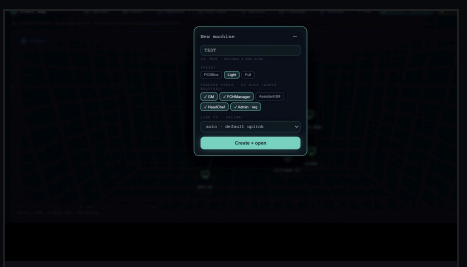

FIG.03 New machine provisioning

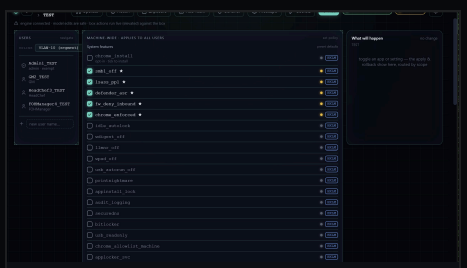

FIG.04 Machine-wide hardening toggles

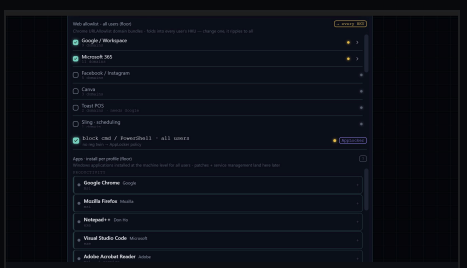

FIG.05 Web allowlist & app installs

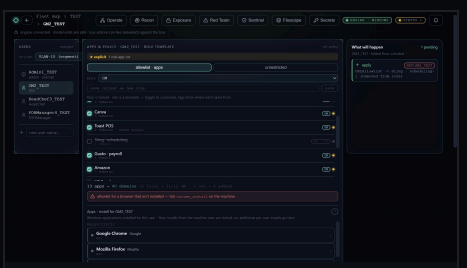

FIG.06 Per-role policy template

FIG.07 Topology with live node

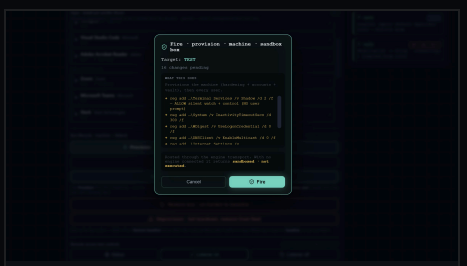

FIG.08 Pre-flight provisioning review

FIG.09 Live execution log

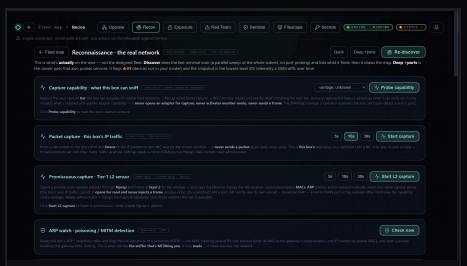

FIG.10 Reconnaissance overview

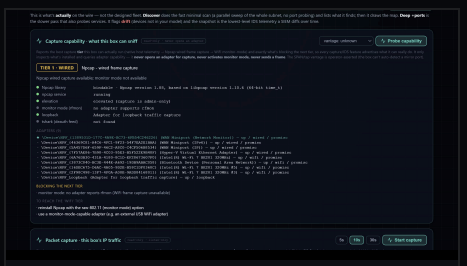

FIG.11 Capture-capability probe

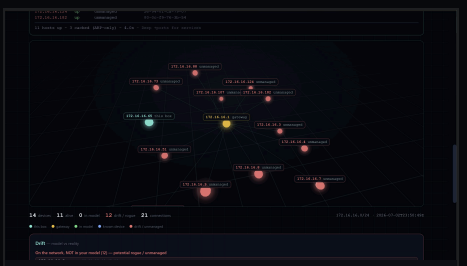

FIG.12 Discovery & drift map

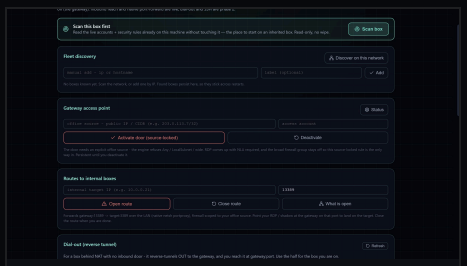

FIG.13 Gateway access door

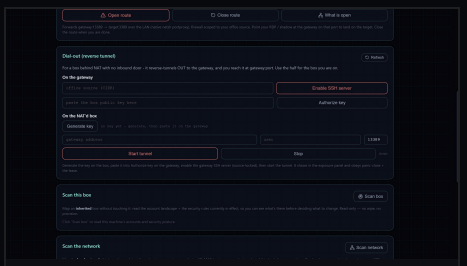

FIG.14 Dial-out reverse tunnel

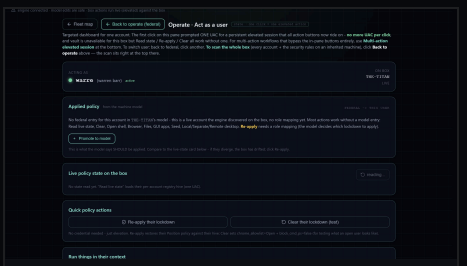

FIG.15 Act-as-user policy pane

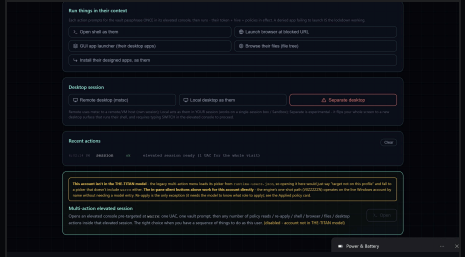

FIG.16 Run in user context

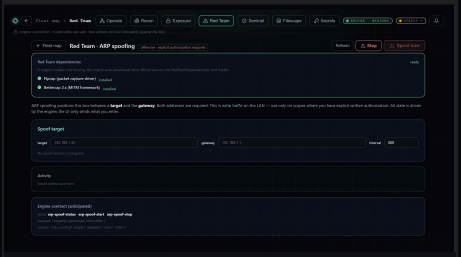

FIG.17 Red-team validation

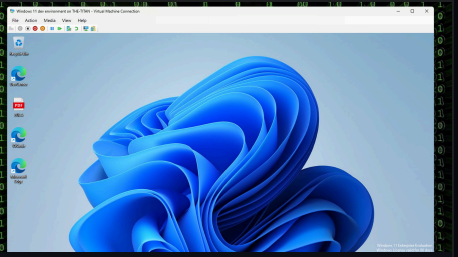

FIG.18 Baseline environment

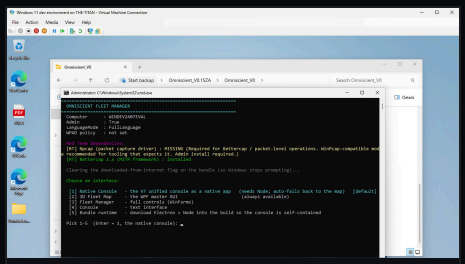

FIG.19 Platform launcher

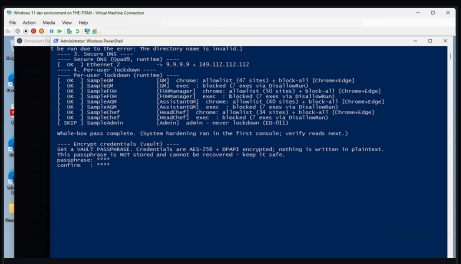

FIG.20 Vault creation

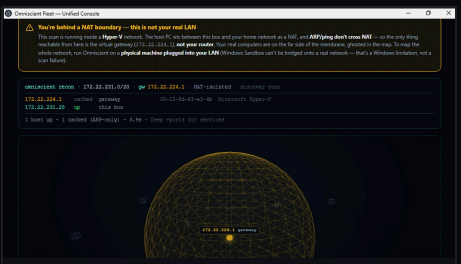

FIG.21 NAT-aware recon, 3D

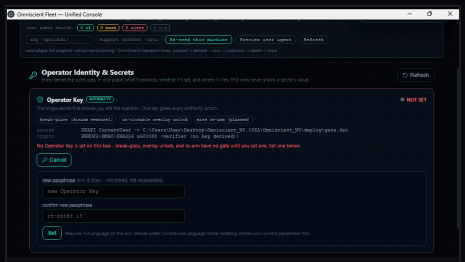

FIG.22 Secrets — two keys

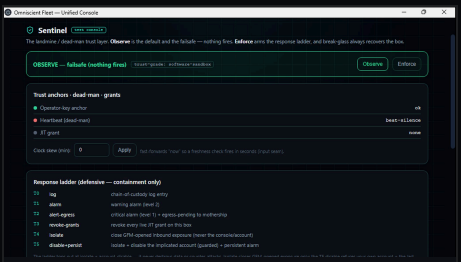

FIG.23 Sentinel response ladder

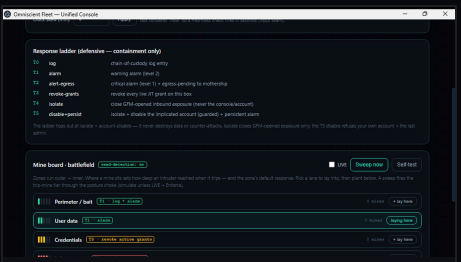

FIG.24 Mine board — five zones

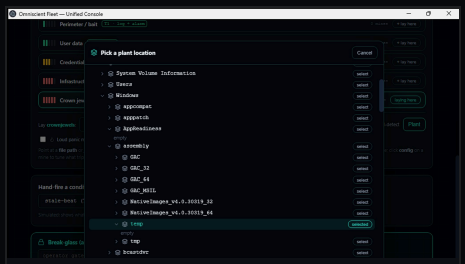

FIG.25 Filescape decoy placement

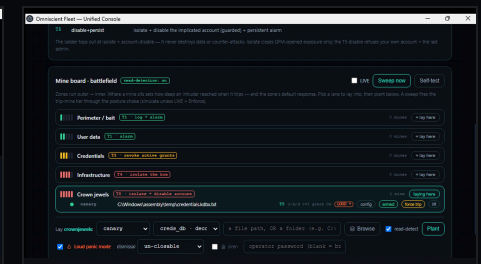

FIG.26 Live armed decoy

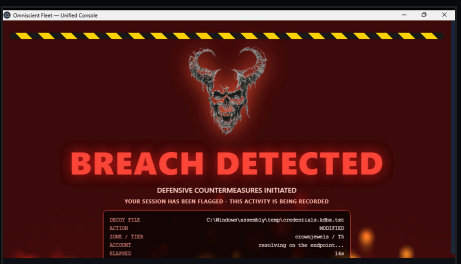

FIG.27 Breach detection takeover

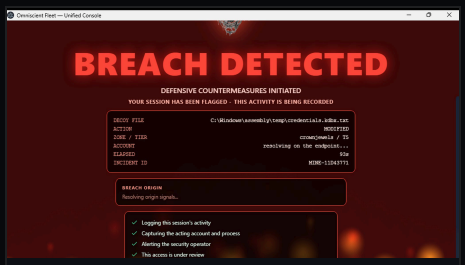

FIG.28 Live containment checklist

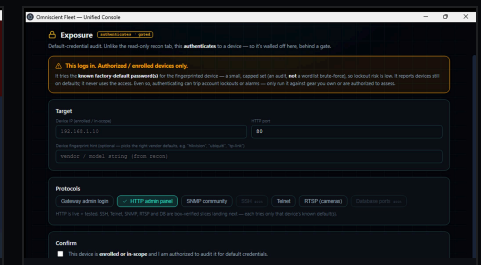

FIG.29 Default-credential audit

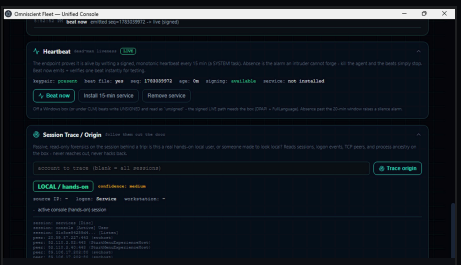

FIG.30 Heartbeat & session trace

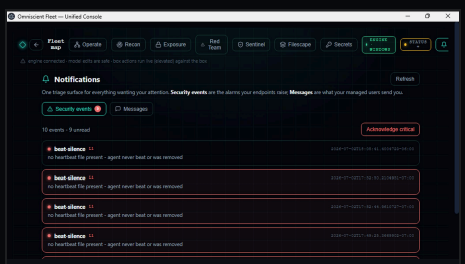

FIG.31 Unified triage feed