SAMPLE REPORT · ANONYMIZED

See what an attacker already knows — before the strike.

Below is a **real, authorized HYDRA engagement** with the client fully anonymized — same depth, same discipline, zero PII. This is exactly the kind of report you receive when you work with us.

Passive only · Receipts-first · GDPR · NDA-ready · 2-3 days

⚠ **SAMPLE · ANONYMIZED · ILLUSTRATIVE.** This is a real HYDRA engagement with the client identity, domains, IP addresses and personnel **fully redacted**. The structure, methodology, lane logic and finding **types** are unchanged — this is exactly what a HYDRA report delivers. No live target is identified.

TIER 3 · FULL INTELLIGENCE · PASSIVE EXTERNAL EXPOSURE

Sample Engagement Report

HYDRA by TIA

Authorized passive external exposure analysis — 5-lane HYDRA methodology, evidence-first.

Method: Passive OSINT Lanes: 5 Findings: 12 Basis: real engagement, de-identified

ENGAGEMENT · SCOPE

Client: [REDACTED] — a European small-to-mid business running several brands and a routine supplier-payment workflow. Sector, location, size, domains, IPs and people are generalized or redacted for this sample.

Authorization: written owner consent (assessment of the client's own public exposure). **Method:** 5-lane HYDRA — PASSIVE only (DNS/CT/DoH, public web, response headers, RDAP/whois, OSINT). No active scanning, no authentication attempts, no exploitation. Receipts-first, zero fabrication.

Declared up front: account-level breach confirmation and active characterization of the on-prem asset are not possible passively → scoped to a paid, signed Phase 2.

EXTERNAL THREAT EXPOSURE

ELEVATED

Four critical exposures on an otherwise healthy perimeter. The realistic money-loss path — **email spoofing → fake supplier invoice** — was confirmed independently across three lanes.

12	4	5	Passive
FINDINGS	CRITICAL / HIGH	RECON LANES	METHOD — 0 TOUCH

EXECUTIVE SUMMARY — 4 REAL RISKS

The rest of the perimeter is reassuringly healthy (see the honest-negatives section). Four exposures carry real, exploitable risk:

1. **On-prem remote-access box exposed to the internet.** A management endpoint on consumer broadband, end-of-life web server, self-signed certificate — the single most attackable real asset. CRITICAL
2. **Email authentication not enforced on both flagship domains** (DMARC p=none). This opens a *credible impersonation path* — mail that can appear to come from leadership or a team address, on the real domain → Business Email Compromise / fake-invoice. *Corroborated across infra + dark-int + people lanes.* CRITICAL
3. **Third-party look-alike domain with live mail.** A recently-registered near-spelling of the brand, behind a privacy proxy, with active mail records — a ready-made phishing channel impersonating the client. CRITICAL
4. **Shared role mailboxes; MFA posture not verifiable passively.** A credential-hygiene pattern that, if unaddressed, makes risk #2 far easier to weaponize. HIGH

FINDINGS BY SEVERITY

Lanes — INFRA infrastructure · DARK-INT breach & dark-web · WEB web surface · PEOPLE human/OSINT · BRAND domains & impersonation. **Severity** — CRITICAL / HIGH exploitable now · MEDIUM hardening.

F1 On-prem remote-access endpoint exposed

[INFRA]

A deliberate DNS record points a management hostname at `consumer broadband (not a datacenter)`. Passive fingerprint: end-of-life web server, self-signed certificate (multi-year). Consistent with an on-prem management / VPN / back-office box — internet-reachable and unadvertised in certificate transparency; its internal network adjacency is *inferred, not passively confirmed.* → identify it, place behind a zero-trust broker, patch/replace, pull off the public net. *What it actually runs = Phase 2 (signed scope).*

F2 DMARC p=none on both flagship domains → BEC

[INFRA · DARK-INT · PEOPLE]

Both flagship domains publish `p=none` (monitor, no enforcement) despite valid SPF and DKIM. Mail can be sent as `From: <anyone>@[client-domain]` with no cryptographic rejection — a *credible impersonation path* (final delivery still depends on the recipient's filtering, but the sender's own domain raises no barrier). For any business that pays suppliers, this is a fake-invoice / payment-redirect vector. → add aggregate reporting, confirm sender alignment, then move `none → quarantine → reject` on both.

F3 Look-alike domain, mail-capable

[BRAND · INFRA · WEB]

A near-spelling of the brand, registered ~6 months ago behind a US privacy proxy, on a different hosting stack than the client — every fingerprint diverges, so it is *not* theirs. Storefront disabled, but `live MX records` present: a disabled store ≠ disabled mail = a ready phishing/BEC channel. → attempt acquisition; else abuse-report + registry dispute; monitor records; warn staff and customers.

F4 Shared role mailboxes + MFA hygiene

[DARK-INT · PEOPLE]

Multiple shared role mailboxes are visible (generic function addresses such as `info@ / orders@`). These tend to be the highest-reuse, lowest-rotation accounts — shared passwords, reused as portal logins; per-user MFA cannot be confirmed passively and is a Phase-2 check. One leaked shared password would let an attacker read supplier correspondence → invoice fraud (amplifies F2). → enforce org-wide 2-step verification; migrate shared → owner+delegation; rotate now.

F5 Expired / wrong-owner TLS + HTTPS→HTTP downgrade

[WEB · INFRA]

Two secondary sites present an expired certificate for an unrelated owner and downgrade `https → plaintext http`. Erodes trust and trains users to click through certificate warnings. → valid certs or TLS-valid edge + clean redirects.

F6 Shop domain with no SPF / no DMARC

[INFRA]

A secondary shop domain is fully spoofable (no SPF, no DMARC) — a clean channel for customer-facing phishing. → `v=spf1 -all` + `p=reject`.

F7 Single publicly-named finance identity

[PEOPLE]

Public professional profiles name exactly one finance decision-maker — both an impersonation *target* and a spoofing *vector* for payment fraud. → authenticate finance comms; callback/codeword for any payment or bank-detail change; reduce public identifiability.

F8 Third-party processor cascade

[DARK-INT]

Several third-party processors sit in the CRM / payments / mail path (customer PII, payment context). A breach at any one leaks client data indirectly; password reuse turns it into a mail-tenant compromise. → inventory staff↔processor logins; unique passwords + MFA; breach-monitor admin accounts.

F9 Content-Security-Policy report-only; weak Referrer-Policy

[WEB · INFRA]

The primary site ships CSP in report-only mode and a permissive Referrer-Policy. → promote CSP to enforced (the sister site already does); tighten Referrer-Policy.

F10 Defensive-registration & listing gaps

[BRAND]

Several obvious brand variants (.com/.eu/.net + alternate spellings) are unregistered and a social handle is unclaimed — open ground for impersonation. → register defensives, claim the handle, verify business listings.

F11 Broken backup MX record [INFRA]

A secondary MX points to a non-resolving host (possibly tied to the F1 on-prem box). → fix or remove.

F12 Account-level breach exposure — UNVERIFIED [DARK-INT]

Free-tier sources show `no public hit` (confirmed negative via dorks + public leak lists), but regional combolists sit behind paid databases. This is absence of public evidence, *not* evidence of absence. → Phase 2 paid pass (breach DBs, Split-Brain on any hit).

HONEST NEGATIVES — WHAT'S ACTUALLY SOLID

A report that is all-red is selling fear, not truth. Here's the genuinely good news, evidence-backed:

- ✓ Managed SaaS platforms eliminate whole classes of self-hosted vulnerability; card data sits with the payment provider.
- ✓ No secrets, stack traces or sensitive paths exposed in page source or robots/sitemap.
- ✓ Mail baseline is healthy — SPF + DKIM present on owned domains; two sub-brands already enforce DMARC.
- ✓ No social-media impersonation found; aggregator listings consistent.
- ✓ The owned domain portfolio is genuinely owned and correctly configured.

REMEDIATION PLAN — SEVERITY-ORDERED

🔴 NOW — QUICK WINS

1. Investigate + isolate the on-prem remote box (zero-trust, patch, off the public net).
2. DMARC `none → quarantine → reject` on both flagship domains (after alignment monitoring).
3. Decide the look-alike domain: acquire, or abuse-report + dispute; warn staff and customers.
4. Enforce org-wide 2-step verification; rotate shared mailbox passwords.
5. Callback-verification + two-person rule for ANY payment or bank-detail change — the single highest-leverage human control, and it defends even against look-alike-domain attacks that DMARC can't stop.

🟡 SOON

1. Fix expired certs / HTTPS downgrade; add SPF + DMARC to the shop domain.
2. Enforce CSP + tighten Referrer-Policy; register defensive domains + claim listings/handle.
3. Inventory staff↔processor logins; unique passwords + MFA everywhere.

🟢 HYGIENE

1. Remove the broken backup MX; consolidate the domain portfolio under one registrar with 2FA + auto-renew.
2. Verify DNSSEC on remaining domains; confirm consent-gating on tracking.

PHASE 2 — WHERE THIS GOES NEXT (PAID, SIGNED SCOPE)

- → Active vulnerability assessment of the on-prem box — what it runs, harden or remove.
- → Paid breach-database pass — confirm or deny account-level credential exposure.
- → DMARC enforcement rollout + staff BEC-awareness training ("finance never changes a bank account by message").
- → Ongoing monitoring — DMARC aggregate reports + brand/domain monitoring.

METHODOLOGY & DISCIPLINE

Receipts-first, zero fabrication. Every 🔴/🟡 is backed by a passive DNS record, RDAP entry, response header or certificate read. Honest negatives are stated as negatives. Explicitly-unverified items (what the on-prem box runs; account-level breach; true owner of the look-alike domain) are flagged as unverified and routed to a signed Phase 2 — never guessed, never inflated.

Passive only. No active exploitation was attempted. All findings derive from publicly available information and open-source intelligence.

HYDRA CUTS HEADS OF DRAGONS — NOT MERCHANTS

We find what an attacker already knows about you, prove it with evidence, and hand you a closed-loop path to fix it. **The first directional signal is free; full proof and remediation are scoped.**