HACKZERO

PENETRATION TEST REPORT · V1

Assessment.

PREPARED FOR
talenthive-eight.vercel.app

ISSUED
2026-06-02

01 · EXECUTIVE SUMMARY

Assessment.

A single page for the leadership read. The verdict, the posture, and the residual risk. Everything that follows substantiates this page.

EXPOSURE

Active critical exposure.

2 critical issues remain open after triage and require immediate attention.

12 OPEN · 0 TRIAGED 0% TRIAGED

Findings at a glance

SEVERITY	TOTAL	OPEN	TRIAGED
CRITICAL	2	2	0
HIGH	3	3	0
MEDIUM	3	3	0
LOW	3	3	0
INFORMATIONAL	1	1	0

Every issue in this report is exploit-validated by Provider's autonomous agents and reproduced at least once before promotion. Each finding has been triaged by a member of talenthive-eight.vercel.app's security team. The decisions, the rationale where required, and the timestamp of each decision are recorded in the Attestation Ledger at the back of this report and constitute the customer's risk-treatment record.

02 · SCOPE AND AUTHORIZATION

In the room.

What we tested, when we tested it, and the methodologies we aligned to. This page is the evidence an auditor reads first for SOC 2 CC4.1 and HIPAA § 164.308(a)(8) purposes.

DOCUMENT	HZ-RPT-0A3CDF4E · v1 Issued 2026-06-02 19:26 UTC
TARGET	talenthive-eight.vercel.app https://talenthive-eight.vercel.app/
ASSESSMENT WINDOW	Started 2026-05-27 03:20 UTC Ended 2026-05-27 03:30 UTC Operation 0a3cdf4e-a519-46b6-9d17-db4bc2916c1f
ENGAGEMENT TYPE	Black-box external penetration test, AI-driven with human oversight Authenticated where credentials were provided in the Rules of Engagement.
METHODOLOGY	OWASP Web Security Testing Guide v4.2 · OWASP Top 10 (2021) · PTES · NIST SP 800-115 Severity classified per CVSS v3.1 mapping; see Methodology page.
VALIDATOR POLICY	Every finding in this report was reproduced by the validator agent at least once before promotion. Unreproducible candidates are dropped silently and do not appear here.
COMPLIANCE OVERLAYS	SOC 2 · HIPAA · PCI DSS · ISO 27001 · GDPR · LGPD
CLASSIFICATION	CONFIDENTIAL For internal use by talenthive-eight.vercel.app and its authorized auditors.

03 · CHAPTER

Method.

How we tested. The six-phase narrative the auditor reads, the methodologies we aligned to, and the validator policy that keeps the report exploit-grounded.

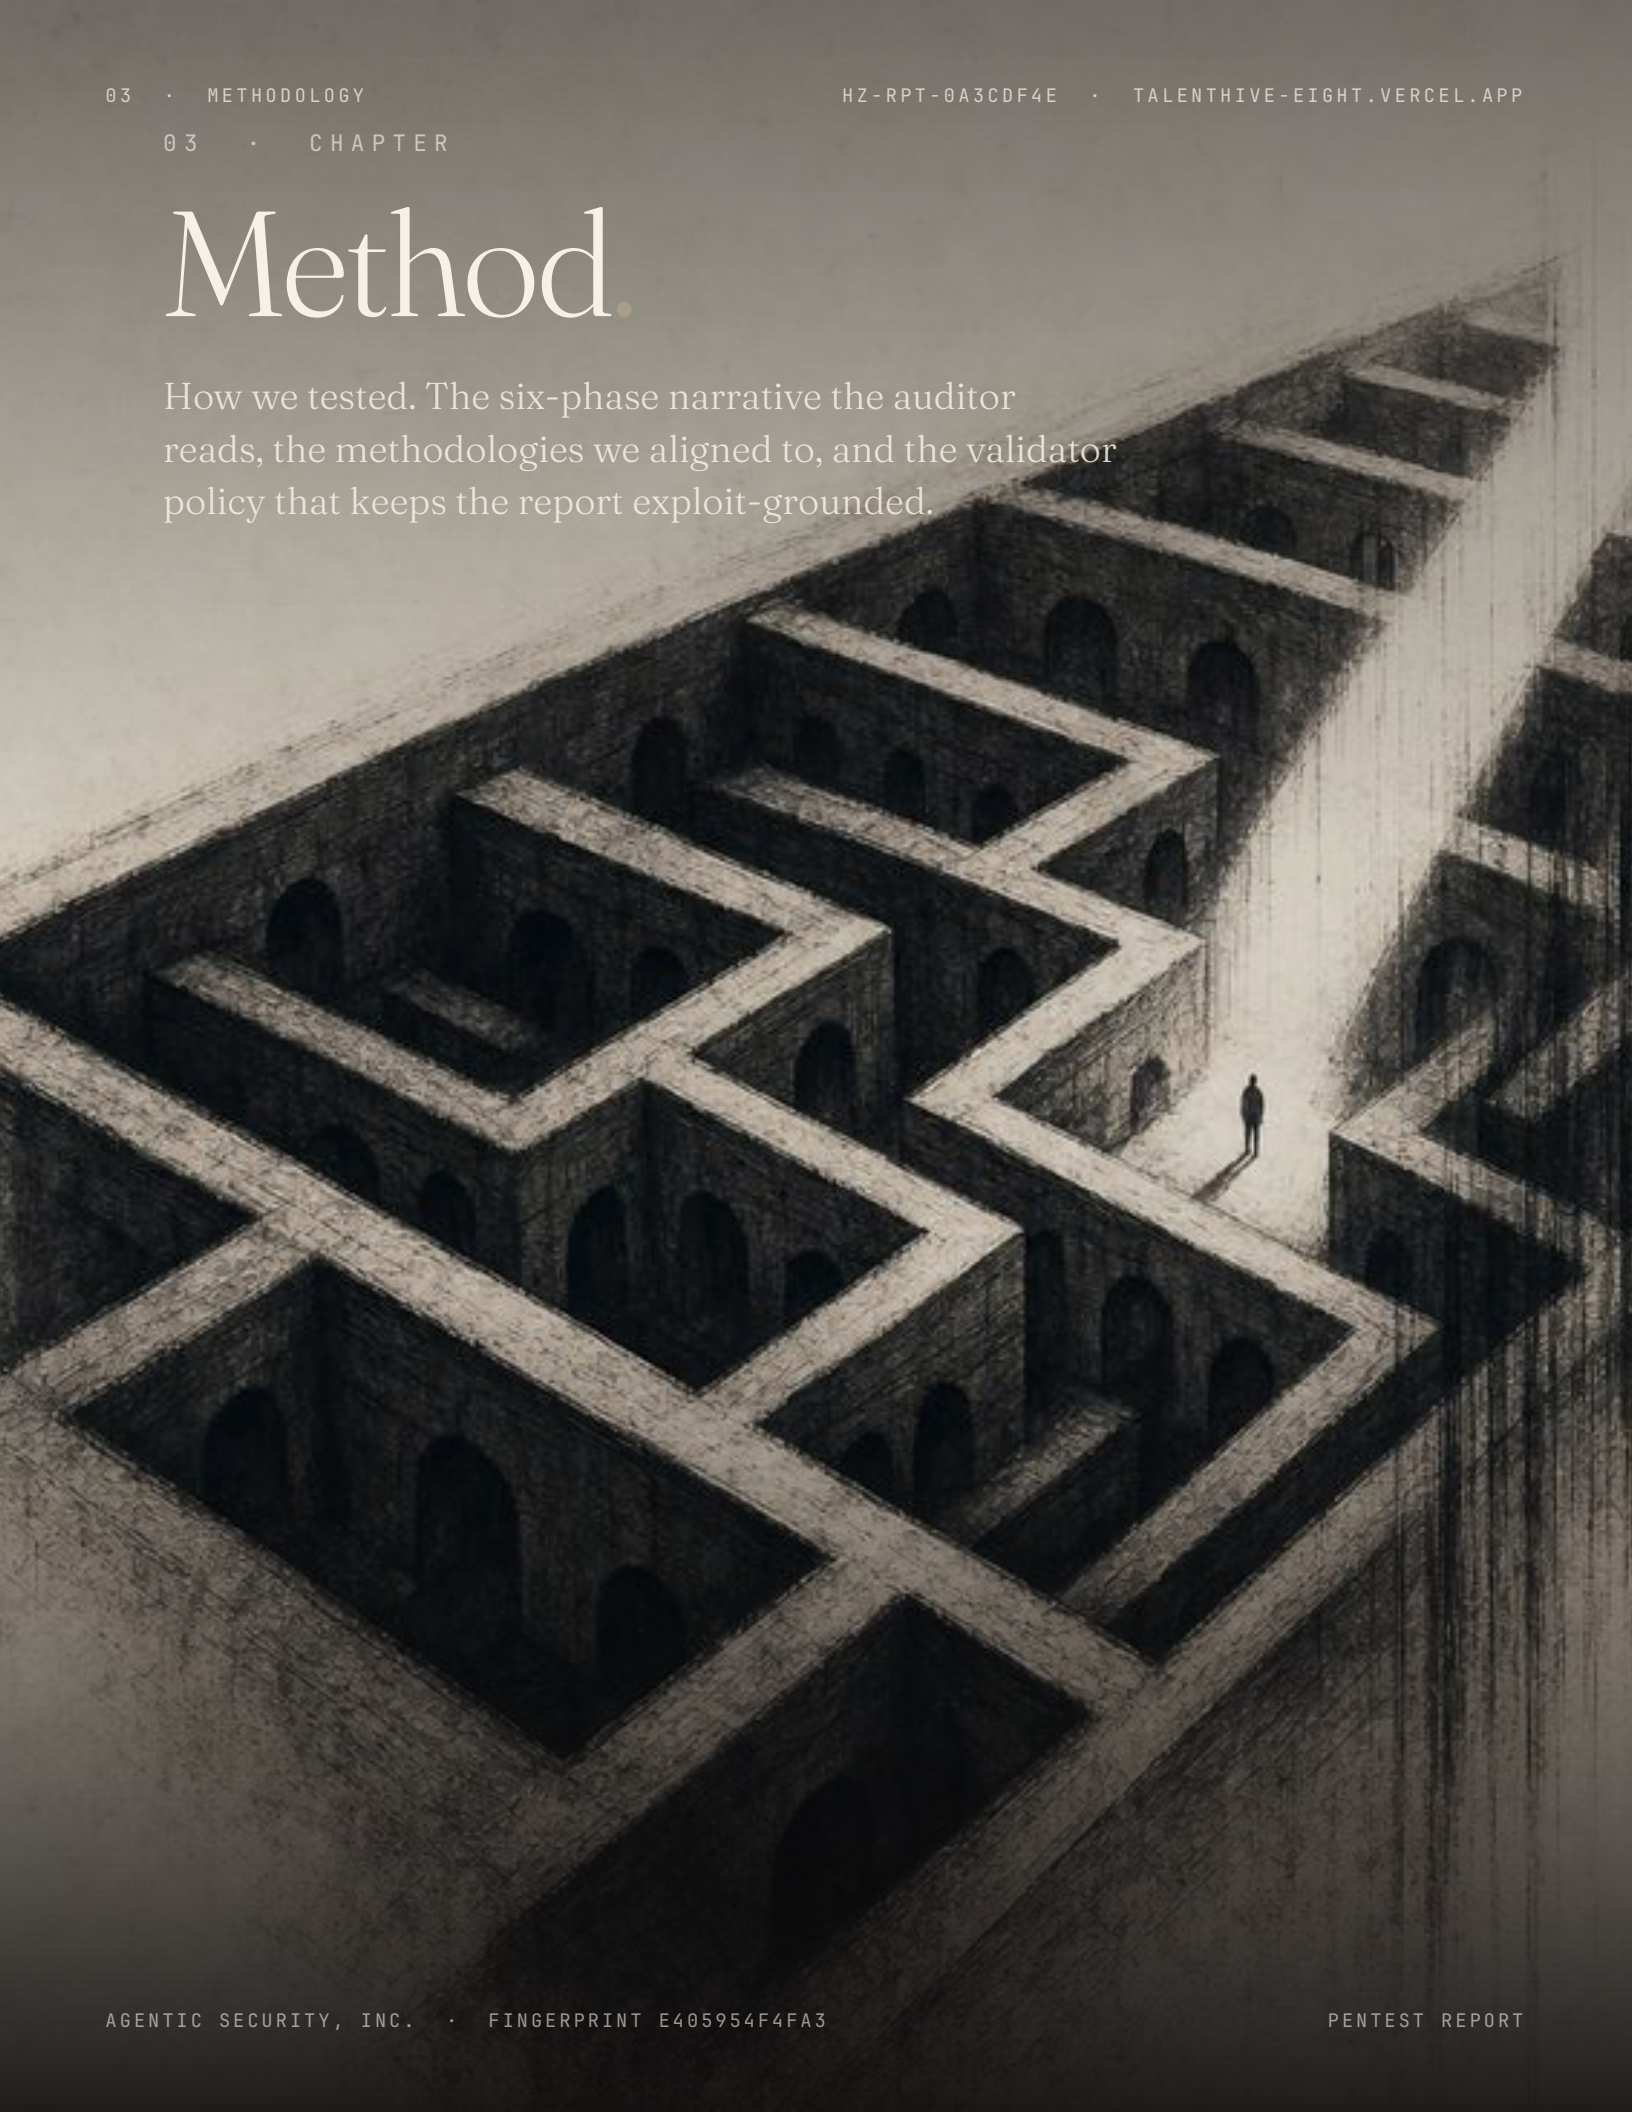

03 · METHODOLOGY

Six phases.

Every Execution moves through six phases. The phase names mirror PTES; the cadence is ours.

01 **Reconnaissance**
External enumeration of the in-scope surface. Subdomain, endpoint, parameter, and authentication discovery without active exploitation.

02 **Surface Mapping**
Classification of discovered surface into testable attack classes (auth, injection, access control, business logic, infrastructure).

03 **Exploitation**
Provider's autonomous agents attempt to exploit hypothesised vulnerabilities, bounded by the Rules of Engagement.

04 **Validation**
Every promoted finding is independently reproduced by a validator agent. Unreproducible candidates are discarded before they reach this report.

05 **Triage**
Customer's security team marks each finding Fixed, False positive, or Risk accepted (with required rationale). Decisions are timestamped and signed.

06 **Reporting**
This document. Generated server-side from the triaged findings, fingerprinted, and delivered for audit and remediation tracking.

Methodology references

The Execution aligned to the **Penetration Testing Execution Standard (PTES)**, the **OWASP Web Security Testing Guide v4.2**, and **NIST SP 800-115**. Findings are categorized against the **OWASP Top 10 (2021)** where the mapping is unambiguous. Severity is assessed per **CVSS v3.1** heuristics combined with exploit reproducibility from the validator pass count.

What is not in this report

Unreproducible candidate findings, internal cost telemetry, raw agent traces, and any artifact that would identify HackZero's internal tooling. Customers may request the raw validator transcript for any specific finding through hackzero.ai/support; it is retained under the audit-log retention specified in the Master Services Agreement.

04 · CHAPTER

What we found.

12 validated findings, ordered worst-first. Critical and High issues get a full detail spread; Medium, Low, and Informational findings are indexed at the end of the chapter for completeness.

04 · A · CRITICAL AND HIGH

Detail.

One spread per finding. Narrative, reproduction, the single best proof artifact, and the customer's triage decision with timestamp.

CRITICAL CRIT-01 OWASP A01 · VALIDATED 0× · CONFIDENCE 55%

Info disclosure

`https://talenthive-eight.vercel.app/api/users`

WHAT IT IS

The `GET /api/users` endpoint returns a full JSON array of all user records—including `id` (UUID), `name`, and `email`—to any unauthenticated caller, constituting a critical information disclosure. This was confirmed by issuing a simple `GET` request to `https://talenthive-eight.vercel.app/api/users` with no authentication headers or cookies, which returned an HTTP `200 OK` response containing 10+ user objects. The exposed UUIDs can be leveraged for Insecure Direct Object Reference (IDOR) attacks against other API endpoints, while the harvested email addresses enable targeted phishing, credential-stuffing, and account enumeration.

HOW TO REPRODUCE

The following steps prove that the `/api/users` endpoint exposes all user records without requiring any authentication.

1. Open a terminal and send an unauthenticated `GET` request to the endpoint:

```
curl -i -X GET https://talenthive-eight.vercel.app/api/users
```

1. Observe the response status is `200 OK` and the body contains a JSON array of all user objects, each with `id`, `name`, and `email` fields:

```
HTTP/2 200
content-type: application/json

[
  { "id": "<uuid>", "name": "...", "email": "..." },
  ...
]
```

1. Note that no `Authorization` header, session cookie, or API key was provided in the request, yet the full user list was returned.

Fix: Require authentication (e.g., a valid session token or API key) on the `/api/users` endpoint and enforce authorization so that only privileged roles (e.g., admins) can list all users. If bulk user listing is not a required feature, remove or disable the endpoint entirely.

STATUS: OPEN

CRITICAL CRIT-02 OWASP A01 · VALIDATED 0× · CONFIDENCE 55%

Info disclosure

`https://talenthive-eight.vercel.app/api/applications`

WHAT IT IS

The `GET /api/applications` endpoint at `https://talenthive-eight.vercel.app/api/applications` returns every application record in the system — including `user_id`, `gig_id`, `message`, `status`, profile `name` / `email`, and `gig title` — without requiring any authentication or authorization. This was confirmed by issuing a simple unauthenticated `GET` request to the endpoint, which returned an HTTP `200` response containing the full dataset of user PII and application details. Any anonymous internet user can harvest all applicant personal information and application metadata, making this a critical information-disclosure vulnerability.

HOW TO REPRODUCE

The following steps prove that an unauthenticated caller can retrieve every application record, including user PII, from the API.

1. Open a terminal and send an unauthenticated `GET` request to the applications endpoint:

```
curl -i -X GET https://talenthive-eight.vercel.app/api/applications
```

1. Observe the response returns HTTP `200` with a JSON body containing all application records. Each record exposes fields such as `user_id`, `gig_id`, `message`, `status`, profile `name`, profile `email`, and `gig title`.

1. Note that no `Authorization` header, session cookie, or API key was provided — the server returned the full dataset to a completely anonymous request.

Fix: Add server-side authentication and authorization checks to the `GET /api/applications` endpoint so that only authenticated users can access it, and scope the returned data so users can only retrieve their own applications (or enforce an admin role for bulk access).

STATUS: OPEN

HIGH HIGH-01 OWASP A01 · VALIDATED 0× · CONFIDENCE 55%

Info disclosure

`https://talenthive-eight.vercel.app/api/gigs`

WHAT IT IS

What this means

The `GET /api/gigs` endpoint at `https://talenthive-eight.vercel.app/api/gigs` requires no authentication and returns the full dataset of all gigs (21+) including sensitive fields such as `owner_id`, owner email addresses, `title`, `description`, and `budget` — constituting a bulk information disclosure of both user PII and business data. This was confirmed by issuing a simple unauthenticated `GET` request to the endpoint, which returned an HTTP `200` with the complete, unfiltered list of records. Additionally, the endpoint accepts `POST` requests for gig creation without authentication (confirmed by the server returning a `400` with `"Missing required fields"`, indicating the route is live and will process a well-formed payload), and previously stored XSS payloads are visible in `description` fields of returned records, meaning unsanitized user input is being persisted and served back to consumers of this API.

HOW TO REPRODUCE

How to reproduce

The following steps prove that the `/api/gigs` endpoint exposes all gig records — including owner PII and budget data — to any unauthenticated caller, and that the endpoint also accepts write operations without credentials.

1. Retrieve all gigs without authentication

Open a terminal and run:

```
curl -i -X GET https://talenthive-eight.vercel.app/api/gigs
```

Observe the `HTTP 200` response containing the full list of gigs (21+ records). Each object includes `owner_id`, owner email, `title`, `description`, and `budget`. No cookies, tokens, or API keys are required.

2. Confirm unauthenticated write access

```
curl -i -X POST https://talenthive-eight.vercel.app/api/gigs \
  -H "Content-Type: application/json" \
  -d '{}'
```

Observe the `HTTP 400` response with the message `"Missing required fields"`. This confirms the `POST` handler is active and will accept gig creation once the correct fields are supplied — again with no authentication.

3. Inspect stored XSS in response data

In the JSON array returned in Step 1, examine the `description` fields. Note that some entries contain unsanitized XSS payloads (e.g., `<script>` tags or event-handler attributes) that were previously submitted and stored as-is, confirming a lack of input sanitization on write and output encoding on read.

Fix: Add authentication and authorization checks to both `GET` and `POST` on `/api/gigs` so that only authenticated users can list or create gigs. Strip or restrict the fields returned to callers (remove `owner_id` and owner email from public-facing responses, or scope results to the authenticated user's own records). Sanitize and validate all user-supplied input on write, and apply context-appropriate output encoding on read, to eliminate the stored XSS vector.

STATUS: OPEN

HIGH HIGH-02 OWASP A01 · VALIDATED 0× · CONFIDENCE 55%

IDOR

`https://talenthive-eight.vercel.app/api/messages?user_id=<uuid>`

WHAT IT IS

The `/api/messages` endpoint suffers from an Insecure Direct Object Reference (IDOR) vulnerability. A `GET` request to `https://talenthive-eight.vercel.app/api/messages?user_id=<uuid>` returns that user's messages with a `200 OK` status and no authentication or authorization check whatsoever. This was confirmed by supplying the `user_id` parameter `97ab28e3-ac42-4897-9ca1-14281fadd841` without any session cookie or token — the server returned a valid messages array. Because user UUIDs are obtainable from the `/api/users` endpoint, an attacker can enumerate every user and read all of their private messages.

HOW TO REPRODUCE

The steps below prove that any unauthenticated caller can read another user's messages by supplying their UUID in the `user_id` query parameter.

1. Obtain a target user's UUID. Browse or call the users listing endpoint to collect a valid `user_id` value (e.g., `97ab28e3-ac42-4897-9ca1-14281fadd841`).

1. Request that user's messages without any credentials. Open a terminal and run:

```
curl -i -X GET 'https://talenthive-eight.vercel.app/api/messages?user_id=97ab28e3-ac42-4897-9ca1-14281fadd841'
```

1. Observe the response. The server returns `HTTP 200 OK` with a JSON body containing the user's messages array — no authentication token, cookie, or API key was required.

Fix: Enforce authentication on the `/api/messages` endpoint and verify that the authenticated user's identity matches the requested `user_id` before returning any data. Reject requests that lack a valid session or attempt to access another user's messages with a `401` or `403` response.

STATUS: OPEN

HIGH HIGH-03 OWASP A01 · VALIDATED 0× · CONFIDENCE 55%

IDOR

`https://talenthive-eight.vercel.app/api/profile?id=<uuid>`

WHAT IT IS

The `GET /api/profile` endpoint returns any user's full profile — including `email`, `name`, `bio`, `skills`, `role`, `resume_url`, and timestamps — when supplied with their UUID via the `id` query parameter, without requiring any authentication or authorization. This was confirmed by sending an unauthenticated `GET` request to `https://talenthive-eight.vercel.app/api/profile?id=97ab28e3-ac42-4897-9ca1-14281fadd841`, which returned a `200 OK` with the complete profile object. Because user UUIDs are obtainable from the `/api/users` endpoint, an attacker can enumerate and read every user's private profile data at scale.

HOW TO REPRODUCE

The following steps prove that any user's full profile can be read by an unauthenticated caller who knows (or enumerates) their UUID.

1. Open a terminal and send the following unauthenticated request to the profile endpoint, supplying a valid user UUID in the `id` parameter:

```
curl -i -X GET 'https://talenthive-eight.vercel.app/api/profile?id=97ab28e3-ac42-4897-9ca1-14281fadd841'
```

1. Observe the response returns HTTP 200 with a JSON body containing the user's `id`, `email`, `name`, `bio`, `skills`, `role`, `avatar_url`, `resume_url`, and timestamps — all without any authentication token or session cookie.

1. Repeat with any other valid UUID (obtainable from `/api/users`) to confirm the issue is not limited to a single record.

Fix: Enforce authentication on the `/api/profile` endpoint (e.g., require a valid session token or JWT) and add an authorization check so that users can only retrieve their own profile — or, where broader access is intentional, return only the minimum publicly necessary fields and strip sensitive data such as `email` and `resume_url`.

STATUS: OPEN

04 · B · MEDIUM, LOW, AND INFORMATIONAL

Index.

Lower-severity findings indexed for completeness. Full evidence for any of these is available on request through the dashboard.

SEV	ID	CLASS	ENDPOINT	TRIAGE
MED	MED-01	Cors	https://talenthive-eight.vercel.app/	Open
MED	MED-02	Clickjacking	https://talenthive-eight.vercel.app/	Open
MED	MED-03	Cors misconfiguration	https://talenthive-eight.vercel.app/	Open
LOW	LOW-01	Missing csp	https://talenthive-eight.vercel.app/	Open
LOW	LOW-02	Missing x content type options	https://talenthive-eight.vercel.app/	Open
LOW	LOW-03	Missing referrer policy	https://talenthive-eight.vercel.app/	Open
INF	INF-01	Info disclosure	https://talenthive-eight.vercel.app/	Open

05 · CHAPTER

Decisions, on the record.

Every triage action is timestamped, attributed, and signed. Risk acceptances carry their written rationale. This is the page an auditor cites.

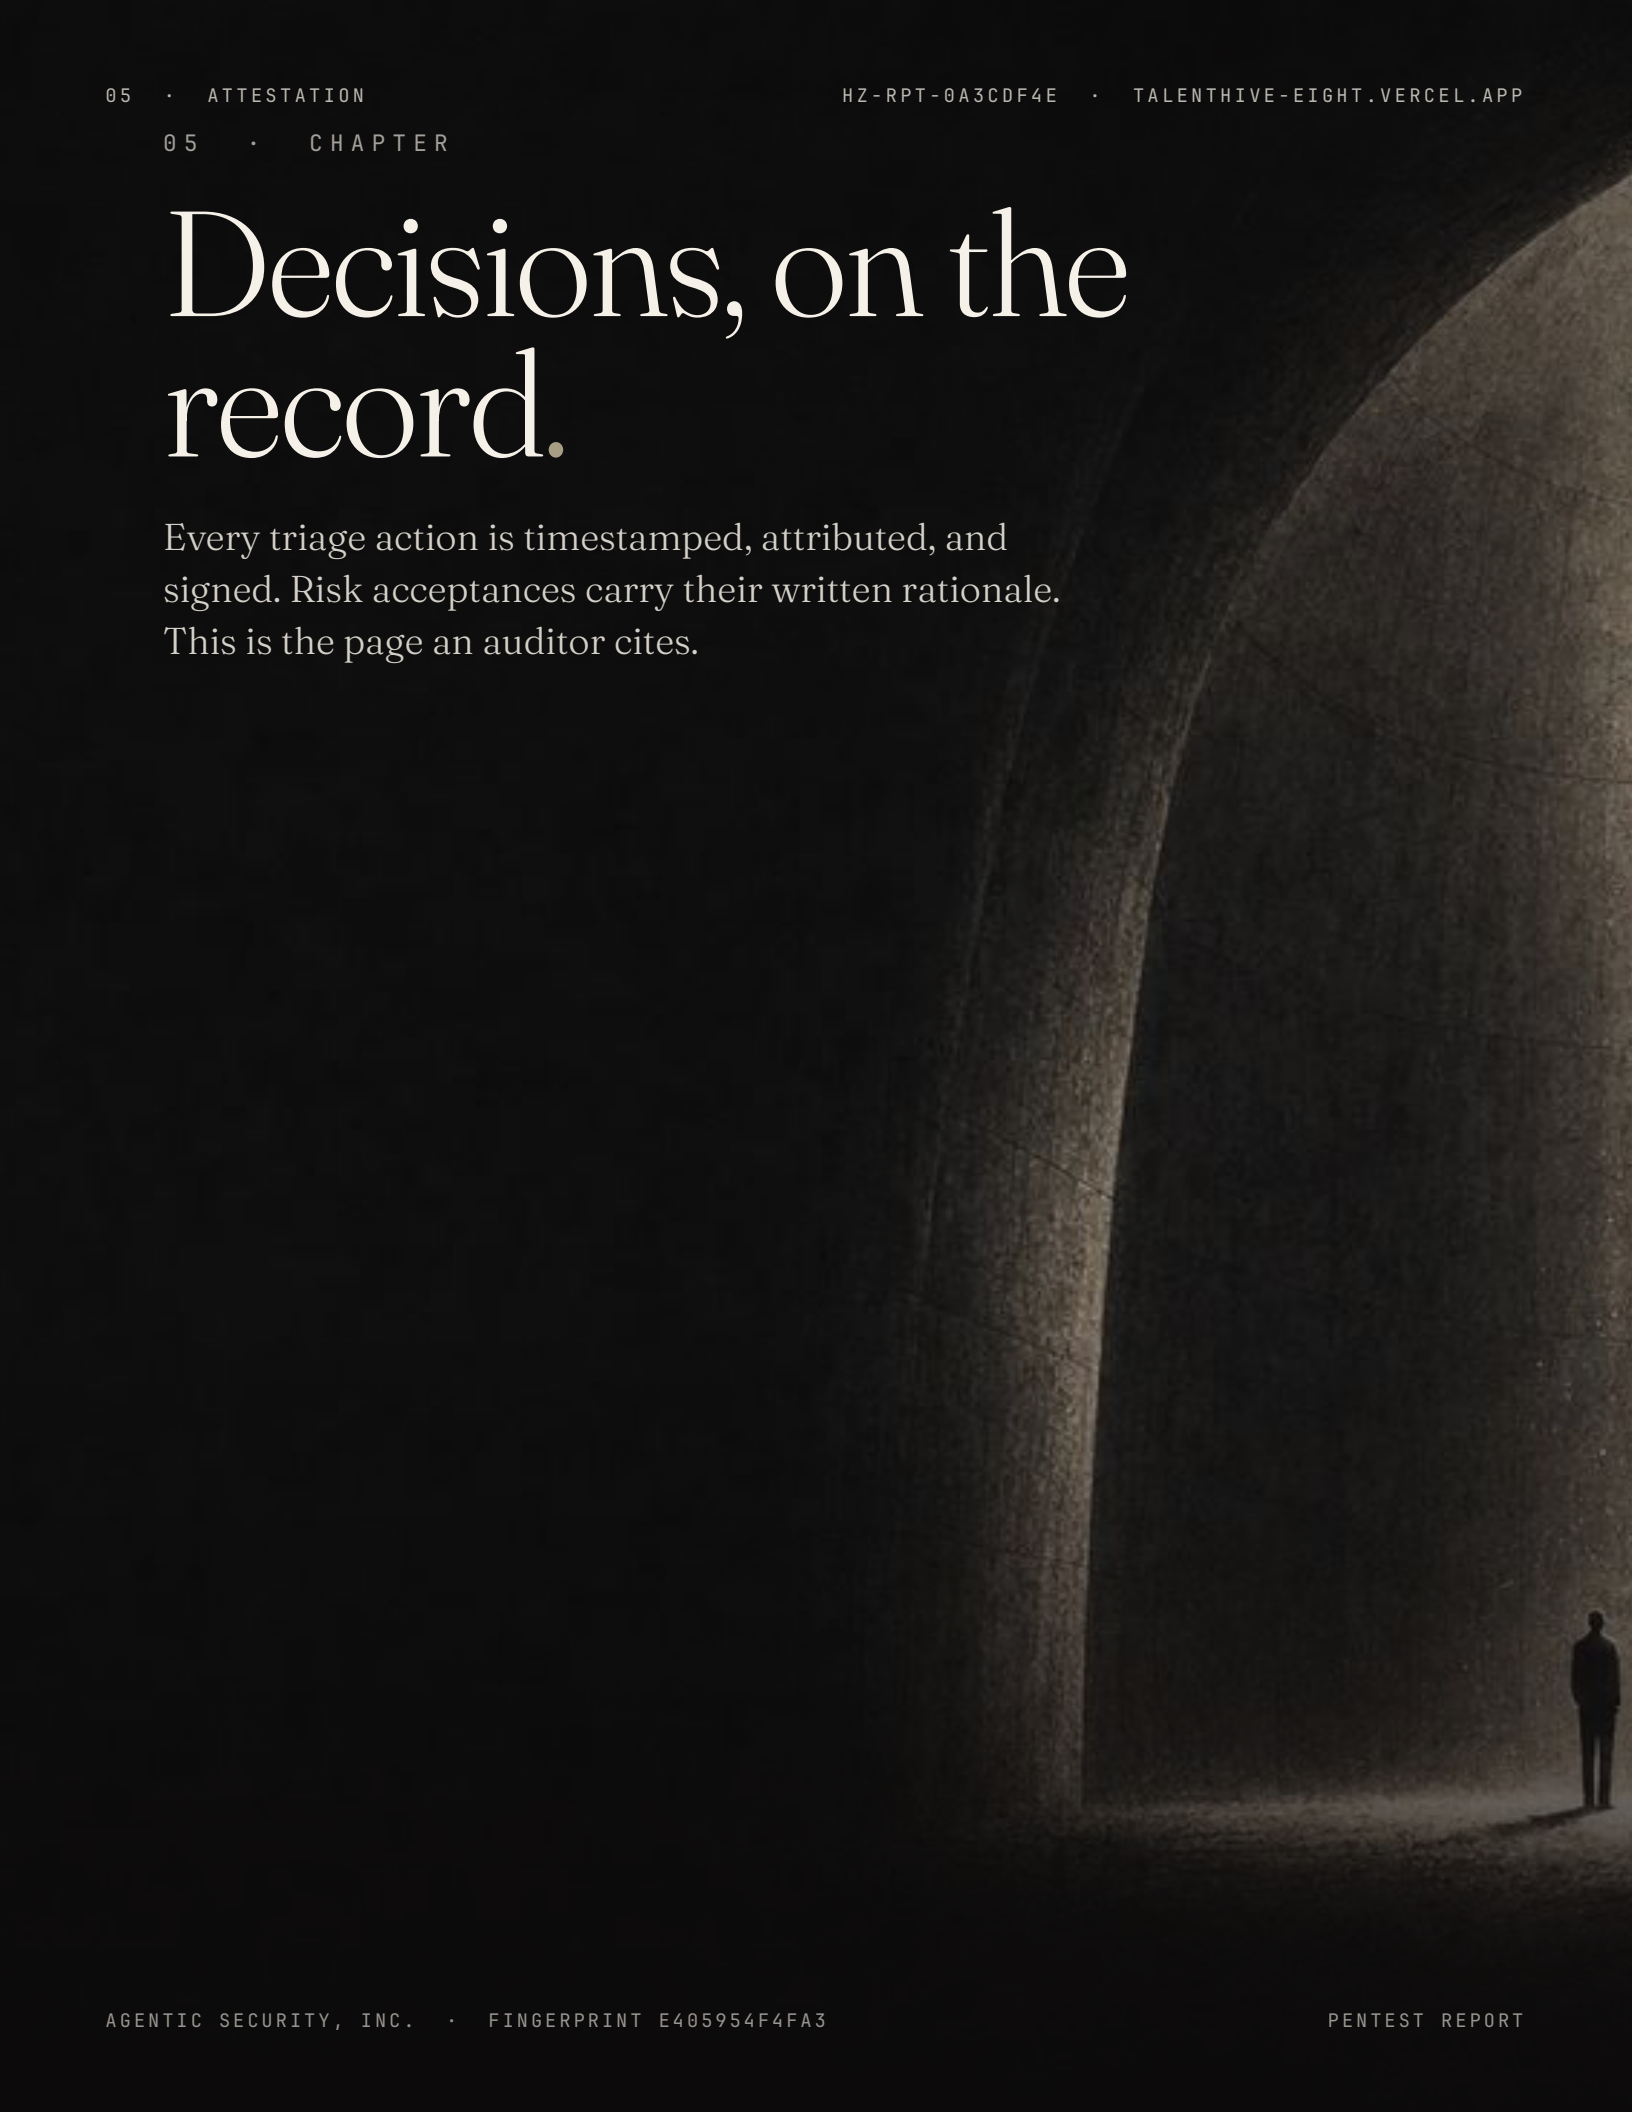

05 · A · RISK TREATMENT LEDGER

Audit trail.

Chronological record of every triage decision made on this assessment.

No triage decisions recorded.

SIGNED BY

[unsigned]

—

NO TRIAGE DECISIONS RECORDED

ELECTRONIC ATTESTATION

By triaging each finding in this assessment through the HackZero dashboard, the named signatory above confirmed: (i) the finding was reviewed by an authorized member of talenthive-eight.vercel.app's security team; (ii) the disposition recorded reflects the customer's decision; and (iii) any rationale supplied for a Risk Accepted decision represents the customer's formal acceptance of residual risk.

This document is the customer's risk-treatment record for the assessment identified above and may be relied upon by talenthive-eight.vercel.app's auditors for SOC 2 Trust Services Criteria CC4.1 and CC7.1 evidence and HIPAA Security Rule § 164.308(a)(8) Evaluation purposes.

Document fingerprint (SHA-256):
e405954f4fa39f5aa0f522ba64e57c419bc4a8794dc4abcca3dd702efd3651b8

06 · APPENDIX

For the record.

Document metadata, fingerprint, and methodology notes.

DOCUMENT	HZ-RPT-0A3CDF4E · version v1 Issued 2026-06-02 19:26 UTC
OPERATION	`0a3cdf4e-a519-46b6-9d17-db4bc2916c1f` Status completed
GENERATED BY	Agentic Security, Inc. hackzero.ai · securing what matters
FINGERPRINT (SHA-256)	`e405954f4fa39f5aa0f522ba64e57c419bc4a8794dc4abcca3dd702efd3651b8` Re-rendering this report against the same triaged findings produces an identical fingerprint. Any divergence indicates the underlying data has changed.

Severity scale

CRITICAL Direct, full-impact exposure. Authentication bypass, remote code execution, full database disclosure, mass account takeover.

HIGH Single-asset compromise or significant business-logic flaw. Authenticated privilege escalation, sensitive data disclosure, scoped account takeover.

MEDIUM Realistic exploitation requires chaining or unusual conditions. Information disclosure of moderate value, weak crypto, rate-limit gaps.

LOW Limited impact on its own. Headers, configuration hygiene, low-value information leaks.

INFORMATIONAL No direct security impact. Observations the customer may still want to address for posture or compliance.