SAMPLE REPORT: FICTIONAL CLIENT AND SYNTHETIC DATA. This document is not a record of a real engagement. No real organization, system or personal data was tested or exposed. Findings, evidence, measurements and named roles are illustrative and are modeled on the structure and rigor of our real reports. Do not cite it as evidence of any actual assessment result.

KEYPREFER TECHNOLOGIES | Offensive Security | Sample deliverable

Penetration Test Report

Web Application and API Security Assessment

Client	Lumenwave HR Cloud (fictional) | Report version 1.6
Target	Lumenwave HR Cloud: web application and public REST API (staging)
Assessment type	Gray-box web application and API penetration test (authenticated; two test tenants; three test accounts in two role types: standard user and HR-admin)
Testing window	September 14–18, 2026 (5 person-days)
Retest window	October 5–6, 2026
Report issued	October 28, 2026, version 1.6 (supersedes 1.5, 1.4, 1.3 and 1.2, all dated October 28, 2026; 1.1 dated October 9, 2026; 1.0 dated September 21, 2026)
Classification	Sample document: fictional client and synthetic data; may be shared with prospective clients. Real KeyPrefer Technologies client reports are classified Confidential (client and named recipients only).
Standards	OWASP WSTG v4.2, OWASP API Security Top 10 (2023), CVSS v3.1

Document control

Role	Holder	Responsibility
Lead assessor	Illustrative role: Senior Security Consultant	Test execution, findings, sign-off
Peer reviewer	Illustrative role: Principal Security Consultant	Independent QA of findings and CVSS metrics (Section 2.2); illustrative role, no review executed on this sample
Client contact	CTO, Lumenwave HR Cloud	Scope authorization, remediation owner

Team and review model. Every engagement has a lead assessor, who executes testing and signs off the report, and a separate peer reviewer, who checks findings and CVSS metrics before issue. In client deliverables this table names each individual and lists their certifications. The roles above are illustrative only, so this sample cannot demonstrate the experience of any individual reviewer, and it is not evidence that the review in Section 2.2 took place on this document.

Revision history

Ver.	Date	Change
1.0	Sep 21, 2026	Initial report.
1.1	Oct 9, 2026	Retest results (Oct 5–6, 2026) added; INF-01 resolution recorded.
1.2	Oct 28, 2026	SEC-04 rescored: Attack Complexity High to Low, CVSS 4.8 to 6.5. Document control roles and responsibilities updated.
1.3	Oct 28, 2026	Review comments addressed: tenant-ownership wording clarified; SEC-04 impact reassessed (C:H/I:H, 9.1); SEC-03 test conditions and score sensitivity added; attack chain labeled hypothetical; staging/production limitation expanded; SEC-01 status updated to core flaw fixed, lockout validation incomplete; Appendix B clarified; document classification updated.
1.4	Oct 28, 2026	SEC-04 demonstrated and derived impact separated; SEC-03 timing summary added; review procedure and team model described; open validation items consolidated; production verification plan and Appendix B reproducibility table added.
1.5	Oct 28, 2026	Roles clarified (three accounts, two role types); SEC-04 headline labeled derived impact; SEC-03 timing table extended with min/p5/p95/max columns; coverage table split (INF-02, INF-03).
1.6	Oct 28, 2026	(1) SEC-04: 6.5 Medium (demonstrated) is now the headline severity throughout; 9.1 Critical is presented as a conditional scenario (HR-admin compromise, not tested). (2) SEC-04 per-IP limiter behavior stated (failures 1 to 10 per minute allowed; 429 from the 11th), giving a theoretical bound of 25 x 10 = 250 per minute versus 200 per minute measured; 225 would apply only if the 10th failure were refused; window type unverified. (3) SEC-03 timing split into an as-reported table and a separate illustrative table. (4) Review status clarified (Section 2.2). (5) 25-IP test parameters documented; MFA wording and status vocabulary standardized; evidence-led summary, coverage table with basis flags and consistency checks (Appendix C) added. No retest observation or result from earlier versions was changed.

1. Executive summary

Sample notice. Lumenwave HR Cloud is fictional. All hosts use the reserved .example domain, and all requests, responses and measurements are synthetic. This summary describes a hypothetical engagement.

KeyPrefer Technologies assessed the Lumenwave HR Cloud staging environment from a gray-box position, using three test accounts (standard user and HR-admin in Tenant A; standard user in Tenant B) in two test tenants provisioned on staging for this engagement and administered by KeyPrefer Technologies testers, so no real customer data was accessed. The client states that staging is configured equivalently to production; KeyPrefer Technologies has not verified this. **All findings and retest results apply to staging builds 2026.09.3 and 2026.10.2 only** and must be confirmed against production (Section 4).

We found four vulnerabilities: one Critical, one High and two Medium, plus three hardening observations. SEC-04 is Medium (6.5) on demonstrated impact, with a conditional Critical (9.1) scenario that was not demonstrated. On October 5 and 6 we retested all four on staging. SEC-02 and SEC-03 are **fixed on staging**. SEC-01 has its **core flaw fixed**, but two lockout checks were not tested. SEC-04 is **partially fixed and open**. No finding is yet Closed (definition below). Hardening observation INF-01 was resolved on staging.

Key evidence at a glance

ID | Severity (CVSS 3.1)	Strongest evidence (staging, synthetic)	Status (staging)
SEC-01 Password reset accepts any OTP **Critical 9.1**	Reset with OTP 000000 returned 200; the new password then gave a full session as the victim. Expired, reused and malformed OTPs were also accepted (E-01 to E-03).	Core flaw fixed; lockout validation incomplete (SEC-01-a, -b not tested)
SEC-02 Cross-tenant access to employee records (BOLA) **High 8.1**	A Tenant A standard user read 12 pay stubs and changed bank details of a Tenant B employee. 2 of 64 endpoint/method combinations affected (E-04, E-05).	Fixed on staging; 64 of 64 returned 404
SEC-03 User enumeration via response timing **Medium 5.3**	Medians of 412 ms (registered) vs 96 ms (unregistered), 200 samples each, no overlap in the samples collected: one request separated the classes (E-07).	Fixed on staging; retest medians 102 ms vs 99 ms, overlapping
SEC-04 No effective login rate limiting **Medium 6.5 (demonstrated)** **Conditional: Critical 9.1** if an HR-admin account were compromised (not tested)	1,000 attempts from one IP drew no throttling; a seeded standard-user test account fell at attempt 212. In retest the per-IP limit works, but a 25-IP run against one account, paced at 8 attempts per IP per minute, drew no delay or challenge. **HR-admin takeover was not tested or demonstrated.**	Partially fixed; open

Severity follows CVSS v3.1 base scores (Critical 9.0 to 10.0, High 7.0 to 8.9, Medium 4.0 to 6.9, Low 0.1 to 3.9); every metric choice is justified in Section 3. SEC-04's severity is its demonstrated 6.5 Medium; its conditional 9.1 Critical scenario assumes HR-admin compromise, which was not tested (Section 3). Hardening observations (Section 5) are not scored.

Status terms used throughout

Term	Meaning
Fixed on staging	The original exploit failed on build 2026.10.2 and valid-flow controls passed. Production unverified. Defense-in-depth recommendations are follow-ups, not conditions.
Core flaw fixed	The main exploit is fixed, but one or more recommended controls were not tested (SEC-01).
Partially fixed, open	Some exploit paths are closed; at least one observed path remains (SEC-04).
Closed	Reserved: fixed on staging, all open validation items verified, and the fix confirmed deployed in production. **No finding is currently Closed.**

Overall posture

Within the tested scope and time, the controls we examined held up: role-based function access, session handling, input validation and file upload (Section 2.3). This is a sampling result from a five-person-day test of staging, not a statement about areas we did not test. The serious weaknesses sit in three places: the password-reset flow, tenant isolation on a subset of employee endpoints, and login protection.

SEC-04 is the only finding with an observed exploit path still open. SEC-01 has no observed open exploit path, but its lockout validation is incomplete (SEC-01-a, SEC-01-b), so it is not closed. SEC-04 is rated **Medium (6.5) on demonstrated impact**: the login endpoint has no per-account control against distributed password guessing, and a seeded standard-user test account was compromised. **Conditional scenario, not demonstrated:** if an HR-admin account were compromised, the impact would score 9.1 (Critical). HR-admin compromise was not tested, and MFA for HR-admin roles is reported as not enforced (not independently verified by KeyPrefer Technologies; SEC-04-c). We keep the 7-day remediation target of the conditional scenario because the missing per-account control is observed. Per-IP throttling now works and raises the cost of single-source attacks. The distributed case drew no control in retest even though every source IP stayed below the per-IP limit (Section 3), and is the main residual risk.

Illustrative attack chain (hypothetical, pre-remediation)

Each step was demonstrated individually against test accounts on staging. KeyPrefer Technologies did not execute the whole chain end to end, and no real customer data was involved. It shows how the findings could have combined on the pre-fix build; it is not a record of an actual compromise.

1. Discover (optional)	2. Take over	3. Pivot
Confirm which emails have accounts (SEC-03, timing).	With a valid user ID or email, reset the victim's password using any OTP (SEC-01).	Use any authenticated account to read and edit other tenants' employee and bank records (SEC-02).

SEC-04 is a separate route to initial access through password guessing. Business impact if these weaknesses were present in production: account takeover of users, potentially including HR administrators, and cross-tenant exposure of salary, bank and identity data (SSNs). Such exposure could amount to a personal-data breach with notification duties under US state breach laws and CCPA/CPRA; legal counsel should assess that.

Recommended actions

Priority	Action	Owner	Target
Now	Close SEC-04: add a per-account control that works against distributed guessing, and enforce MFA for HR-admin roles (currently reported as not enforced; unverified) as an interim control.	Platform / Identity	7 days
Now	Confirm staging/production parity (authentication flows, tenant scoping, gateway rate limits), confirm build 2026.10.2 fixes are deployed in production, and request a targeted production verification retest (Section 4).	CTO office / Platform	14 days

Priority	Action	Owner	Target
Now	Complete open validation items: SEC-01-a and -b; SEC-04-a to -d, including a test against an authorized HR-admin test account (Section 4).	Platform / KeyPrefer Technologies retest	14 days
Next	Add breached-password checks; alert on bursts of failed logins and resets.	Security / Ops	60 days
Next	Add automated cross-tenant authorization tests to CI so SEC-02 cannot regress.	Backend	60 days
Later	Schedule an annual retest and a retest after any rewrite of authentication flows.	CTO office	12 months

Positive observations

- Session tokens are short-lived, rotated on privilege change and invalidated on logout.
- Function-level authorization (HR-admin vs. standard user) was enforced server-side on every tested endpoint (112 endpoint/role combinations).
- No injection, SSRF or insecure file-handling issues were found across the surface and time we tested.
- The core flaws in SEC-01 (Critical) and SEC-02 (High) were fixed on staging before the October 5–6 retest.

2. Scope, methodology and limitations

2.1 Scope and rules of engagement

Asset	Location	Access given
Web app (SPA)	https://staging-app.lumenwave.example	Standard user and HR-admin in Tenant A; standard user in Tenant B
REST API	https://api-staging.lumenwave.example/v2/*	Same credentials plus API documentation (OpenAPI)
Out of scope	Production, cloud management plane, DoS/DDoS, social engineering, third-party payroll provider	n/a

Accounts used: Tenant A standard user, Tenant A HR-admin and Tenant B standard user (three accounts, two role types). Tenants A and B are test tenants provisioned on the client's staging environment, administered by KeyPrefer Technologies testers, and hold synthetic data only. Testing was non-destructive. Credential-guessing was limited to the Tenant A standard-user test account (seeded with a weak password); no HR-admin account was targeted. Request volume was capped at 15 requests per second. Critical issues were reported to the client within 24 hours of confirmation.

2.2 Approach and review procedure

Testing followed OWASP WSTG v4.2 and the OWASP API Security Top 10 (2023), combining manual testing, intercepting-proxy tooling (Burp Suite Professional) and custom scripts for authorization matrix and timing analysis. Each finding was reproduced at least three times by the lead assessor and scored with a documented CVSS rationale.

Peer review (procedure in client engagements). Before issue, a second assessor who did not own the finding reviews it for (a) whether the evidence supports the stated behavior and impact; (b) an independent re-derivation of each CVSS metric, with disagreements resolved in writing; and (c) whether limitations and untested aspects are stated. **Status for this sample:** revisions 1.2 to 1.6 of this document were editorial iterations made to improve the sample. They are not evidence that an independent security review was executed, and no reviewer sign-off exists for this sample. Because the roles are illustrative, this sample documents the procedure only; it cannot show who reviewed it or their qualifications. In client deliverables the reviewer is named and their certifications listed.

Illustrative CVSS peer-review checklist (template only). The layout below shows what the reviewer records per finding. It is an example of format, deliberately left uncompleted: it records no outcome for any finding in this report and is not evidence that a review was performed.

Metric / check	Reviewer question	Illustrative outcome field
AV, AC, PR, UI	Does the evidence show the position, preparation and privileges the attacker needs? Is effort being mistaken for complexity?	Agree / disagree + written note
S (Scope)	Does the impact cross an authority boundary, or stay within the same application authority?	Agree / disagree + written note
C, I, A	Is the impact demonstrated or derived? If derived, is the basis stated and is the demonstrated-only score also shown?	Agree / disagree + written note
Vector and score	Does the recomputed base score match the displayed vector and severity band?	Match / mismatch
Limitations	Are untested aspects and environment differences stated next to the finding?	Stated / missing

2.3 Coverage and results

The table lists the primary test-guide references exercised, representative test cases and feature groups, and outcomes. It is **not** a complete mapping to WSTG v4.2 or the OWASP API Top 10: items not listed were not claimed. Basis: **R** = figures stated elsewhere in this report; **I** = illustrative synthetic cases for this sample, with no counts claimed.

Area	Reference	Representative test cases / feature groups	Outcome	Basis
Account enumeration	WSTG-IDNT-04	Forgot-password response body, status and timing for registered vs unregistered emails	Finding SEC-03	R
Weak lockout / rate limiting	WSTG-ATHN-03	Login: repeated failures from 1 IP, many IPs and many identifiers (/v2/auth/login)	Finding SEC-04	R

Area	Reference	Representative test cases / feature groups	Outcome	Basis
Password reset logic	WSTG-ATHN-09	Reset with wrong, expired, reused, malformed and omitted OTP (/v2/auth/reset-password)	Finding SEC-01	R
Object-level authorization (BOLA)	API1:2023	Tenant A tokens against Tenant B identifiers on employee-scoped routes (64 endpoint/method combinations)	Finding SEC-02 (2 of 64)	R
Excessive data in responses	API3:2023	Login response fields versus what the UI needs	Observation INF-01	R
Function-level authorization (BFLA)	API5:2023	Standard-user tokens on HR-admin functions (112 endpoint/role combinations)	No finding	R
Mass assignment / property tampering	API3:2023	Extra or privileged properties in write requests (31 write endpoints)	No finding	R
Session management and token handling	WSTG-SESS-01 to 07 (subset)	Token lifetime, rotation on privilege change, invalidation on logout, replay of old tokens	No finding	I
Injection	WSTG-INPV (subset)	SQL/NoSQL metacharacters in search and filter parameters; template syntax in document fields; shell characters in import filenames	No finding	I
Server-side request forgery	WSTG-INPV; API7:2023	URL-based fields in document import; internal-address targets	No finding	I
File upload (pay stub and document import)	WSTG-BUSL-08/ 09	Unexpected extensions and content types, oversized files, mismatched file headers	No finding	I
CORS, security headers, TLS configuration	WSTG-CONF (subset)	Origin handling and preflight, response header set, TLS protocol and cipher review	INF-02 (HSTS)	I
Error handling and information disclosure	WSTG-ERRH-01	Malformed JSON, invalid identifiers and oversized bodies on API routes	INF-03 (framework version in a 500)	I
Business logic (leave, expense approvals)	WSTG-BUSL (subset)	Approval by a non-approver, self-approval, resubmission after approval	No finding	I

2.4 Limitations

- Point-in-time assessment of staging builds 2026.09.3 and 2026.10.2; later releases are not covered.
- Production was not tested. Staging-to-production equivalence is a client statement KeyPrefer Technologies has not verified. Differences in configuration, in particular gateway or WAF rate limits, tenant configuration and email infrastructure, could change exposure and could affect SEC-03 and SEC-04 most directly.
- Fix verification was performed on staging build 2026.10.2 only. A finding marked Fixed on staging is unverified in production until the fix is confirmed deployed there.
- Gray-box testing with five person-days is sampling-based and cannot prove the absence of vulnerabilities. "No finding" means none was found in the tested scope.
- Timing measurements (SEC-03) came from a single vantage point at a capped request rate; server load was neither controlled nor measured.
- Password guessing was demonstrated against a standard-user test account only. HR-admin impact is derived from role privileges, not demonstrated (SEC-04).
- The 25-IP distributed test (SEC-04) used 25 source addresses from one test harness at a deliberately low pace; a real distributed attacker may use more addresses and different pacing.
- Mobile apps, SSO/SAML integrations, internal admin tooling and the payroll-provider integration were out of scope.

3. Detailed findings

SEC-01 | Password reset accepts any OTP	CRITICAL 9.1
Vector	CVSS:3.1/AV:N/AC:L/PR:N/UI:N/S:U/C:H/I:H/A:N
Affected	POST /v2/auth/reset-password
Mapping	CWE-640, CWE-287 | API2:2023 Broken Authentication
Status (staging)	**Core flaw fixed** (verified Oct 5, 2026); lockout validation incomplete, so not fully closed; production unverified

Description

The reset endpoint checks that an otp field is present but never compares it to the code issued to the user. The OTP is verified by a separate step whose result is not bound to the reset request on the server, so a caller can skip it.

Impact

Any unauthenticated attacker who knows a valid user ID or email can set a new password and sign in as that user, including HR administrators. C:H and I:H are justified because a full session with the victim's privileges was obtained in testing.

Evidence (synthetic, sanitized)

```
E-01 POST /v2/auth/reset-password HTTP/2
Host: api-staging.lumenwave.example
```

SAMPLE

```
{"userId":"usr_t1_0042","resetMode":"otp","otp":"000000","newPassword":"[REDACTED]"}
-> HTTP/2 200 {"status":"password_updated"}
E-02 POST /v2/auth/login with the new password -> HTTP/2 200, session issued.
Audit log shows no OTP verification event for this reset.
```

Negative control (E-03)	Result on vulnerable build
Expired OTP (issued 30 min earlier)	200 accepted
Previously used OTP	200 accepted
Malformed OTP ("abc")	200 accepted
OTP field omitted	400 field required (presence check only)

Remediation

- Issue an opaque, random, single-use reset token only after the OTP is verified server-side; bind it to the account and a 10-minute expiry.
- Make the reset endpoint accept only that token, never client-supplied identity or verification flags.
- Invalidate the token atomically on use; email the user on every reset and revoke existing sessions.
- Attempt limit: a challenge is the single OTP issued for one account by one reset request. Count invalid attempts against the challenge, regardless of source IP or session, and invalidate it after the fifth invalid attempt, so further attempts are refused even with the correct code. Do not lock the account itself (lockout-based denial of service).
- Rate-limit issuance of new challenges per account (for example, three per hour) so the challenge limit cannot be bypassed by requesting fresh codes.

Retest (Oct 5, 2026, build 2026.10.2, staging)

Invalid OTP returned 403 invalid_reset_token; expired token 410; reused token 409; the legitimate flow returned 200 and the old password stopped working. After five invalid attempts against one challenge, a sixth attempt was refused, confirming that failures accumulate against the challenge across successive requests. Reset notification emails and session revocation were observed.

Not covered by the retest (open items): (a) whether the failure counter is shared across source IPs and sessions (SEC-01-a); (b) whether issuance of new challenges is rate-limited (SEC-01-b). Until verified, SEC-01 is recorded as **core flaw fixed, lockout validation incomplete**, not closed.

SEC-02 | Cross-tenant access to employee records (BOLA)	HIGH 8.1
Vector	CVSS:3.1/AV:N/AC:L/PR:L/UI:N/S:U/C:H/I:H/A:N
Affected	GET /v2/employees/{id}/paystubs, PUT /v2/employees/{id}/bank-details
Mapping	CWE-639 | API1:2023 Broken Object Level Authorization
Status (staging)	**Fixed on staging** (verified Oct 6, 2026); production unverified

Description

Two employee sub-resources authorized the caller's role but not the tenant that owns the requested object. Employee identifiers are sequential integers, so identifiers in other tenants can be enumerated. Of 64 employee-scoped endpoint and method combinations tested, 2 were affected.

Impact

Any authenticated user could read pay stubs and overwrite bank details of employees in other tenants. Tenant B was a test tenant administered by KeyPrefer Technologies, so no real customer data was touched. Both read and write were demonstrated, hence C:H and I:H. Scope is Unchanged because the tenant boundary is enforced by the same application authority.

Evidence (synthetic, sanitized)

```
E-04 GET /v2/employees/20417/paystubs HTTP/2 (token: Tenant A standard user)
-> HTTP/2 200 [Tenant B employee, 12 pay stub records]
E-05 PUT /v2/employees/20417/bank-details HTTP/2
{"routing":"[TEST-ABA]","account":"[TEST-ACCT]"} -> HTTP/2 200; change confirmed from Tenant B admin view
```

Remediation

- Scope every data-layer query by tenant (tenant ID predicate or row-level security), not only by role.
- Centralize object-level authorization in middleware so new endpoints inherit it by default.
- Return 404 for objects outside the caller's tenant; use non-sequential identifiers as defense in depth.
- Add automated two-tenant authorization tests for every employee-scoped route to CI (follow-up, Section 4).

Retest (Oct 6, 2026, build 2026.10.2, staging)

All 64 combinations were repeated with Tenant A tokens against Tenant B identifiers: every request returned 404 and no state changed. Same-tenant requests still succeeded. **Fixed on staging.** The CI tests and non-sequential identifiers were not observed in retest and remain follow-ups.

SEC-03 | User enumeration via response timing	MEDIUM 5.3
Vector	CVSS:3.1/AV:N/AC:L/PR:N/UI:N/S:U/C:L/I:N/A:N

SAMPLE

Affected	POST /v2/auth/forgot-password
Mapping	CWE-208, CWE-204 | API2:2023 Broken Authentication
Status (staging)	**Fixed on staging** (verified Oct 5, 2026); production unverified

Description

Response bodies and status codes were identical for registered and unregistered emails, which is correct. However, the server sent the email synchronously for registered accounts. Over 200 samples each, median response time was 412 ms for registered and 96 ms for unregistered addresses, and in the samples collected the two sets did not overlap (illustrative slowest unregistered sample 139 ms; illustrative fastest registered sample 371 ms). On staging, under our test conditions, a single request was therefore sufficient to separate the two classes.

Timing summary (staging, single vantage point, capped request rate)

Two tables, two data statuses. Table A holds the values as originally reported (sample size and medians). **Table B** (amber) holds **illustrative synthetic Min, p5, p95 and Max values, not measured engagement data**, added in version 1.5 to show how the separation would look. Do not read Table B as empirical evidence.

Table A. As originally reported

Build	Class	n	Median
2026.09.3 (initial)	Registered email	200	412 ms
2026.09.3 (initial)	Unregistered email	200	96 ms
2026.10.2 (retest)	Registered email	200	102 ms
2026.10.2 (retest)	Unregistered email	200	99 ms

Table B. ILLUSTRATIVE ONLY: synthetic values, not measured

Build	Class	Min (synthetic)	p5 (synthetic)	p95 (synthetic)	Max (synthetic)
2026.09.3	Registered	371 ms	380 ms	451 ms	478 ms
2026.09.3	Unregistered	84 ms	87 ms	118 ms	139 ms
2026.10.2	Registered	81 ms	86 ms	131 ms	152 ms
2026.10.2	Unregistered	80 ms	85 ms	127 ms	149 ms

On the illustrative figures, the registered minimum (371 ms) sits 232 ms above the unregistered maximum (139 ms) before the fix, and the ranges overlap almost entirely after it. In a real engagement the raw samples, a histogram and the collection script are delivered in the evidence pack (E-07; Appendix B).

Test conditions and reliability

Measurements came from a single vantage point at a capped request rate against staging; server load was neither controlled nor measured, and we did not test from multiple networks or under load. "No overlap" describes the 200 samples collected per class, not every future request. Network jitter, queueing or faster email dispatch in production could narrow the gap; an attacker would then need repeated samples per address, which costs time but does not remove the weakness.

Scoring rationale

Attack Complexity is Low: on staging the difference was consistent and reproducible and needed no special access or preparation. Collecting measurements is effort, not complexity under the CVSS definition. Confidentiality is Low because the response discloses account existence, a limited piece of information. Sensitivity: if production conditions made the signal unreliable enough to justify AC:H, the base score would be 3.7 (Low). We score the observed behavior and rely on the 30-day remediation window rather than an adjusted score.

Remediation and retest

Email dispatch was moved to an asynchronous queue. Retest over 200 samples each (same single vantage point) gave medians of 102 ms and 99 ms with overlapping distributions; bodies, status codes and headers remained identical. **Fixed on staging.**

SEC-04 | No effective login rate limiting	DEMONSTRATED 6.5 MEDIUM | CONDITIONAL 9.1 CRITICAL
Severity (demonstrated)	CVSS:3.1/AV:N/AC:L/PR:N/UI:N/S:U/C:L/I:L/A:N = **6.5 Medium** (standard-user test account only)
Conditional scenario	CVSS:3.1/AV:N/AC:L/PR:N/UI:N/S:U/C:H/I:H/A:N = **9.1 Critical** if an HR-admin account were compromised (not tested)
Affected	POST /v2/auth/login
Mapping	CWE-307 | API2:2023 Broken Authentication
Score history	v1.0 and 1.1: 4.8 Medium (AC:H, C:L/I:L). v1.2: 6.5 Medium (AC:L, C:L/I:L). v1.3 to 1.5: headline 9.1 Critical on derived impact (AC:L, C:H/I:H). v1.6: 6.5 Medium (demonstrated) is the headline and 9.1 is the conditional scenario; no vector changed.
Status (staging)	**Partially fixed; open** (retest Oct 5, 2026); production unverified

Description

The login endpoint applied no throttling, delay or challenge to repeated failures. From one IP, 1,000 attempts against a single test account completed in 4 min 10 s (250 s, about 4 attempts per second, inside the 15 requests-per-second cap), every failure

returning an identical 401. The Tenant A standard-user test account, with a deliberately seeded weak password, was compromised on attempt 212, roughly 53 seconds in at that rate, giving access to that user's own records only. The seeded password is a test condition that demonstrates the control gap; it does not indicate how many real accounts are exposed. No HR-admin account was targeted.

What was demonstrated and what is conditional

Basis	Statement	Status
Demonstrated (original test)	No throttling, delay or challenge from a single IP; seeded standard-user test account compromised at attempt 212; access limited to that user's own records.	Observed
Demonstrated (retest)	Per-IP limit works: failures 1 to 10 in a minute return 401 (allowed); from the 11th, 429. A 25-IP run of 400 attempts against one account (16 per IP, 8 per IP per minute, 2 minutes; measured 200 per minute) drew no delay or challenge (see reconciliation below).	Observed
Conditional, not demonstrated	Compromise of an HR-admin account by the same method. Whether any real HR-admin password is guessable was not tested.	Inferred from role privileges
Severity score (demonstrated impact)	AV:N/AC:L/PR:N/UI:N/S:U/C:L/I:L/A:N	**6.5 Medium**
Conditional scenario (HR-admin compromise)	AV:N/AC:L/PR:N/UI:N/S:U/C:H/I:H/A:N	**9.1 Critical** (conditional)

Why the conditional Critical scenario is plausible

- **Observed:** the missing per-account control is a property of the endpoint, so it applies to every account that signs in there, including HR-admin accounts.
- **Observed:** HR-admin and standard-user roles are distinct, with HR-admin functions enforced server-side (112 endpoint/role combinations, Section 2.3). HR-admin accounts exist in the tenants tested.
- **Inferred from role privileges, not tested:** an HR-admin account holds elevated access to employee, payroll and bank data in its tenant, so its compromise would raise C and I to High.
- **Reported, not verified:** MFA for HR-admin roles is reported as not enforced; this has not been independently verified by KeyPrefer Technologies (SEC-04-c). If the report is accurate, there is no second factor to limit a successful guess.
- **Not tested:** whether any real HR-admin password is guessable. The 9.1 scenario is therefore a statement about exposure if that happens, not a claim that it has or will.

Scoring rationale

Attack Complexity: Low. The vulnerability is the absence of any throttling, delay or challenge. Exploitation needs no preparation, special network position or timing and succeeds repeatably. Password strength decides which accounts are exposed, not how hard the attack is. This is consistent with SEC-03. **C and I: Low (demonstrated), High (conditional).** The severity is scored on demonstrated impact (C:L/I:L, 6.5 Medium). The conditional scenario (C:H/I:H, 9.1 Critical) applies the reasonably expected worst-case impact of compromising a privileged account, which the same weakness makes possible. Versions 1.3 to 1.5 headlined the 9.1 figure; version 1.6 headlines 6.5 so the severity does not imply a demonstrated admin takeover. Readers whose policy scores the reasonably expected worst case may use 9.1. The remediation priority (7 days) follows the conditional scenario because the missing per-account control is observed; either figure justifies action.

What the score does not say. CVSS base scores do not model likelihood. The chance that an HR-admin password is guessed depends on password hygiene (addressed by breached-password checks and MFA) and on whether production adds gateway-level controls (Section 2.4). The conditional Critical scenario drives urgency of the per-account fix and an interim MFA control for privileged roles.

Remediation

- Layered limits: per-IP at the gateway and per-account in the application, with progressive delays rather than hard lockouts. The per-account limit must apply regardless of the number of source IPs.
- Interim control: enforce MFA for HR-admin roles now, and cover privileged accounts first in the fix.
- Block known-breached passwords at registration and reset. Alert on failed-login bursts and on one source attempting many accounts.
- Verification: with the client's written authorization, repeat the guessing test against a dedicated HR-admin test account with a seeded weak password. This converts the conditional scenario into a demonstrated or refuted result. The score will be reassessed once HR-admin MFA is verified as enforced.

Retest (Oct 5, 2026, build 2026.10.2, staging)

Test	Observed	Result
1 IP, 1 account, 60 attempts	Attempts 1 to 10 returned 401 (allowed); attempts 11 to 60 returned 429 with Retry-After	Pass
1 IP, 40 different account identifiers	429 from the 11th failure within one minute	Pass
25 IPs, 1 account, 400 attempts (16 per IP, 8 per IP per minute, 2 min)	No delay or challenge; 401 returned throughout; no IP exceeded 8 failures in any minute	Fail (no per-account control)
Legitimate login during tests	Succeeded	Pass

Note on the 40-identifier row: the 40 identifiers were the three provisioned test accounts plus 37 non-existent identifiers. The row is read only as evidence that the per-IP limit counts failures across identifiers, not as a test of 40 real accounts.

SAMPLE

Reconciling the 25-IP result with the per-IP limit

Per-IP limiter behavior. The limiter allows 10 failures per IP per minute: in the 1-IP test (synthetic detail) failures 1 to 10 returned 401 and attempts 11 onward returned 429. The 10th failure is therefore allowed, and the limit triggers on the 11th. The 25-IP run returned 401 throughout because it was paced under that limit. **Synthetic test parameters:** 25 distinct source addresses, each seen by the gateway as a separate client; 16 attempts per IP (400 in total); 8 attempts per IP per minute, so no IP exceeded 8 failures in any minute (below the 10 allowed); total duration 2 minutes (about 3.3 attempts per second overall, inside the 15 requests-per-second cap).

Measured versus theoretical rate. The measured rate in this test was **200 attempts per minute** (400 attempts in 2 minutes; 25 IPs x 8 per minute). The figure of **250 attempts per minute** is a **theoretical upper bound**, not a measurement: 25 IPs x 10 allowed failures per minute, with the 11th refused. It was not tested. (Had the 10th failure itself been refused, the bound would be 25 x 9 = 225 per minute.) Whether the per-IP limiter uses a fixed or a rolling window was **not verified**. If it is a fixed window, an attacker could place attempts either side of a window boundary and briefly exceed 10 per IP; under a rolling window the bound holds in any 60 seconds. Both rates scale linearly with the number of source addresses.

What the result shows: with every IP kept below the per-IP limit, one account absorbed 400 failed attempts with no delay or challenge, so no per-account control is present. The test does not measure the fastest possible distributed attack, and the 25 addresses came from one harness, so real-world pacing may differ.

Residual risk while open. The principal attack path, distributed guessing against a single account, remains open: no per-account control was observed, even with every IP below the per-IP limit. Until one is deployed and verified, practical protection for privileged accounts is password strength, HR-admin MFA once enforced and verified, breached-password checks and alerting. Production gateway or WAF controls, if present, may reduce exposure but were not tested.

4. Retest summary and closure

Builds tested: 2026.09.3 (initial) and 2026.10.2 (retest), both on staging. Retest was performed October 5–6, 2026 by the original lead assessor using the original reproduction steps, plus negative and valid-flow controls. A finding is marked **Fixed on staging** only when the original exploit was observed to fail and valid-flow controls succeeded, not inferred from a status code. Where a recommended control was not tested (SEC-01) or an exploit path remains (SEC-04), the status says so. Retest results apply to staging; deployment to production was not verified. The SEC-04 rescorings in 1.2 and 1.3 changed its score and severity only; 1.4 to 1.6 add explanation only. No retest results were altered.

ID	Finding	Original	Retest outcome (staging)	Follow-up
SEC-01	OTP bypass	Critical 9.1	Core flaw fixed; lockout validation incomplete	Verify SEC-01-a and -b; monitor reset volume
SEC-02	BOLA	High 8.1	Fixed on staging	CI tenant-isolation tests (Backend, 60 days); confirm in production
SEC-03	Timing enumeration	Medium 5.3	Fixed on staging	Confirm in production
SEC-04	Login rate limiting	Medium 6.5 demonstrated; conditional Critical 9.1	Partially fixed; open	Per-account control and HR-admin MFA (Identity, 7 days); retest on request

Open validation items

Ref	Item	Why it matters	Status
SEC-01-a	Failure counter shared across source IPs and sessions	If counted per IP or session, an attacker can rotate to avoid the five-attempt limit	Not tested
SEC-01-b	Rate limit on issuing new reset challenges per account	Without it, fresh codes can be requested to bypass the challenge limit	Not tested
SEC-04-a	Per-account control against distributed guessing	Main residual attack path; no control observed in retest	Open
SEC-04-b	Guessing test against an authorized HR-admin test account	Converts the conditional HR-admin scenario into a demonstrated or refuted result	Not tested; needs client authorization
SEC-04-c	HR-admin MFA status: reported as not enforced; verify independently, then enforce as interim control	Score to be reassessed once MFA enforcement is verified	Reported absent; not independently verified
SEC-04-d	Per-IP limiter window type (fixed or rolling) and boundary behavior	Sets the real per-IP ceiling; 250 per minute is a theoretical bound (10 allowed, 11th refused), 200 per minute was measured	Not verified

Production verification plan

Production was out of scope and no result in this report applies to it. The checks below each need a written scope extension and rules of engagement, and should use client-provisioned production test tenants and accounts so no real customer data is touched.

Check	Findings	Method
Fixes deployed: build 2026.10.2 changes present in production	All	Client supplies release record; KeyPrefer Technologies confirms by behavior test
Gateway/WAF rate limits and production-only login controls	SEC-04	Compare configuration with staging; repeat single-IP and distributed login tests

SAMPLE

Check	Findings	Method
Tenant scoping on employee sub-resources	SEC-02	Repeat the 64-combination cross-tenant test with two production test tenants
Email dispatch path and response timing	SEC-03	Repeat 200-sample timing test, ideally from more than one vantage point
Reset flow, counter sharing and challenge-issuance limit	SEC-01	Repeat retest steps plus SEC-01-a and SEC-01-b

Compliance context

This report is point-in-time technical evidence that can support security-control reviews such as SOC 2 and ISO/IEC 27001. It is not an attestation or certification and does not guarantee ongoing security.

5. Hardening observations (not scored)

Observations describe weaknesses that did not cross a security boundary or whose exploitability was not demonstrated. They carry no CVSS score.

Ref	Observation	Recommendation / status
INF-01	The login response returned a three-level reporting chain, internal employee code, salary band and a client-side permissions object, more than the UI needs. Data stayed within the caller's own tenant, no authorization boundary was crossed, and server-side function authorization was confirmed independent of the permissions object (BFLA, Section 2.3).	Use explicit response DTOs. Resolved on staging: fields removed and permissions object reduced to UI hints (verified Oct 5, 2026).
INF-02	HSTS present but without includeSubDomains.	Add the directive after confirming subdomain readiness.
INF-03	One API 500 response disclosed the framework version.	Return generic error bodies; log details server-side.

Appendix A. Rating methodology

Rating	CVSS v3.1	Meaning
Critical	9.0 to 10.0	Exploitable remotely with little effort; severe data or system compromise. Fix immediately (target 7 days) and apply compensating controls until fixed.
High	7.0 to 8.9	Significant compromise likely; fix within 14 days.
Medium	4.0 to 6.9	Exploitation needs preconditions or impact is limited; fix within 30 days.
Low	0.1 to 3.9	Minor impact or hard to exploit; fix within 90 days.
Info	n/a	Hardening observation; no demonstrated security impact.

Where a finding carries both a demonstrated score and a conditional higher-impact scenario (SEC-04), the demonstrated score is the finding's severity and both are shown; the remediation target may follow the conditional scenario, and the report says so explicitly.

Appendix B. Evidence index and reproducibility

ID	Content	Shown in this sample
E-01 to E-03	SEC-01 requests, responses, negative controls	Yes (abbreviated excerpts)
E-04, E-05	SEC-02 cross-tenant read and write	Yes (abbreviated excerpts)
E-06	SEC-04 attempt log with timestamps	No; summarized in the finding text
E-07	SEC-03 timing samples (200 per class)	Summary statistics only (partly synthetic); raw samples not shown
E-08	INF-01 response diff before and after fix	No
E-09	Retest captures for all findings (including SEC-04 25-IP per-IP counts and timing)	No; results summarized in the findings

This is a fictional sample: no evidence pack exists. In a real engagement, full captures, timing datasets and scripts are delivered in an encrypted evidence pack. The excerpts here are abbreviated and synthetic, so a reader cannot reproduce the findings from this document alone.

Finding	What a reader can check from this sample	What a real evidence pack adds
SEC-01	Request/response excerpt (E-01, E-02); negative controls (E-03); retest status codes	Full captures, audit-log extracts, reproduction script
SEC-02	Read and write excerpts (E-04, E-05); endpoints tested (64) and affected (2)	Full endpoint/method matrix and per-request results
SEC-03	Summary statistics (medians as reported; min, p5, p95, max synthetic), sample counts, test conditions	Raw samples, percentile table, histogram, timing script
SEC-04	Attempt counts, timings, retest table, both CVSS vectors	Timestamped attempt log (E-06), retest captures, scripts
INF-01	Description and resolution status	Before/after response diff (E-08)

Appendix C. Numerical and cross-section consistency checks (v1.6)

Item	Check	Result
SEC-04 run rate	1,000 attempts in 250 s = 4.0 per second; cap is 15 per second	Consistent
SEC-04 attempt 212	212 / 4.0 = 53 s	Consistent

Item	Check	Result
SEC-04 25-IP test	400 / 25 = 16 attempts per IP over 2 minutes. At 8 attempts per IP per minute, each IP remained below the limit.	Consistent with the retest results; 3.3 per second overall is under the 15 per second cap
SEC-04 40 identifiers	3 provisioned accounts + 37 non-existent identifiers	Consistent; read as limiter behavior only
SEC-03 separation	371 - 139 = 232 ms; medians sit between each class's min and max	Consistent (min/max synthetic)
SEC-02 counts	2 of 64 affected; 64 of 64 returned 404 on retest	Consistent
CVSS recomputation	SEC-01 9.1; SEC-02 8.1; SEC-03 5.3 (3.7 if AC:H); SEC-04 9.1 derived, 6.5 demonstrated, 4.8 in v1.0 (AC:H)	All match displayed vectors and bands
Severity counts	One Critical (SEC-01), one High (SEC-02), two Medium (SEC-03; SEC-04 on demonstrated impact, with a conditional Critical scenario)	Consistent across Sections 1, 3, 4
Dates	Sep 14–18, 2026 is Monday to Friday (5 days); retest Oct 5–6; issue 28 Oct	Consistent
Removed claim	"Fixed within 8 working days" (no notification date in the sample)	Removed

End of sample report. Fictional client; synthetic data.

SAMPLE