Bachao.AI

BACHAO.AI · SECURITY ASSESSMENT REPORT

A****** — Production Security Assessment (Sample · Client Identity Redacted)

Vulnerability Assessment & Penetration Test — SaaS Platform (Black-box + Grey-box)

Sample deliverable — a real Bachao.AI engagement with the client name, domains and IP addresses redacted (shown as asterisks). Multi-tenant SaaS platform.

TLP:AMBER · Confidential

Prepared for: **A******

Prepared by: **Bachao.AI — Dhisattva AI Pvt Ltd (DPIIT-Recognised Startup)**

Assessment: **4–9 August 2026** · Remediation re-tested: **16 August 2026**

Certificate No.: **BVPT-SAMPLE-2026-0001** · Classification: **TLP:AMBER — Confidential**

✔ Remediation verified — all Critical, High & Medium findings resolved

Assessment identified 8 findings; after re-test no open Critical / High / Medium issue remains.

Sample deliverable — do not scan

This is a masked demonstration report; the certificate reference is **not a live verification**. Real client certificates carry a scannable, OTP-authenticated QR at **bachao.ai/verify**.

Contents

The report is organised as follows.

1 Document Control

Report title	A****** — Production Security Assessment (Sample · Client Identity Redacted)
Client	A****** Technologies Private Limited
Certificate / Reference No.	BVPT-SAMPLE-2026-0001

Engagement	Vulnerability Assessment & Penetration Test — SaaS Platform (Black-box + Grey-box)
Assessment type	Black-box + Authenticated Grey-box, with remediation verification
Assessment window	4–9 August 2026
Remediation re-test	16 August 2026
Report version	2.0 — Final (post-retest)
Classification	CONFIDENTIAL — restricted to A****** & Bachao.AI
Assessor	Bachao.AI — Dhisattva AI Pvt Ltd (DPIIT-Recognised Startup)
Standards	CERT-In-aligned · ISO/IEC 27001:2022 · SOC 2 · GDPR · DPDP · OWASP · CVSS v3.1
Lead assessor	Shouvik Mukherjee — Lead Security Assessor
Verification	sample deliverable — reference not live

2 About This Assessment

Bachao.AI (Dhisattva AI Pvt Ltd) is a DPIIT-recognised cybersecurity and fraud-prevention company. This engagement was an authorised Vulnerability Assessment & Penetration Test (VAPT) of A******'s internet-facing production surface, followed by an independent verification of the client's remediation. Testing combined unauthenticated black-box probing with authenticated grey-box testing (capture-and-replay of the application's real API calls under a valid session). The assessment was conducted to a **CERT-In-aligned methodology** and its findings are mapped to **ISO/IEC 27001:2022, SOC 2, GDPR and the DPDP Act 2023** control frameworks, following the **OWASP** Top 10 and API Security Top 10 testing standards and scored with **CVSS v3.1** — meeting Indian and international security-assessment benchmarks.

This report is confidential and intended for A******. It may be shared with the client's customers, investors, auditors and regulators as independent evidence of the security testing performed.

3 Executive Summary

This is a sample deliverable: a real Bachao.AI engagement with the client name, domains and IP addresses redacted (shown as asterisks) to protect confidentiality — the findings, evidence and structure are the genuine report. Bachao.AI performed a black-box and authenticated grey-box VAPT of the A****** production platform — the customer web application, the multi-tenant API and the email domain. Testing was weighted to access control (RBAC, object-level authorization and tenant isolation), API security and misconfiguration, and every automated signal was manually verified before inclusion. The initial assessment identified one Critical (a CORS misconfiguration enabling credentialed cross-origin reads and account takeover), one High (a cross-tenant BOLA/IDOR on organization objects), three Medium (a spoofable email domain from missing DMARC enforcement, missing security headers, and no rate limiting on the authentication endpoint) and three Low hardening issues. Strong controls were also confirmed: enforced authentication, SPF hard-fail with DKIM, and no confirmed injection across 120 checks. A****** remediated every finding; on the included re-test all issues were verified resolved — no open Critical, High or Medium remains — and the all-green security certificate was issued for the assessed scope. This masked sample illustrates the depth, evidence and end-to-end arc of a Bachao.AI engagement.

1	1	3	3	0
CRITICAL	HIGH	MEDIUM	LOW	INFO

ORIGINAL — AT ASSESSMENT
4–9 August 2026

1 Critical | 1 High | 3 Medium | 3 Low

→

CURRENT — AFTER RE-TEST
16 August 2026

✔ 0 open Critical / High / Medium

Resolved & re-verified

4 Scope of Work

The following assets were in scope. Testing was non-destructive and read-only: no data was deleted or altered, no active denial-of-service or load testing was performed, and no real personal data was exfiltrated. All activity is recorded in a timestamped forensic log.

HOST	ROLE	RESOLVES TO	STATUS
dashboard.a******.ai	Production web application (customer console)	20.***.***.**	In scope
api.prod.a******.com	Production API (multi-tenant)	20.***.***.**	In scope
a******.com	Marketing site + primary email domain	185.***.***.***	In scope

Authenticated testing. Grey-box testing used client-provided tenant test accounts across two organizations and role levels (member and admin), plus an API bearer token, to exercise RBAC, object-level authorization (IDOR/BOLA) and tenant isolation.

5 Methodology

The engagement follows **OWASP ASVS Levels 1–2**, the **PTES** phases and **NIST SP 800-115** technical testing guidance, aligned to **CERT-In** reporting expectations.

1 · Reconnaissance & asset discovery

Passive discovery (Certificate Transparency, DNS) and active enumeration of the in-scope hosts, sub-domains, API base URLs and technology stack; TLS/transport and email-security posture.

2 · Automated assessment (curated, AI-native)

Our platform runs a curated suite covering the OWASP Top 10 and OWASP API Security Top 10 in isolated micro-VMs; every candidate finding is re-validated by a second pass to strip false positives before a human sees it.

3 · Authenticated grey-box (capture-replay)

Real sessions at each privilege level are captured and replayed with access-control tampering (BOLA/IDOR, function-level authorization, mass-assignment) — the class of flaw black-box structurally cannot find.

4 · Senior-led manual exploitation

Hands-on verification of business logic, tenant isolation and result-integrity abuse; every reported finding is manually reproduced with request/response evidence.

5 · Risk rating & reporting

Each finding is scored with CVSS v3.1, mapped to OWASP/CWE and the client's compliance frameworks, and written with a concrete, developer-ready remediation.

6 · Remediation re-test

After the client applied fixes, every finding was live re-tested and confirmed closed with before/after evidence.

Tooling by phase

PHASE	REPRESENTATIVE TOOLING / CHECKS
Reconnaissance	Certificate Transparency (crt.sh), DNS / DNS-over-HTTPS, port & service enumeration, TLS inspection
Transport & TLS	openssl s_client, cipher/protocol & certificate-chain analysis
Email security	SPF / DKIM / DMARC record analysis via authoritative DoH lookups
Web & API	OWASP Top 10 / API Top 10 curated suite, HTTP method & header inspection, CORS probing
Authenticated	Session capture-replay harness with BOLA/IDOR/BFLA and mass-assignment tampering
Verification	Manual reproduction with curl / HTTP proxy; CVSS v3.1 scoring

6 Risk Rating Methodology

Severity is derived from the **CVSS v3.1** base score, adjusted by the assessor for business context (data sensitivity, exposure, exploitability).

SEVERITY	CVSS BAND	MEANING & EXPECTED ACTION
CRITICAL	9.0 – 10.0	Immediate exploitation risk; direct compromise of data or systems. Fix immediately.
HIGH	7.0 – 8.9	Serious impact; realistic exploitation path. Fix urgently.
MEDIUM	4.0 – 6.9	Moderate impact or conditional exploitation. Fix in the normal cycle.
LOW	0.1 – 3.9	Limited impact; defense-in-depth / hardening.
INFO	0.0	No direct security impact; informational / good-practice note.

7 Test Coverage

The assessment executed **622 individual checks** across the categories below; “observations” are candidate signals raised for analyst review (the confirmed subset, after verification, appears in §Detailed Findings).

CATEGORY	WHAT IT CHECKS	CHECKS	OBS.
Access Control (RBAC / IDOR / BOLA)	Object- and function-level authorization, tenant isolation, privilege escalation.	88	2
API Security (OWASP API Top 10)	Auth on API routes, mass-assignment, excessive data exposure, rate limiting.	104	4
Attack-Surface & Recon	Subdomain discovery, DNS/TLS posture, technology fingerprinting, exposed services.	74	7
Authentication & Session	Login flows, token handling, session fixation, cookie flags, MFA behaviour.	96	3
Injection & Input Handling	SQLi, NoSQLi, command/template injection, XSS, deserialization, SSRF.	120	0

CATEGORY	WHAT IT CHECKS	CHECKS	OBS.
Security Misconfiguration	Headers, CORS, verbose errors, default/exposed endpoints, cache policy.	82	5
Transport / DNS / Email Hygiene	TLS config, HSTS, SPF/DKIM/DMARC, DNSSEC, open-resolver checks.	58	3
Total		**622**	**24**

8 Findings Summary

ID	SEVERITY	FINDING	HOST	CVSS
F-01	CRITICAL	CORS reflects any Origin with Allow-Credentials on authenticated API	api.prod.a******.com	9.3
F-02	HIGH	Broken Object-Level Authorization (BOLA/IDOR) on organization objects	api.prod.a******.com	8.1
F-03	MEDIUM	No DMARC enforcement — domain spoofable for phishing	a******.com	4.3
F-04	MEDIUM	Missing security response headers (HSTS / CSP / X-Frame-Options)	dashboard.a******.ai	5.3
F-05	MEDIUM	No rate limiting on authentication endpoints	api.prod.a******.com	5.3
F-06	LOW	Session cookie missing SameSite / hardened flags	dashboard.a******.ai	3.7
F-07	LOW	Technology / version information disclosed in headers	dashboard.a******.ai / api.prod.a******.com	3.7
F-08	LOW	DNSSEC not implemented	a******.com	3.7

9 Detailed Findings

F-01 CRITICAL ✓ FIXED

CORS reflects any Origin with Allow-Credentials on authenticated API

api.prod.a******.com · /api/v1/* (authenticated)

CVSS 9.3 · AV:N/AC:L/PR:N/UI:R/S:C/C:H/I:H/A:N · A05:2021 — Security Misconfiguration · CWE-942

Background. When an API reflects an arbitrary request Origin back in Access-Control-Allow-Origin and also sets Access-Control-Allow-Credentials: true, any attacker-controlled web page a logged-in user visits can make credentialed cross-origin requests and READ the authenticated responses — session tokens, profile and tenant data — enabling session/token theft and account takeover.

Description. The authenticated production API reflected any supplied Origin (e.g. https://attacker.example) in Access-Control-Allow-Origin together with Access-Control-Allow-Credentials: true, and returned readable JSON to the cross-origin caller. Chained with a logged-in session this yields full account takeover (CORS → credentialed read → session hijack → takeover).

Impact. Critical: a hostile page could exfiltrate a logged-in user's tokens and tenant data and act as that user. For a multi-tenant platform this is a cross-customer data-exposure and account-takeover risk.

Evidence.

```
$ curl -sI -H 'Origin: https://attacker.example' -b '<session>' \
    https://api.prod.a******.com/api/v1/users/me
  HTTP/2 200
  access-control-allow-origin: https://attacker.example
  access-control-allow-credentials: true
  content-type: application/json   (body returned to cross-origin caller)
```

Remediation. Replace Origin reflection with a strict server-side allowlist of trusted first-party origins (dashboard.a******.ai). Never combine a reflected/wildcard Origin with Allow-Credentials: true. Apply uniformly across root, error and data responses.

Compliance. ISO 27001 A.8.26 (application security) · SOC 2 CC6.6 · DPDP S.8(4) reasonable security safeguards

References. MITRE CWE-942 · OWASP A05:2021 · PortSwigger: Exploiting CORS misconfigurations

Re-test. Re-tested after fix: the API now returns Access-Control-Allow-Origin only for the dashboard.a******.ai allowlist; disallowed origins receive no ACAO header and cannot read responses. Cross-origin credentialed read is blocked. RESOLVED.

F-02 HIGH ✔ FIXED

Broken Object-Level Authorization (BOLA/IDOR) on organization objects

api.prod.a******.com · /api/v1/organization/{orgId}/*

CVSS 8.1 · AV:N/AC:L/PR:L/UI:N/S:U/C:H/I:H/A:N · API1:2023 — Broken Object Level Authorization · CWE-639

Background. Object-level authorization must confirm that the authenticated principal owns (or is entitled to) the specific object referenced by an ID. Where the server trusts a client-supplied object/tenant ID without an ownership check, any authenticated user can read or modify another tenant's objects — a BOLA / IDOR.

Description. An authenticated member of one organization could substitute a different {orgId} in /api/v1/organization/{orgId}/members and /.../billing and receive another tenant's data. The server authorized on authentication alone, not object ownership.

Impact. High: cross-tenant read of member lists and billing data. A malicious or curious tenant could enumerate other customers' records — a serious multi-tenant isolation and privacy breach.

Evidence.

```
Logged in as tenant A (orgId=1001):
$ curl -s -b '<session A>' https://api.prod.a******.com/api/v1/organization/1002/members
  HTTP/2 200  → returns tenant B's member records (orgId 1002 ≠ 1001)
```

> **Remediation.** Enforce object ownership server-side on every object-scoped route: derive the tenant from the session/token and verify the requested {orgId} belongs to it; reject mismatches with 403/404. Add automated multi-tenant authorization tests.

Compliance. ISO 27001 A.8.3 (information access restriction) · SOC 2 CC6.3 · DPDP S.8(4) reasonable security safeguards

References. OWASP API Security Top 10 — API1:2023 · MITRE CWE-639

> **Re-test.** Re-tested after fix: object-scoped routes now verify the {orgId} against the session principal; cross-tenant requests return 403 and no data. RESOLVED.

F-03 MEDIUM ✔ FIXED

No DMARC enforcement — domain spoofable for phishing

a******.com · DNS (email authentication)

CVSS 4.3 · AV:N/AC:L/PR:N/UI:R/S:U/C:N/I:L/A:N · A05:2021 — Security Misconfiguration · CWE-16

Background. SPF and DKIM authenticate mail, but only a DMARC policy tells receivers what to do with mail that fails — and produces reporting. Without an enforcing DMARC record a domain can be spoofed in phishing that appears to come from the company.

Description. The primary domain published SPF (hard-fail) and DKIM but had no enforcing DMARC policy (p=none / absent), leaving the domain spoofable in phishing campaigns against customers and staff.

Impact. Medium: for a platform whose customers trust its email, attackers could send convincing spoofed mail (invoices, password resets) from the real domain. Rated Medium as exploitation requires the recipient to act on the spoofed message.

Evidence.

```
$ dig +short TXT _dmarc.a******.com
  (no enforcing record — p=none/absent)
$ dig +short TXT a******.com | grep spf
  "v=spf1 ... -all"   (SPF hard-fail present)
```

Remediation. Publish an enforcing DMARC record (start p=quarantine with rua/ruf reporting, move to p=reject once aligned). Confirm SPF and DKIM alignment for all sending sources.

Compliance. ISO 27001 A.5.14 (information transfer) · SOC 2 CC6.7 · DPDP —

References. RFC 7489 (DMARC) · CERT-In email-security advisories

Re-test. Re-tested after fix: _dmarc.a******.com now publishes p=quarantine with aggregate reporting; alignment confirmed for all senders. RESOLVED.

F-04 MEDIUM ✔ FIXED

Missing security response headers (HSTS / CSP / X-Frame-Options)

dashboard.a******.ai · all application responses

CVSS 5.3 · AV:N/AC:L/PR:N/UI:N/S:U/C:L/I:N/A:N · A05:2021 — Security Misconfiguration · CWE-693

Background. Defence-in-depth headers reduce the blast radius of other bugs: HSTS forces HTTPS, Content-Security-Policy constrains script sources (limiting XSS), and X-Frame-Options / frame-ancestors prevents clickjacking.

Description. Application responses lacked HSTS, a Content-Security-Policy and X-Frame-Options.

Impact. Medium: increases exploitability of any XSS, enables clickjacking, permits HTTPS downgrade.

Evidence.

```
Response headers on dashboard.a******.ai carried no Strict-Transport-Security, Content-Security-Policy or X-Frame-Options.
```

Remediation. Add HSTS (max-age ≥ 6 months, includeSubDomains), a scoped Content-Security-Policy, X-Frame-Options: DENY (or frame-ancestors 'none'), X-Content-Type-Options: nosniff.

Compliance. ISO 27001 A.8.26 · SOC 2 CC6.6 · DPDP —

References. OWASP Secure Headers Project

Re-test. Re-tested after fix: HSTS, CSP and X-Frame-Options are present on all app responses. RESOLVED.

F-05 MEDIUM ✔ FIXED

No rate limiting on authentication endpoints

api.prod.a******.com · /api/v1/auth/login

CVSS 5.3 · AV:N/AC:L/PR:N/UI:N/S:U/C:N/I:N/A:L · API4:2023 — Unrestricted Resource Consumption · CWE-307

Background. Login and OTP endpoints without throttling permit credential-stuffing and brute-force at scale, and can be abused to exhaust resources.

Description. The login endpoint accepted rapid repeated attempts with no rate limit or lockout.

Impact. Medium: enables credential-stuffing / brute-force against user accounts.

Evidence.

```
200+ sequential login attempts within a minute were accepted without throttling or lockout.
```

Remediation. Add per-IP and per-account rate limiting with progressive back-off / lockout and CAPTCHA on repeated failures; alert on anomalous login volume.

Compliance. ISO 27001 A.8.5 (secure authentication) · SOC 2 CC6.1 · DPDP —

References. OWASP API Security Top 10 — API4:2023 · MITRE CWE-307

Re-test. Re-tested after fix: login now rate-limits and locks out after repeated failures. RESOLVED.

F-06 LOW ✔ FIXED

Session cookie missing SameSite / hardened flags

dashboard.a******.ai

CVSS 3.7 · AV:N/AC:H/PR:N/UI:N/S:U/C:L/I:N/A:N · A05:2021 — Security Misconfiguration · CWE-1275

Background. SameSite, Secure and HttpOnly on session cookies reduce CSRF and token-theft exposure.

Description. The session cookie did not set SameSite (and Secure was inconsistent across paths).

Impact. Low: marginally increased CSRF / cross-site token exposure.

Evidence.

```
Set-Cookie for the session lacked a SameSite attribute on some responses.
```

Remediation. Set SameSite=Lax (or Strict), Secure and HttpOnly on all session cookies.

Compliance. ISO 27001 A.8.26 · SOC 2 CC6.6 · DPDP —

References. OWASP Session Management Cheat Sheet

Re-test. Re-tested after fix: session cookies now set SameSite=Lax, Secure and HttpOnly. RESOLVED.

F-07 LOW ✔ FIXED

Technology / version information disclosed in headers

dashboard.a******.ai / api.prod.a******.com

CVSS 3.7 · AV:N/AC:H/PR:N/UI:N/S:U/C:L/I:N/A:N · A05:2021 — Security Misconfiguration · CWE-200

Background. Server / framework banners help an attacker map the stack to known CVEs.

Description. Responses disclosed server / framework identifiers and an X-Powered-By banner.

Impact. Low: information leakage that assists targeting; needs a matching CVE to exploit.

Evidence.

```
Server and X-Powered-By headers present on application / API responses.
```

Remediation. Suppress or genericise Server / X-Powered-By and framework version headers at the edge.

Compliance. ISO 27001 A.8.9 (configuration management) · SOC 2 CC6.1 · DPDP —

References. MITRE CWE-200

Re-test. Re-tested after fix: version banners are suppressed at the edge. RESOLVED.

F-08 LOW ✓ FIXED

DNSSEC not implemented

a******.com

CVSS 3.7 · AV:N/AC:H/PR:N/UI:N/S:U/C:N/I:L/A:N · A05:2021 — Security Misconfiguration · CWE-350

Background. Without DNSSEC, resolvers cannot cryptographically validate DNS answers for the domain.

Description. The zone is not signed with DNSSEC.

Impact. Low: marginally increased exposure to DNS spoofing / cache-poisoning (defence-in-depth).

Evidence.

```
No DS/RRSIG records observed for a******.com.
```

Remediation. Enable DNSSEC signing and publish the DS record at the registrar.

Compliance. ISO 27001 A.5.14 · SOC 2 CC6.7 · DPDP —

References. RFC 4033

Re-test. Re-tested after fix: the zone is DNSSEC-signed and the DS record is published. RESOLVED.

10 Positive Assurance (Controls Verified)

The following controls were tested and found to be working correctly — evidence that the platform is soundly engineered, not merely an absence of findings:

- Authentication is enforced on data endpoints — unauthenticated API calls are rejected; the API does not leak data to anonymous callers.
- SPF is published with a hard-fail (-all) on the primary domain, and DKIM signing is correctly configured — a strong email-authentication baseline.
- No open DNS resolver and no exposed administrative interfaces were found on the in-scope hosts.
- Injection testing (SQLi/NoSQLi/command/template/SSRF, 120 checks) surfaced no confirmed injection — input handling on the tested surface is sound.

11 Excluded / Verified False Positives

In the interest of an honest, vendor-grade report, the following automated signals were investigated and **excluded** after manual verification:

- **SQL Injection 'possible' on the search API (heuristic signals).** Verified false positive — an injection payload and a benign baseline returned identical responses with no differential behaviour; the flag is a scanner heuristic on uniform error handling, not an injection.
- **Open redirect on the marketing site.** Not reproduced — the redirect target is validated against an allowlist; the flagged parameter could not be steered off-domain.

12 Compliance Mapping

The methodology and every finding are tested, rated and reported against the frameworks the client's enterprise customers and auditors ask about — one engagement, evidence reusable across all.

FRAMEWORK	HOW THIS ENGAGEMENT ALIGNS
CERT-In (India)	Report produced in CERT-In-aligned format with full evidence and proof-of-testing; findings carry CVSS v3.1, affected asset, reproduction and remediation.
DPDP Act 2023 — Schedule I	Findings referenced to the S.8(4) reasonable-security-safeguards obligation for a platform processing personal data.
OWASP Top 10 & API Security Top 10	Mapped to A05:2021 misconfiguration and API1:2023 BOLA / API4:2023 resource consumption, among others.
ISO/IEC 27001:2022 Annex A	Referenced to A.8.3 access restriction, A.8.9 configuration, A.8.26 application security and A.5.14 information transfer.
SOC 2 (Trust Services)	Findings cross-referenced to CC6.x logical-access and CC7.x monitoring criteria to support a SOC 2 evidence trail.

13 Standards Coverage & Control Cross-Reference

OWASP Top 10:2021 — coverage & result

OWASP TOP 10:2021	COVERAGE	RESULT
A01:2021 Broken Access Control	Tested (authenticated BOLA/IDOR/BFLA)	No issue
A02:2021 Cryptographic Failures	Tested (TLS, cert, transport)	No issue (strong TLS)
A03:2021 Injection	Tested (SQLi/XSS/command/template)	No issue
A04:2021 Insecure Design	Reviewed (business logic, workflow)	No issue
A05:2021 Security Misconfiguration	Tested (headers, CORS, DMARC, exposure)	Findings — see §Findings
A06:2021 Vulnerable Components	Tested (version fingerprinting, known-CVE)	Informational (version disclosure)
A07:2021 Auth & Identification Failures	Tested (auth, session, MFA, origin)	Low (defense-in-depth)
A08:2021 Software & Data Integrity	Reviewed	No issue
A09:2021 Logging & Monitoring	Out of test scope (internal)	N/A
A10:2021 SSRF	Tested (candidate endpoints)	No issue (production)

Per-finding control cross-reference

Each finding mapped to the specific ISO/IEC 27001:2022 Annex A control, SOC 2 Common Criteria and DPDP Act 2023 obligation an auditor will ask about.

FINDING	OWASP	CWE	ISO 27001:2022	SOC 2	DPDP 2023
F-01	A05:2021 — Security Misconfiguration	CWE-942	A.8.26 (application security)	CC6.6	S.8(4) reasonable security safeguards

FINDING	OWASP	CWE	ISO 27001:2022	SOC 2	DPDP 2023
F-02	API1:2023 — Broken Object Level Authorization	CWE-639	A.8.3 (information access restriction)	CC6.3	S.8(4) reasonable security safeguards
F-03	A05:2021 — Security Misconfiguration	CWE-16	A.5.14 (information transfer)	CC6.7	—
F-04	A05:2021 — Security Misconfiguration	CWE-693	A.8.26	CC6.6	—
F-05	API4:2023 — Unrestricted Resource Consumption	CWE-307	A.8.5 (secure authentication)	CC6.1	—
F-06	A05:2021 — Security Misconfiguration	CWE-1275	A.8.26	CC6.6	—
F-07	A05:2021 — Security Misconfiguration	CWE-200	A.8.9 (configuration management)	CC6.1	—
F-08	A05:2021 — Security Misconfiguration	CWE-350	A.5.14	CC6.7	—

14 Remediation Roadmap

No open findings requiring remediation.

Recommended remediation SLAs (CERT-In / RBI-aligned)

SEVERITY	RECOMMENDED REMEDIATION SLA	GUIDANCE
CRITICAL	≤ 7 days	Emergency change; fix before anything else.
HIGH	≤ 14 days	Prioritised fix in the current cycle.
MEDIUM	≤ 30 days	Normal change management.
LOW	≤ 90 days	Next maintenance window.
INFO	Best-effort	Good-practice hardening.

One post-fix re-test is included in this engagement. On confirmation that the Critical, High and Medium findings are remediated and re-tested clean, the signed all-green security certificate is issued for the assessed scope.

15 Next Steps & Contact

1. Review this report with your engineering and compliance teams.
2. Remediate the findings using the developer-ready guidance in §Detailed Findings, prioritised by the Remediation Roadmap SLAs.
3. Notify Bachao.AI when the fixes are deployed; we run the **included post-fix re-test**.
4. On a clean re-test, the signed **all-green security certificate** is issued for the assessed scope — the artefact you share with customers, investors and auditors.

Point of contact. Bachao.AI · Dhisattva AI Pvt Ltd · ceo@bachao.ai · Sample deliverable — certificate reference not live

16 Annex — Evidence Pack (Live Test Captures)

Verbatim commands and responses captured during testing, reproducing the confirmed findings and the key positive-assurance controls. Redacted of any sensitive values.

CORS credentialed read on authenticated API (F-01, before fix)

```
curl -sI -H 'Origin: https://attacker.example' -b '<session>' \
  https://api.prod.a******.com/api/v1/users/me
  access-control-allow-origin: https://attacker.example
  access-control-allow-credentials: true
```

Cross-tenant BOLA (F-02, before fix)

```
curl -s -b '<session tenant A: orgId 1001>' \
  https://api.prod.a******.com/api/v1/organization/1002/members
  HTTP/2 200  → returns tenant B (orgId 1002) member records
```

Post-remediation re-test (all findings)

```
Re-test window: disallowed CORS origins receive no ACAO header; cross-tenant object requests return 403;
_dmarc publishes p=quarantine; HSTS/CSP/X-Frame-Options present. No open Critical / High / Medium remains.
```

17 Annex — Test Register (Proof of Testing)

Representative enumeration of the 622 checks executed, by category, with the observed result. “Observation(s)” denotes a candidate signal raised for analyst review; the verified subset appears in §Detailed Findings.

Access Control (RBAC / IDOR / BOLA)

Security checks in this category.

TEST	RUNS	RESULT
Cross-tenant object access on /api/v1/organization/{id}/*	34	Observation — verified BOLA (fixed)
Horizontal privilege escalation (peer user records)	22	No issue
Vertical privilege escalation (member → admin)	18	No issue
Function-level authorization on admin API routes	14	No issue

Security Misconfiguration

Security checks in this category.

TEST	RUNS	RESULT
CORS Origin reflection with Allow-Credentials	12	Observation — verified Critical (fixed)
Security response headers (HSTS / CSP / X-Frame-Options)	20	Observation — hardened
Verbose error / stack-trace disclosure	16	No issue

Transport / DNS / Email Hygiene

Security checks in this category.

TEST	RUNS	RESULT
DMARC enforcement policy	6	Observation — added (fixed)
SPF hard-fail on primary domain	4	Pass (correct)
DKIM signing present	4	Pass (correct)
DNSSEC signing	3	Observation — recommended

18 Annex — Tools & Techniques

- Bachao.AI VAPT engine (curated OWASP/API test suite, isolated micro-VM execution)
- Authenticated capture-replay harness (SPA session capture, BOLA/IDOR tampering)
- Manual tooling: curl, OpenSSL, DNS-over-HTTPS resolvers, HTTP proxy

19 Annex — Glossary

TERM	MEANING
BOLA / IDOR	Broken Object Level Authorization — a user reaches another user's object by id.
BFLA	Broken Function Level Authorization — access to functions above one's privilege.
CVSS v3.1	Common Vulnerability Scoring System — the standard 0–10 severity score.
CORS	Cross-Origin Resource Sharing — browser policy controlling cross-site reads.
DMARC	Domain-based Message Authentication — email anti-spoofing policy.
SSRF	Server-Side Request Forgery — coercing the server to make attacker-chosen requests.
Soft-404	A page that returns HTTP 200 for URLs that do not exist.

20 Annex — Standards & References

- OWASP Top 10:2021 · OWASP API Security Top 10:2023
- OWASP Application Security Verification Standard (ASVS) v4.0, Levels 1–2
- PTES — Penetration Testing Execution Standard
- NIST SP 800-115 — Technical Guide to Information Security Testing
- CVSS v3.1 — Common Vulnerability Scoring System
- CERT-In — reporting format and methodology alignment
- ISO/IEC 27001:2022 Annex A · SOC 2 (AICPA TSC) · GDPR Article 32 · DPDP Act 2023

21 Revision History

VERSION	DATE	NOTE
1.0	12 August 2026	Initial assessment report issued (Critical/High/Medium findings open).
2.0	16 August 2026	Post-remediation re-test — all findings resolved; all-green certificate issued.

22 Disclaimer & Confidentiality

This assessment reflects the state of the in-scope systems during the test window (4–9 August 2026). Security is time-variant: changes to the environment after the test window may introduce new risks. A penetration test demonstrates the presence of vulnerabilities, not their absence; a clean result reduces but does not eliminate risk.

Bachao.AI (Dhisattva AI Pvt Ltd) is a technology provider. The compliance mappings are technical observations to support A******'s own compliance work and are not a legal opinion. This document is confidential to A****** and Bachao.AI.

Classification — TLP:AMBER. This report is marked under the Traffic Light Protocol (TLP) and contains security-sensitive information, including exploitable vulnerability detail. It may be shared only within A****** on a need-to-know basis and must not be disclosed to third parties without Bachao.AI's written consent. Store and transmit it accordingly.

Shouvik Mukherjee — Lead Security Assessor

Bachao.AI · Dhisattva AI Pvt Ltd · 16 August 2026

Certificate No. BVPT-SAMPLE-2026-0001 · Sample — reference not live