Bachao.AI

BACHAO.AI · SECURITY ASSESSMENT REPORT

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

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

Sample deliverable — a real Bachao.AI engagement with the client name, domains and IP addresses redacted (shown as asterisks). AI supply-chain / logistics platform — application, API and corporate domain.

TLP:AMBER · Confidential

Prepared for: **M**********

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

Assessment: **15 August 2026**

Certificate No.: **BVPT-SAMPLE-2026-0004** · Classification: **TLP:AMBER — Confidential**

✓ No Critical or High severity findings

8 findings identified (0 Critical · 0 High · 1 Medium · 5 Low · 2 Informational). Remediation guidance and a path to the all-green certificate are provided.

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	M******** — Production Security Assessment (Sample · Client Identity Redacted)
Client	M******** Technologies Private Limited
Certificate / Reference No.	BVPT-SAMPLE-2026-0004

Engagement	Vulnerability Assessment & Penetration Test — Production (Black-box + Grey-box)
Assessment type	Black-box + Authenticated Grey-box
Assessment window	15 August 2026
Report version	1.0 — Final
Classification	CONFIDENTIAL — restricted to M******** & 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 M********'s internet-facing production surface, with prioritised remediation guidance. 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 M********. 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 M******** production platform — the application, its API and the corporate domain. Testing followed OWASP ASVS L1–L2, the OWASP Top 10 and API Security Top 10, PTES and NIST practice, and every automated result was manually verified before inclusion. The production platform is in good security posture: authenticated testing with an Admin account found no broken access control, IDOR or tenant-isolation issue; authentication and input validation are enforced; and transport uses strong modern cryptography. The findings are a single Medium domain-level email-authentication (DMARC) misconfiguration and a set of Low / Informational hardening items — no Critical or High severity issues were confirmed on the production scope. This masked sample illustrates the depth, evidence and structure of a Bachao.AI engagement.

0	0	1	5	2
CRITICAL	HIGH	MEDIUM	LOW	INFO

Posture: no Critical or High severity vulnerabilities were confirmed. The residual items are addressed with concrete remediation and a defined path to the all-green certificate (remediation + one re-test).

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
app.m********.com	Production application / login (grey-box)	20.***.***.***	In scope
api.m********.com	Production API	20.***.***.***	In scope
m********.com	Corporate / marketing domain	185.***.***.*	In scope
www.m********.com	Marketing site (www)	185.***.***.*	In scope

Authenticated testing. Grey-box testing used a client-provided Admin session on the production application; the authenticated crawler exercised the production application and API (phases 0–7). Full-depth tenant-isolation testing requires a second throwaway tenant, recommended for a follow-up.

> Note on coverage terminology: passive resilience checks (e.g. rate-limit / abuse headers) are counted in Test Coverage but were performed read-only — no active denial-of-service or load testing was carried out.

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.

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 **240 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.
API Security	Security checks in this category.	180	2
Reconnaissance & Discovery	Security checks in this category.	60	4
Total		**240**	**6**

8 Findings Summary

ID	SEVERITY	FINDING	HOST	CVSS
F-01	MEDIUM	DMARC not enforced — two conflicting policy records	m********.com (email domain)	5.3
F-02	LOW	SSH (port 22) exposed to the public internet	app.m********.com / api.m********.com	3.7

ID	SEVERITY	FINDING	HOST	CVSS
F-03	LOW	DNSSEC not implemented	m********.com	3.7
F-04	LOW	Server / framework version disclosed in headers	api.m********.com / app.m********.com	3.7
F-05	LOW	Missing HTTP security headers / clickjacking protection	app.m********.com / api.m********.com	3.1
F-06	LOW	API accepts requests with no Origin header (defense-in-depth gap)	api.m********.com	3.1
F-07	INFO	Arbitrary Host header accepted (no reflection observed)	api.m********.com	—
F-08	INFO	Response caching / ETag hygiene	app.m********.com	—

9 Detailed Findings

F-01 MEDIUM OPEN

DMARC not enforced — two conflicting policy records

m********.com (email domain) · _dmarc.m********.com (DNS TXT)

CVSS 5.3 · AV:N/AC:L/PR:N/UI:N/S:U/C:N/I:L/A:N · A05:2021 — Security Misconfiguration · CWE-16

Background. DMARC binds SPF/DKIM results to the visible From: domain and tells receivers what to do with mail that fails alignment. RFC 7489 requires exactly one DMARC record per domain; when more than one is published, conforming receivers apply no DMARC policy at all.

Description. The domain publishes two DMARC TXT records (p=none and p=quarantine). Because more than one record is present, receivers treat the domain as having no DMARC policy, so enforcement does not apply — and one record requests no enforcement regardless.

Impact. An attacker can spoof @m********.com sender addresses in phishing to the company's customers and partners; recipients will not reject or quarantine the forged mail — a direct business-email-compromise and brand-abuse risk for a logistics platform whose partners act on emailed instructions.

Evidence.

```
$ dig +short TXT _dmarc.m********.com   (via DoH)
  "v=DMARC1; p=none;"
  "v=DMARC1; p=quarantine; adkim=r; aspf=r; rua=mailto:dmarc_rua@o**************.net;"
=> two records present => receivers apply NO DMARC policy (RFC 7489 §6.6.3)
```

> **Remediation.** Publish a single DMARC record and enforce it. Start at p=quarantine with aggregate reporting, then p=reject once SPF/DKIM alignment is confirmed. Remove the stray p=none record. Recommended: v=DMARC1; p=reject; adkim=s; aspf=s; rua=mailto:dmarc@m********.com; fo=1;

Compliance. ISO 27001 A.5.14 (information transfer) · SOC 2 CC6.7 · DPDP S.8(4) reasonable safeguards

References. RFC 7489 §6.6.3 · MITRE CWE-16

F-02 LOW ● OPEN

SSH (port 22) exposed to the public internet

app.m********.com / api.m********.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-284

Background. Management interfaces exposed to the whole internet are a standing target: SSH sees continuous credential-stuffing, and any future SSH-daemon CVE becomes remotely reachable. Defense-in-depth restricts such interfaces to trusted networks.

Description. TCP/22 is reachable from any internet host on the production nodes.

Impact. Increases attack surface and automated brute-force noise; a single weak/compromised key or an SSH CVE becomes internet-reachable.

Evidence.

```
Port scan: 22/tcp open on the production Azure nodes (20.***.***.***).
```

> **Remediation.** Restrict SSH to a bastion/VPN CIDR via an Azure NSG, disable password auth (keys only), prefer Azure Bastion / brokered access over a public listener.

Compliance. ISO 27001 A.8.9 (configuration mgmt) · SOC 2 CC6.1 · DPDP S.8(4) reasonable safeguards

References. MITRE CWE-284 · NIST SP 800-115

F-03 LOW ● OPEN

DNSSEC not implemented

m********.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 verify that a DNS answer genuinely came from the zone owner, so answers can be forged in networks lacking other protections.

Description. The zone is not signed with DNSSEC; resolvers cannot validate answers cryptographically.

Impact. Marginally increased exposure to DNS cache-poisoning / spoofing. Defense-in-depth.

Evidence.

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

> **Remediation.** Enable DNSSEC signing at the DNS provider and publish the DS record at the registrar.

Compliance. ISO 27001 A.5.14 (information transfer) · SOC 2 CC6.7 · DPDP S.8(4) reasonable safeguards

References. RFC 4033 · MITRE CWE-350

F-04 LOW OPEN

Server / framework version disclosed in headers

api.m********.com / app.m********.com

CVSS 3.7 · AV:N/AC:H/PR:N/UI:N/S:U/C:L/I:N/A:N · A06:2021 — Vulnerable & Outdated Components · CWE-200

Background. Version banners let an attacker map the exact software in use and go straight to a matching public exploit, shortening the reconnaissance phase.

Description. HTTP responses advertise server/framework identifiers useful for CVE targeting.

Impact. Information leakage that assists targeting; exploitation requires a matching CVE (AC:H).

Evidence.

```
$ curl -sI https://app.m********.com/ | grep -i server
  server: nginx/1.27.5
$ curl -sI https://api.m********.com/ | grep -i server
  server: uvicorn
```

Remediation. Suppress or genericise the Server header at the edge/app (e.g. nginx server_tokens off).

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

References. MITRE CWE-200

F-05 LOW ● OPEN

Missing HTTP security headers / clickjacking protection

app.m********.com / api.m********.com

CVSS 3.1 · AV:N/AC:H/PR:N/UI:N/S:U/C:N/I:L/A:N · A05:2021 — Security Misconfiguration · CWE-1021

Background. Modern browsers rely on response headers to enforce framing, transport and content policies. Absent X-Frame-Options and a CSP frame-ancestors directive, the authenticated dashboard can be framed by a hostile site for clickjacking; absent HSTS, a first request can be stripped to HTTP.

Description. Neither the application nor the API returns Strict-Transport-Security, Content-Security-Policy, X-Frame-Options, X-Content-Type-Options, Referrer-Policy or Permissions-Policy. The app therefore has no anti-clickjacking control.

Impact. Clickjacking of the authenticated dashboard (UI-redress), MIME-sniffing exposure, and no HSTS pinning of HTTPS. Individually Low; collectively a hardening gap a security questionnaire will flag.

Evidence.

```
$ curl -sI https://app.m********.com/ | grep -iE 'strict-transport|content-security|x-frame|x-content-type|referrer-policy|permissions-policy'
  (no matching headers returned)
=> HSTS, CSP, X-Frame-Options, X-Content-Type-Options, Referrer-Policy, Permissions-Policy all ABSENT
```

> **Remediation.** Add at the edge: Strict-Transport-Security (max-age>=31536000; includeSubDomains), Content-Security-Policy with frame-ancestors 'self', X-Frame-Options: DENY, X-Content-Type-Options: nosniff, a strict Referrer-Policy and a Permissions-Policy.

Compliance. ISO 27001 A.8.9 / A.8.26 (application security requirements) · SOC 2 CC6.1 · DPDP S.8(4)

References. OWASP Secure Headers Project · MITRE CWE-1021

F-06 LOW OPEN

API accepts requests with no Origin header (defense-in-depth gap)

api.m********.com

CVSS 3.1 · AV:N/AC:H/PR:N/UI:N/S:U/C:N/I:L/A:N · A05:2021 — Security Misconfiguration · CWE-346

Background. Validating the Origin on browser-facing endpoints is a cheap extra barrier. It is not a primary control here because the API authenticates with a bearer token, not an ambient cookie, so cross-site request forgery does not apply.

Description. The authenticated API served a valid-bearer request that carried no Origin header.

Impact. Low — no direct exploit given bearer auth (not CSRF-exploitable); noted for completeness.

Evidence.

```
$ curl -s -o /dev/null -w '%{http_code}' -H 'authorization: Bearer <redacted>' \
    https://api.m********.com/api/partners?limit=500   # no Origin header
  200
```

Remediation. Validate/allowlist Origin on browser-facing API routes; keep bearer auth as primary.

Compliance. ISO 27001 A.8.26 (application security requirements) · SOC 2 CC6.6 · DPDP —

References. MITRE CWE-346

F-07 INFO OPEN

Arbitrary Host header accepted (no reflection observed)

api.m********.com

A05:2021 — Security Misconfiguration · CWE-20

Background. Host-header injection becomes exploitable only when the injected value is reflected into links, redirects or password-reset emails. No such reflection was observed, so this is recorded as informational.

Description. The server returns HTTP 200 to a request with a spoofed Host header.

Impact. Informational — confirm the Host is not used in email links or absolute-URL generation.

Evidence.

```
$ curl -sI -H 'Host: evil-bachao-test.com' --resolve ... https://api.m********.com/
  HTTP/2 200   (no Location/Link reflecting the injected host)
```

Remediation. Enforce an allowlist of expected Host values at the ingress; reject unknown Hosts (421).

Compliance. ISO 27001 A.8.9 · SOC 2 CC6.1 · DPDP —

References. MITRE CWE-20

F-08 INFO ● OPEN

Response caching / ETag hygiene

app.m********.com

A05:2021 — Security Misconfiguration · CWE-525

Background. ETags derived from filesystem inode/size can leak minor backend metadata; explicit cache directives avoid caching authenticated content.

Description. Cache-Control is broad and ETag values expose backend file metadata on some responses.

Impact. Informational hardening.

Evidence.

```
$ curl -sI https://app.m********.com/ | grep -iE 'cache-control|etag'
  cache-control: no-cache
  etag: "6a7be87b-3d0"
```

Remediation. Set explicit Cache-Control on authenticated responses and normalise/disable weak ETags.

Compliance. ISO 27001 A.8.9 · SOC 2 CC6.1 · DPDP —

References. MITRE CWE-525

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:

- Authenticated API enforces authentication — requests without a valid token are rejected (HTTP 401).
- Pagination is validated and capped — an oversized limit is rejected (HTTP 422); no unbounded mass-export.
- TLS certificate uses strong modern cryptography — ECDSA with SHA-384 (P-256).
- Authenticated grey-box (Admin session, phase 7) — the access-control candidate signals raised by the engine (IDOR / privilege) were manually reviewed and cleared: no broken object/function-level authorization, IDOR or tenant-isolation break was reproduced on the production application.

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:

- **XML external-entity / 'XML bomb' on m********.com.** Engine flagged application/xml accepted with HTTP 200 on the marketing host, but no evidence of entity expansion or XML parsing; unconfirmed and excluded (recommend a quick confirmation). Note: the demo/staging cluster's higher-severity engine signals (SQLi, Adminer, SSRF, CSRF) were also verified as false positives and are documented in the separate Demo/Staging Advisory, not this production report.

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 format following CERT-In-aligned methodology with full evidence and proof-of-testing; findings carry CVSS v3.1, affected asset, reproduction and remediation.
SOC 2 Type II (AICPA TSC)	Findings mapped to the Security / Common Criteria (CC6 logical access, CC7 operations) — penetration-test evidence for a SOC 2 trail.
ISO/IEC 27001:2022 Annex A	Referenced to A.8.9 configuration, A.8.24 cryptography, A.8.26 application security and A.5.14 information transfer for the ISO 27001 programme.
GDPR Article 32	Security-of-processing observations (transport, spoofing exposure) for EU-facing assurance.
DPDP Act 2023 (India)	Findings referenced to the S.8(4) reasonable-security-safeguards obligation for the personal and commercial data the platform processes.
OWASP · CVSS · PTES · NIST	Testing backbone: OWASP Top 10 & API Security Top 10, ASVS L1–L2, PTES phases, NIST SP 800-115; every finding scored with CVSS v3.1.

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-16	A.5.14 (information transfer)	CC6.7	S.8(4) reasonable safeguards
F-02	A05:2021 — Security Misconfiguration	CWE-284	A.8.9 (configuration mgmt)	CC6.1	S.8(4) reasonable safeguards
F-03	A05:2021 — Security Misconfiguration	CWE-350	A.5.14 (information transfer)	CC6.7	S.8(4) reasonable safeguards
F-04	A06:2021 — Vulnerable & Outdated Components	CWE-200	A.8.9 (configuration mgmt)	CC6.1	—
F-05	A05:2021 — Security Misconfiguration	CWE-1021	A.8.9 / A.8.26 (application security requirements)	CC6.1	S.8(4)
F-06	A05:2021 — Security Misconfiguration	CWE-346	A.8.26 (application security requirements)	CC6.6	—
F-07	A05:2021 — Security Misconfiguration	CWE-20	A.8.9	CC6.1	—
F-08	A05:2021 — Security Misconfiguration	CWE-525	A.8.9	CC6.1	—

14 Remediation Roadmap

SEVERITY	TIMELINE	FINDINGS	ITEMS
MEDIUM	Within 2–4 weeks	F-01	DMARC not enforced — two conflicting policy records
LOW	Next maintenance cycle	F-02, F-03, F-04, F-05, F-06	SSH (port 22) exposed to the public internet; DNSSEC not implemented; Server / framework version disclosed in headers; Missing HTTP security headers / clickjacking protection; API accepts requests with no Origin header (defense-in-depth gap)
INFO	Optional / good practice	F-07, F-08	Arbitrary Host header accepted (no reflection observed); Response caching / ETag hygiene

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.

DMARC — two conflicting records (F-01)

```
dig +short TXT _dmarc.m********.com   # via DoH
  "v=DMARC1; p=none;"
  "v=DMARC1; p=quarantine; adkim=r; aspf=r; rua=mailto:dmarc_rua@o**************.net;"
```

Positive control — authentication enforced (HTTP 401)

```
curl -s -o /dev/null -w '%{http_code}' https://api.m********.com/api/partners?limit=500
  401   # no token => rejected
```

Positive control — pagination capped (HTTP 422)

```
curl ... -H 'authorization: Bearer <redacted>' 'https://api.m********.com/api/partners?limit=100000'
  422   # oversized limit rejected — no unbounded mass-export
```

Positive control — strong TLS certificate

```
openssl s_client -connect api.m********.com:443 </dev/null 2>/dev/null | openssl x509 -noout -text | grep -i
'Signature Algorithm'
  Signature Algorithm: ecdsa-with-SHA384   (P-256)
```

Missing security headers (F-05)

```
curl -sI https://app.m********.com/ | grep -iE 'strict-transport|content-security|x-frame|x-content-type|
referrer-policy|permissions-policy'
  (no output — all six absent)
```

17 Annex — Test Register (Proof of Testing)

Representative enumeration of the 240 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.

API Security

Security checks in this category.

TEST	RUNS	RESULT
Authentication enforced on data endpoints (no-token rejection)	24	Pass (correct)
Pagination bounds / mass-export guard on list endpoints	18	Pass (correct)
Broken object-level authorization (IDOR / BOLA) — authenticated	20	No issue
Origin validation on browser-facing API routes	10	Observation — defense-in-depth gap

Security Misconfiguration

Security checks in this category.

TEST	RUNS	RESULT
Security response headers (HSTS / CSP / X-Frame-Options)	16	Observation — hardening
Server / framework version banner disclosure	12	Observation — hardening
Management-interface exposure (SSH / TCP-22)	8	Observation — restrict

Transport / DNS / Email Hygiene

Security checks in this category.

TEST	RUNS	RESULT
DMARC enforcement policy	6	Observation — single enforcing record needed
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.

TERM	MEANING
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	15 August 2026	Final production assessment report issued.

22 Disclaimer & Confidentiality

This assessment reflects the state of the in-scope systems during the test window (15 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 M********'s own compliance work and are not a legal opinion. This document is confidential to M******** 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 M******** 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 · 15 August 2026

Certificate No. BVPT-SAMPLE-2026-0004 · Sample — reference not live