SECURITY ASSESSMENT REPORT SAMPLE TLP:AMBER

Web Application
Penetration Testing

PREPARED EXCLUSIVELY FOR

Sample Client (anonymized engagement)

A real engagement, anonymized for sharing: the client's name, application, hosts, people and identifiers are withheld; findings, ratings and figures are as delivered.

DATE OF ISSUANCE	REPORT PERIOD	REVISION	PREPARED BY
October 8, 2026	Q4 2026	Sample edition	N30.ai

STRICTLY CONFIDENTIAL — This document contains sensitive security and technical information. Reproduction, distribution, or disclosure to unauthorized parties without explicit written consent is strictly prohibited.

TABLE OF **CONTENTS**

PAGE	SECTION
03	Executive Summary
04	Cybersecurity Maturity Score
05	Scope & Methodology
06	Risk Matrix
07	Testing Phases
08	Findings Overview
09	Authorship Forgery on Document Upload
10	Proof of Concept: Document Upload
11	Out-of-Scope Buildings Readable by Identifier
12	Proof of Concept: Building Access
13	Entity and Super-Entity Records Readable Globally
14	Proof of Concept: Entity Records
15	Resident Personal Data Readable by Identifier
16	Proof of Concept: Resident Personal Data
17	Identity Asserted by a Client-Supplied Header
18	Proof of Concept: Header Identity
19	Per-Control Permission Flags Readable for Any User
20	Staff Directory Disclosed to Any Caller
21	Interactive API Documentation Exposed Without Authentication
22	Cross-Origin Policy Set to Wildcard
23	Verbose Errors Disclose Internal Detail
24	Stack and Version Disclosure on Every Response
25	Unhandled Exception on a Reachable Parameter
26	Unbounded Result Size and Negative Values Crash the Service
27	Authorization Failures Are Not Centrally Logged
28	Vulnerabilities Overview
29	Exploitability Overview
30	Company Overview

STRICTLY CONFIDENTIAL — This document contains sensitive security and technical information. Reproduction, distribution, or disclosure to unauthorized parties without explicit written consent is strictly prohibited.

EXECUTIVE **SUMMARY**

N30.ai tested the client's accounting application during Q4 2026. Testing covered the application without an account and authenticated with the engagement's own session, following the OWASP Top 10.

Fourteen findings are set out in full in this report, five of them High. Their root cause is shared: the application authenticates a token but does not authorise the object it returns. Any signed-in user can read building, entity and resident records belonging to entities outside their own scope simply by changing a number in the request, and can forge the authorship of a document record. The identity used to build the caller's permission menu is asserted by a header the client supplies.

The application's own defences held up where they were tested: the login, the token's signature, the transport encryption and the tenant boundary all worked, and no SQL injection, cross-site scripting or path traversal was found on the reachable parameters. The weakness sits in authorisation below that layer. All object-level tests were read-only, and the single write test created a clearly named document that was removed afterwards.

A single change carries most of the gain: applying the entity-scope check that already exists on one route to every object detail route clears the five High findings at once. The remaining findings are configuration, error handling and monitoring matters. On that basis, N30.ai rates the application:



RISK RATING

SATISFACTORY
Security meets industry standards. Minimal to no actionable findings identified.

NEEDS IMPROVEMENT
Protects some areas, but moderate changes and configuration updates are required.

NON SATISFACTORY
Significant security deficiencies exist requiring immediate attention.

STRICTLY CONFIDENTIAL — This document contains sensitive security and technical information. Reproduction, distribution, or disclosure to unauthorized parties without explicit written consent is strictly prohibited.

CYBERSECURITY
MATURITY SCORE

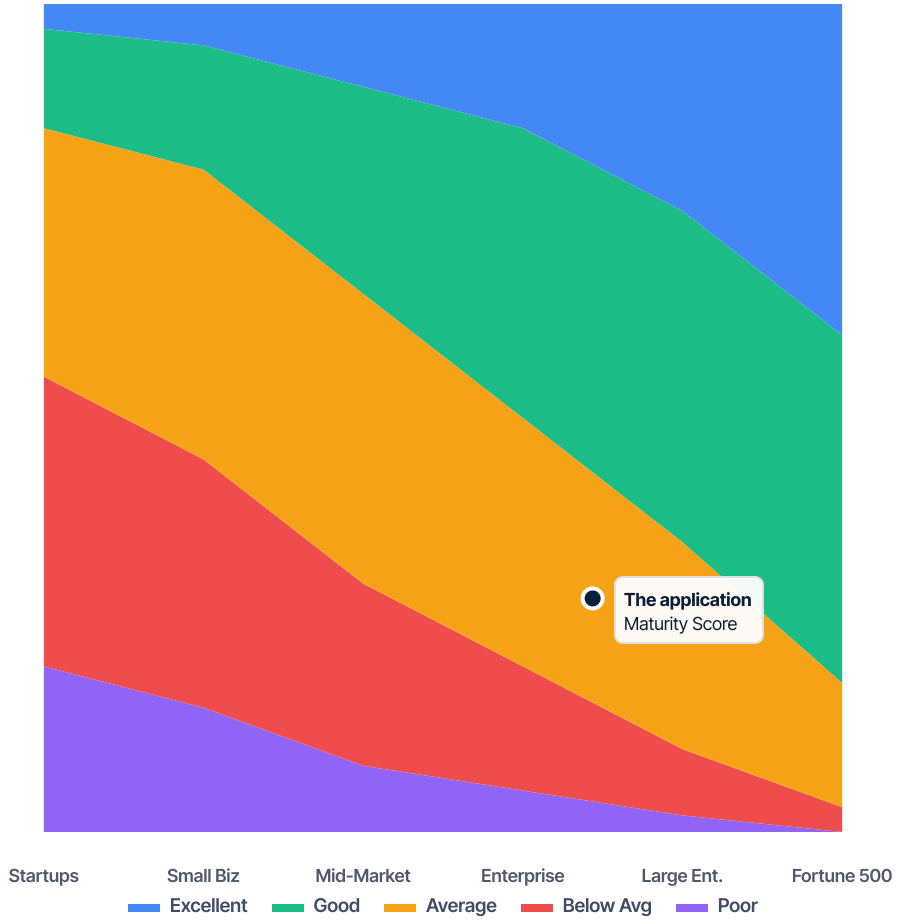
EXECUTIVE TAKEAWAY

The application sits in the lower half of the Average tier, held there by one area: access control. Two changes carry most of the gain, and both are inexpensive.

The client's accounting application shows partial controls: the login, the token's signature and the transport encryption held up, and the tenant boundary is validated. The weakness sits below that layer: the API authenticates a token but does not authorise the object it returns, so building, entity and resident records outside the caller's scope are readable by changing a number in the request. Fourteen findings were confirmed, five of them High, and they cluster around one root cause. Applying the entity-scope check that already exists on one route to every detail route moves the application into the Good tier.

Maturity Scores by Company Size

Illustrative maturity bands by company size, with the application's positioning



UNDERSTANDING THIS CHART

This chart compares cybersecurity maturity across company sizes, with the indicator dot marking the application's position from the findings in this assessment.

STRICTLY CONFIDENTIAL — This document contains sensitive security and technical information. Reproduction, distribution, or disclosure to unauthorized parties without explicit written consent is strictly prohibited.

SCOPE & **METHODOLOGY**

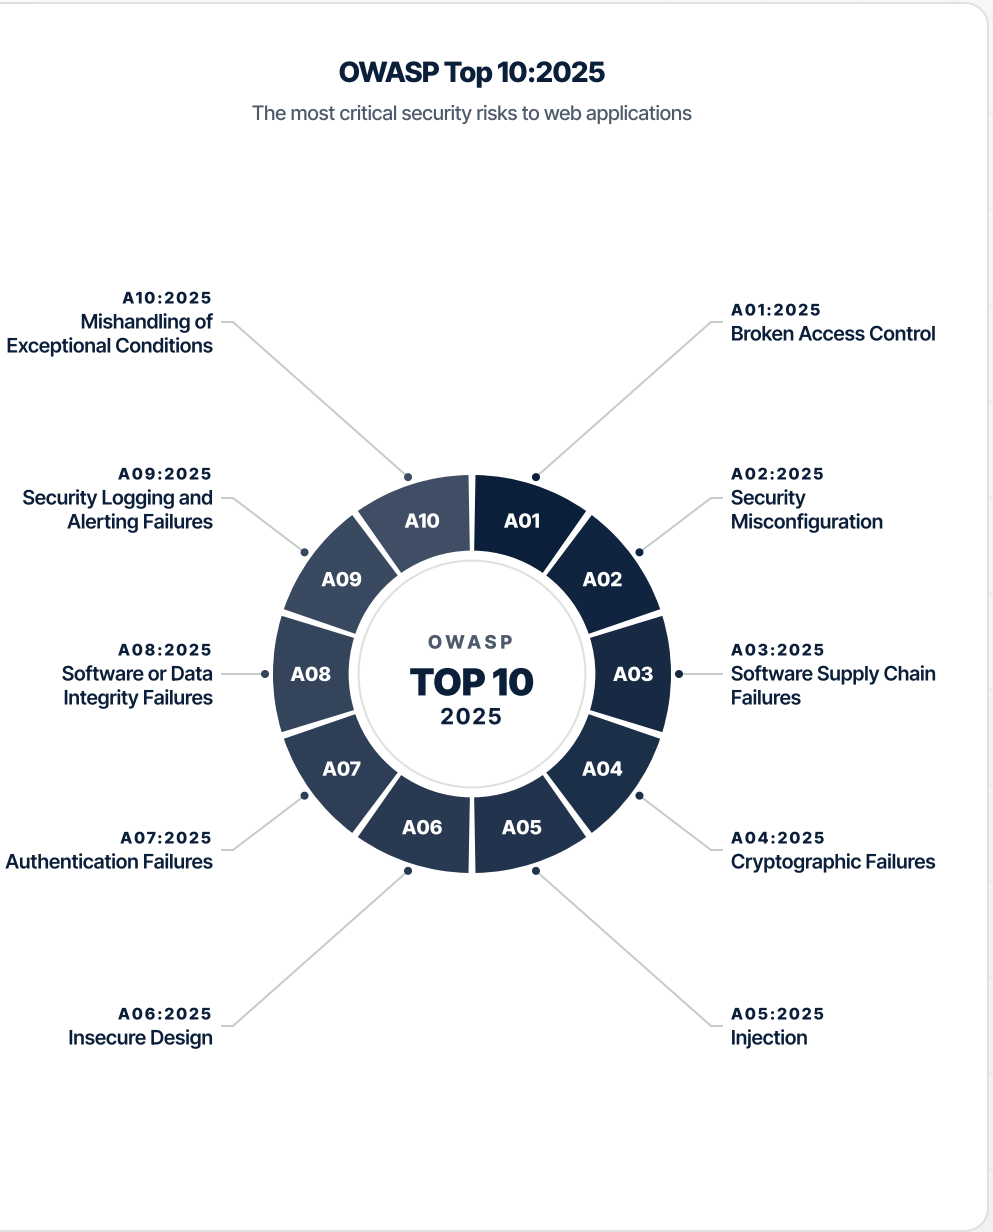
SCOPE HIGHLIGHTS

- ✓ **Accounting web application at https://app.sample-client.example/**
- ✓ **JSON API at https://api.sample-client.example/api/**
- ✓ **Uncredentialed testing, without an account**
- ✓ **Credentialed testing with the engagement's own session token**
- ✓ **Methodology based on the OWASP Top 10:2025, mapped to ASVS 5.0**

The web application penetration test, conducted during Q4 2026, encompassed the client's accounting application (community association accounting), its web front end at https://app.sample-client.example/ and the JSON API that backs it at https://api.sample-client.example/api/.

N30.ai tested the application from two perspectives: uncredentialed, as an anonymous user, and credentialed, authenticated with the engagement's own session. The interactively reachable surface was mapped from the recorded session; the API surface was tested live against the running target.

The methodology is based on the OWASP Top 10:2025, the standard awareness document that represents a broad consensus about the most critical security risks to web applications, and each finding is mapped to the corresponding OWASP ASVS 5.0 requirement.



STRICTLY CONFIDENTIAL — This document contains sensitive security and technical information. Reproduction, distribution, or disclosure to unauthorized parties without explicit written consent is strictly prohibited.

RISK **MATRIX**

PRIORITIZATION STRATEGY

The top-left cells carry the most gain: the most serious findings that are also the least work to correct. The access-control findings sit there, and all five share one root cause.

Each finding is plotted against two dimensions: the **Level of Risk** it carries, and the **Level of Effort Required** to put it right. The severity of each finding is the one shown on its own page; the effort is N30.ai's assessment of the work involved.

Read from the top left, the matrix is a work plan. The findings at the top are the ones that matter most; the ones on the left are the cheapest to close. The access-control findings concentrate there, and they share a single root cause: the API trusts the caller. Correcting that one layer clears several cells at once, which is why it is the first recommendation in this report.

Finding Prioritization Matrix

Every finding in this report, plotted by severity and the effort to remediate



STRICTLY CONFIDENTIAL — This document contains sensitive security and technical information. Reproduction, distribution, or disclosure to unauthorized parties without explicit written consent is strictly prohibited.

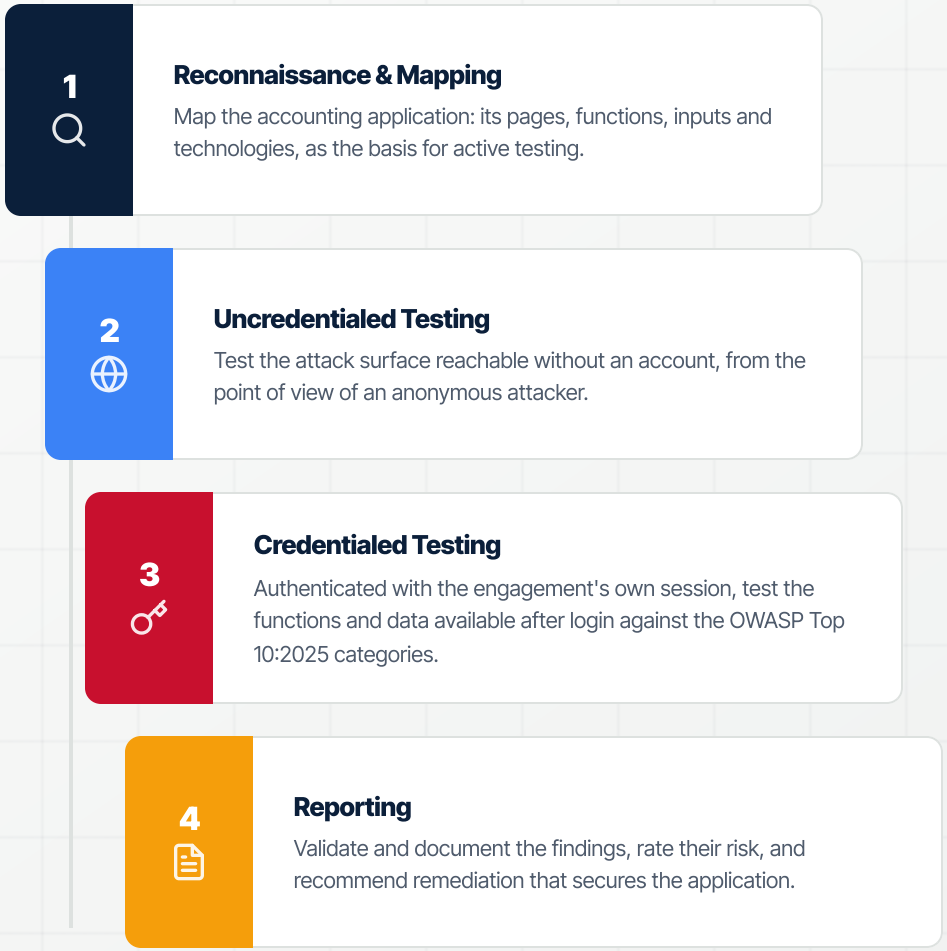
TESTING **PHASES**

METHODOLOGY APPROACH

A four-phase approach based on the OWASP Top 10:2025, exercised against the accounting application both without an account and authenticated.

N30.ai performed the penetration test of the accounting application in four phases, from mapping the application to reporting. The purpose of this approach is to identify and validate as many business-critical vulnerabilities as possible within the defined scope of this engagement.

Testing moved from the attack surface any anonymous visitor can reach to the functions and data available after authentication. Every finding was validated to eliminate false positives and mapped to its OWASP Top 10:2025 category, so the results show which risk categories affect the application.



STRICTLY CONFIDENTIAL — This document contains sensitive security and technical information. Reproduction, distribution, or disclosure to unauthorized parties without explicit written consent is strictly prohibited.

FINDINGS **OVERVIEW**

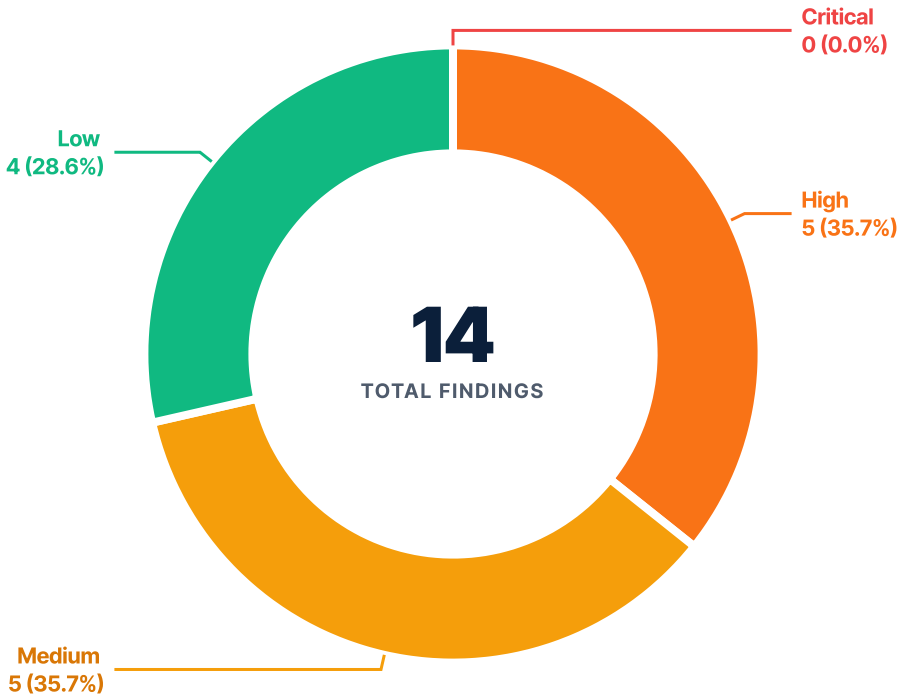
HOW EACH FINDING WAS RATED

Every finding was rated on what an attacker could reach, what the data is worth, how easily it can be exploited, and how much work it takes to correct. Each finding states its rating on its own page, and the identifier used there is the one used in this report.

The pages that follow set out each finding in turn: what was found, what it means in business terms, how it was confirmed, and how to correct it. Findings from the same report are grouped by their identifier, so a finding and its evidence pages sit together.

The assessment identified **14 findings**, each with a page of its own in this report: five High, five Medium and four Low. No Critical finding was identified. The five High findings share a single root cause in the API's authorisation below the authentication layer.

Risk Rating Overview



0	5	5	4
CRITICAL	HIGH	MEDIUM	LOW

STRICTLY CONFIDENTIAL — This document contains sensitive security and technical information. Reproduction, distribution, or disclosure to unauthorized parties without explicit written consent is strictly prohibited.

FINDING 1

Authorship Forgery on Document Upload

RISK LEVEL

High

AFFECTED SECURITY CONTROL

Data Integrity

COMPLIANCE IMPACT

`OWASP A08:2025`

AFFECTED ENDPOINTS

1



7.5



Affected Hosts

`api.sample-client.example`

`/api/globalDocumentFiles/upload GlobalDocumentFile`

Vulnerability Description

N30.ai identified that the Global Documents upload endpoint, `POST /api/globalDocumentFiles/uploadGlobalDocumentFile`, records the author of the stored document from the `createdBy` field supplied in the multipart body. The value is written verbatim; the identity of the signed-in caller is not used.

Authenticated as the operator's own account (user 3318), an upload sent with `"createdBy": 1745` was accepted and stored with `createdBy: 1745`. The record therefore names a different user as its author than the one who created it.

The application's own audit trail records the action under the real session user (the operator's own account), while the document record stores the forged identifier. The two now disagree, so the document's authorship can no longer be established from the record itself. In an accounting platform this undermines the integrity of the audit trail for any document whose author is disputed.

Remediation & Solution Strategy

N30.ai recommends the following corrective actions in order of priority:

1. Take the author from the session, not the body

Set `createdBy`, `modifiedBy` and any similar field from the authenticated principal on the server, and ignore the values sent by the client.

2. Bind the model to a fixed field list

Reject or drop unknown and server-owned fields on write endpoints (an explicit request model or an allow-list), so a caller cannot set authorship, timestamps or identifiers.

3. Apply the same fix to every write endpoint

Enumerate the write endpoints from the API documentation and correct each one that accepts `createdBy`, `modifiedBy`, `userId` or a similar field in its body.

STRICTLY CONFIDENTIAL — This document contains sensitive security and technical information. Reproduction, distribution, or disclosure to unauthorized parties without explicit written consent is strictly prohibited.

PROOF OF **CONCEPT**

● EXPLOIT CONFIRMED

The request below uploaded a document to the staging application while signed in as the operator's own account. The author stored on the record is the value supplied in the request body, not the signed-in user.

EVIDENCE 1: THE REQUEST, EXACTLY AS SENT

Console - upload, signed in as user 3318

Text selectable: copy it straight from this page

```
> curl -s "$API/api/globalDocumentFiles/uploadGlobalDocumentFile" \
    -H "Authorization: Bearer $TOKEN" -H "market: ***" -H "userid: 331
  8" \
    -F 'viewModelJsonFormControlName={"globalDocumentId":1,"name":"ZZ-A
  SVS",\
    "fileName":"a.txt","path":"\\\\[file-server]\\[share]\\[folder]",
    "createdBy":1745}' \
    -F 'file=@a.txt'
```

```
< {"globalDocumentFileId":21,"globalDocumentId":1,"name":"ZZ-ASVS","cre
  atedBy":1745,"modifiedBy":null}
```

The value 1745 is a different user identifier, sent by the client. The stored record names that user as the author.

EVIDENCE 2: THE RESPONSE, AND THE AUTHOR IT STORED

Terminal - the stored record names a different author

```
# uploaded as user 3318, stored as user 1745
Authenticated as user 3318 (the operator's account)
createdBy sent 1745
createdBy stored 1745 <- the client's value
[*] [*] The audit trail logged the real session user, the
operator's account
[*] [*] The record and the audit trail now disagree on the author

[!] Cleanup: the ZZ-ASVS document was deleted through the
Global Documents UI
```

STRICTLY CONFIDENTIAL — This document contains sensitive security and technical information. Reproduction, distribution, or disclosure to unauthorized parties without explicit written consent is strictly prohibited.

FINDING 2

Out-of-Scope Buildings Readable by Identifier

RISK LEVEL

High

AFFECTED SECURITY CONTROL

Access Control

COMPLIANCE IMPACT

`OWASP A01:2025`

AFFECTED ENDPOINTS

1



8.2

RISK SCORE

Affected Hosts

`api.sample-client.example` `/api/buildings/{id}`

Vulnerability Description

The detail route `GET /api/buildings/{id}` returns a building's complete record for any identifier. The caller's entity scope is not consulted: the operator's scope contains 89 entities, yet buildings tied to entities 1, 3, 7, 8, 9, 10, 11, 12, 15 and others outside that scope are returned in full.

A sweep of the identifier range from 1 to 259 returned 240 readable buildings, of which at least 32 map to entities outside the caller's scope. Each record carries the general-ledger entity, bank code, corporate billing identifier, fee and late-fee schedules and the managed address: business data for properties the caller is not entitled to see.

The control exists elsewhere in the same application. The entity-scoped route `/api/balancesheetschedules/type/1/entity/{id}` returns 404 for an entity outside the scope, so object-level access is enforced on that path and simply was not applied to these detail routes. The same class was confirmed on `/api/superentities/{id}` and `/api/entities/{id}`.

Remediation & Solution Strategy

N30.ai recommends the following corrective actions in order of priority:

1. Check ownership on the server for every detail route

Resolve the object, then verify it belongs to an entity in the caller's scope before returning it; return 404 or 403 otherwise. Use the entity check already implemented on the balance-sheet schedule route as the reference.

2. Apply the control centrally, not per controller

Enforce entity scope in one place (a filter, pipeline behaviour or repository layer) so a new detail route cannot be added without it, rather than repeating the check per endpoint.

3. Sweep the identifier space in testing

Add an automated check that walks each object detail route and asserts a 404 for an object outside the test account's scope.

STRICTLY CONFIDENTIAL — This document contains sensitive security and technical information. Reproduction, distribution, or disclosure to unauthorized parties without explicit written consent is strictly prohibited.

PROOF OF **CONCEPT**

EXPLOIT CONFIRMED

The requests below named buildings belonging to entities outside the caller's scope. The application answered in full each time, so the identifier in the address alone decides whose records come back.

EVIDENCE 1: A SINGLE OUT-OF-SCOPE BUILDING, AS SELECTABLE TEXT

Console - one request, scope not consulted

Text selectable: copy it straight from this page

```
> curl -s "$API/api/buildings/1" \
    -H "Authorization: Bearer $TOKEN" -H "market: ***" -H "userid: 3318" \
    | head -c 200

< {"id":1,"name":"...","entity":{...},"address":{...},"glEntityID":1,...}
```

Building 1 belongs to entity 1, which is outside the operator's 89-entity scope. The correct answer is 404 or 403.

EVIDENCE 2: THE IDENTIFIER WALK, AND THE SCOPE BOUNDARY

Terminal - the sweep, and the control that holds

```
# building 1 belongs to entity 1, outside the caller's scope
buildings swept identifiers 1 to 259
readable 240 of 259
out of scope at least 32 map to entities outside the scope
[+] 200 OK with full detail for every out-of-scope building
[*] [*] Control that held:
/api/balancesheetschedules/type/1/entity/500 -> 404
[*] [*] The scope check exists elsewhere; it is simply not applied here
```

STRICTLY CONFIDENTIAL — This document contains sensitive security and technical information. Reproduction, distribution, or disclosure to unauthorized parties without explicit written consent is strictly prohibited.

FINDING 3

Entity and Super-Entity Records Readable Globally

RISK LEVEL

High

AFFECTED SECURITY CONTROL

Access Control

COMPLIANCE IMPACT

`OWASP A01:2025`

AFFECTED ENDPOINTS

2



8.0

RISK SCORE

Affected Hosts

`api.sample-client.example`

`/api/superentities/{id}`

`/api/entities/{id}`

Vulnerability Description

The detail routes `GET /api/superentities/{id}` and `GET /api/entities/{id}` return records for identifiers outside the caller's entity scope. A sweep of `/api/superentities/{id}` returned 113 of 119 identifiers, and `/api/entities/{id}` returned entities that are not in the caller's scope selector.

These records describe the management structure of the portfolio: entity names, their relationships and identifiers. Combined with the buildings route in F02 and the resident profiles in F04, a signed-in user can reconstruct the organisational and property structure of entities they do not administer.

As with F02, no path-level scope check exists on these routes. The application enforces entity scope on other paths, so this is an inconsistent application of an existing control rather than an absent design.

Remediation & Solution Strategy

N30.ai recommends the following corrective actions in order of priority:

1. Enforce entity scope on the entity detail routes

Verify the requested entity or super-entity is within the caller's scope before returning it, and return 404 or 403 when it is not.

2. Centralise the scope decision

Move the entity-scope check into shared middleware so the entity, super-entity and building routes cannot diverge.

3. Reduce the fields returned

Return only the fields each caller needs for the selected entity, rather than the full record, to limit what a single out-of-scope lookup discloses.

STRICTLY CONFIDENTIAL — This document contains sensitive security and technical information. Reproduction, distribution, or disclosure to unauthorized parties without explicit written consent is strictly prohibited.

PROOF OF **CONCEPT**

● EXPLOIT CONFIRMED

The entity and super-entity detail routes returned records for identifiers outside the caller's scope, with no path-level check.

EVIDENCE 1: THE REQUEST, AS SELECTABLE TEXT

Console - entity detail, scope not consulted

Text selectable: copy it straight from this page

```
> curl -s "$API/api/superentities/1" \
    -H "Authorization: Bearer $TOKEN" -H "market: ***" -H "userid: 331
  8" \
    | head -c 200
```

```
< {"id":1,"name":...,"entities":[...]}
```

The same answer comes back for entity identifiers that are not in the caller's scope selector.

EVIDENCE 2: THE IDENTIFIER COVERAGE

Terminal - coverage of both routes

```
# identifiers outside the scope selector are returned in full
/api/superentities/{id} 113 of 119 identifiers returned
/api/entities/{id} entities outside the scope returned
[*] [*] The caller's own scope lists 89 entities

[!] Correct behaviour: 404 or 403 for an out-of-scope
identifier
```

STRICTLY CONFIDENTIAL — This document contains sensitive security and technical information. Reproduction, distribution, or disclosure to unauthorized parties without explicit written consent is strictly prohibited.

FINDING 4

Resident Personal Data Readable by Identifier

RISK LEVEL

High

AFFECTED SECURITY CONTROL

Access Control

COMPLIANCE IMPACT

`OWASP A01:2025`

AFFECTED ENDPOINTS

2



8.5



Affected Hosts

`api.sample-client.example`

`/api/residents/{id}/profile`

`/api/residents/search`

Vulnerability Description

The route `GET /api/residents/{id}/profile` returns a resident's name, mailing address and, where present, tax identifier for any identifier, with no check against the caller's entity scope. A sample of identifiers returned real people: resident 1 is a named individual with a former account number and a street address.

The name search `/api/residents/search?residentName=smith` returned 59 matching residents across entities and is not filtered by scope, so records can be found by name as well as by identifier. The configuration endpoint `/api/taxprofile` returns a company tax identifier with a named contact, telephone number and email address.

This is the most sensitive of the access-control findings. The data is personal information about residents of properties the caller does not administer, and it is reachable by any signed-in account by changing a number in the request.

Remediation & Solution Strategy

N30.ai recommends the following corrective actions in order of priority:

1. Scope resident lookups to the caller's entities

Verify the resident belongs to an entity in the caller's scope before returning the profile, and apply the same check to the name search.

2. Mask sensitive fields by default

Return tax identifiers and full contact details only to the roles that need them, and mask them otherwise, so a lookup discloses no more than the task requires.

3. Log and alert on bulk lookups

Alert on rapid sequences of resident identifiers from one account, which is the signature of automated enumeration.

STRICTLY CONFIDENTIAL — This document contains sensitive security and technical information. Reproduction, distribution, or disclosure to unauthorized parties without explicit written consent is strictly prohibited.

PROOF OF **CONCEPT**

EXPLOIT CONFIRMED

The requests below returned residents of properties the caller does not administer, by identifier and by name, with no scope check.

EVIDENCE 1: A RESIDENT PROFILE BY IDENTIFIER

Console - resident profile, any identifier

Text selectable: copy it straight from this page

```
> curl -s "$API/api/residents/1/profile" \
    -H "Authorization: Bearer $TOKEN" -H "market: ***" -H "userid: 331
  8" \
    | head -c 200
```

```
< {"residentId":1,"name":"...","mailingAddress":"...","taxId":...}
```

Values are masked on this page; the raw capture is in the report's evidence folder. Resident 1 is a real person returned for a single identifier change.

EVIDENCE 2: COVERAGE, AND THE NAME SEARCH

Terminal - how far the disclosure reaches

```
# residents of any entity, by identifier or by name
by identifier /api/residents/{id}/profile returns any resident
by name /api/residents/search?residentName=smith -> 59 matches
tax profile /api/taxprofile returns a company tax id and contact

[!] Names, addresses and tax identifiers for out-of-scope
residents

[*] [*] No scope filter applies to the search either
```

STRICTLY CONFIDENTIAL — This document contains sensitive security and technical information. Reproduction, distribution, or disclosure to unauthorized parties without explicit written consent is strictly prohibited.

FINDING 5

Identity Asserted by a Client-Supplied Header

RISK LEVEL

High

AFFECTED SECURITY CONTROL

Authentication

COMPLIANCE IMPACT

`OWASP A07:2025`

AFFECTED ENDPOINTS

1



7.0



Affected Hosts

`api.sample-client.example` `/api/useraccess`

Vulnerability Description

The bearer token issued after sign-in is self-contained but carries no user claim: its payload holds only issuer, audience, client identifier and scope. The API therefore derives the caller's function menu from a `userid` header supplied by the client. Changing only that header returns a different user's permission map with the same token.

One token, five header values: `userid: 3318` returned 285 functions, `5` returned 280, `6` returned 109, `9` returned 107, and `999999` returned an error. The menu is entirely client-asserted.

The boundary is important and was measured: on every data endpoint tested, changing the header changed nothing; sixteen endpoints were compared and all were identical. The data is bound to the token; only the menu is not. This is confirmed identity spoofing against the permission-map path, and the step that makes the user directory (F07) and the per-control flags (F06) walkable to find and map a privileged account.

Remediation & Solution Strategy

N30.ai recommends the following corrective actions in order of priority:

1. Bind identity into the token

Include a user claim (subject, object identifier or user principal name) in the token at issue time and derive the caller's identity from it alone.

2. Stop trusting the header for identity

Until identity is in the token, use the header only as a display hint, never to select a permission set; if it is needed, validate it against the token's own claim.

3. Reject mismatched requests

If a header identity disagrees with the token, reject the request, so the mismatch is visible rather than silently honoured.

STRICTLY CONFIDENTIAL — This document contains sensitive security and technical information. Reproduction, distribution, or disclosure to unauthorized parties without explicit written consent is strictly prohibited.

PROOF OF **CONCEPT**

● EXPLOIT CONFIRMED

One bearer token, five values of the userid header. The token carries no user claim, so the permission map returned follows the header the client sends.

EVIDENCE 1: THE REQUEST, AS SELECTABLE TEXT

Console - the permission map follows the header

Text selectable: copy it straight from this page

```
> curl -s "$API/api/useraccess" \
    -H "Authorization: Bearer $TOKEN" -H "market: ***" -H "userid: 6" \
    | wc -c

< 20895
```

Changing only the userid header returned a different user's permission map: 20,895 bytes and 109 functions for user 6, against 53,510 bytes and 285 functions for the operator.

EVIDENCE 2: ONE TOKEN, FIVE PERMISSION MAPS

Terminal - the token carries no user claim

```
# the same token, five header values
userid 3318 53,510 bytes, 285 functions
userid 5 52,495 bytes, 280 functions
userid 6 20,895 bytes, 109 functions
userid 9 20,730 bytes, 107 functions
userid 999999 400 - no such user
[*] [*] Token claims: iss, aud, client_id, scope only - no user
claim
[*] [*] Boundary: on 16 data endpoints the header changed nothing
```

STRICTLY CONFIDENTIAL — This document contains sensitive security and technical information. Reproduction, distribution, or disclosure to unauthorized parties without explicit written consent is strictly prohibited.

FINDING 6

Per-Control Permission Flags Readable for Any User

RISK LEVEL

Medium

AFFECTED SECURITY CONTROL

Access Control

COMPLIANCE IMPACT

`OWASP A01:2025`

AFFECTED ENDPOINTS

1



5.0



Affected Hosts

`api.sample-client.example`

`/api/users/{userId}/navigationlinks/{linkId}/pagecontrols/access`

Vulnerability Description

The route `GET /api/users/{userId}/navigationlinks/{linkId}/pagecontrols/access` returns the control-level permissions of whichever user is named in the path. The caller's own identity is not compared with the user in the path.

For navigation link 304, the operator's account returned `hasAccess: true` for the AddHoliday control, while users 5, 6, 9, 1 and 2 returned false and user 3319 returned true. The user directory (F07) supplies the identifiers, so the flags can be read for every staff member in turn.

No data is changed, but the result is a complete entitlement map of the user base: combined with the directory, it shows who holds which controls, which is the reconnaissance needed to choose a target for an escalation attempt.

Remediation & Solution Strategy

N30.ai recommends the following corrective actions in order of priority:

1. Compare the path user with the caller

Return a user's permission flags only to that user or to an administrator entitled to see them; reject any other caller.

2. Draw the flags from the authenticated identity

Where the caller is asking about themselves, take the identifier from the token rather than the path.

3. Treat the entitlement map as sensitive

Restrict and log requests that read permission data for identifiers other than the caller's own.

STRICTLY CONFIDENTIAL — This document contains sensitive security and technical information. Reproduction, distribution, or disclosure to unauthorized parties without explicit written consent is strictly prohibited.

FINDING 7

Staff Directory Disclosed to Any Caller

RISK LEVEL

Medium

AFFECTED SECURITY CONTROL

Data Protection

COMPLIANCE IMPACT

`OWASP A01:2025`

AFFECTED ENDPOINTS

1



5.5

RISK SCORE

Affected Hosts

`api.sample-client.example` `/api/users`

Vulnerability Description

The route `GET /api/users` returns the complete user list: 3,377 records in a response of 399,667 bytes, each with user identifier, first and last name, username and email address. No paging, filtering or scope applies.

The disclosure is itself a data-protection issue, and it is also the enabling step for the identity findings: it supplies the identifiers that make the permission-flags route (F06) and the header-spoofing reconnaissance (F05) walkable. One authenticated caller can enumerate the entire workforce.

Remediation & Solution Strategy

N30.ai recommends the following corrective actions in order of priority:

1. Restrict the directory to the roles that need it

Return the full user list only to administrative callers; other roles should see the users they work with, not the whole organisation.

2. Page and filter the endpoint

Require a filter or paging so the directory cannot be pulled in a single request, and omit email addresses and usernames where they are not needed.

3. Alert on full-directory reads

Treat a request that returns the whole list as unusual and record it.

STRICTLY CONFIDENTIAL — This document contains sensitive security and technical information. Reproduction, distribution, or disclosure to unauthorized parties without explicit written consent is strictly prohibited.

FINDING 8

Interactive API Documentation Exposed Without Authentication

RISK LEVEL

Medium

AFFECTED SECURITY CONTROL

Security Configuration

COMPLIANCE IMPACT

`OWASP A05:2025`

AFFECTED ENDPOINTS

1



5.5



Affected Hosts

`api.sample-client.example` `/help`

Vulnerability Description

The route `GET /help` returns the ASP.NET Web API help page, 390,866 bytes describing approximately 110 controllers, every route and verb (including all delete routes) and sample request bodies. It requires no token: a bad or absent token still returns the page.

The page is a complete map of the API surface, including administrative and setup routes that the application interface never shows. It shortens the work of an attacker considerably: every endpoint named in this report was reached from this map, and any endpoint not corrected can be found the same way.

Remediation & Solution Strategy

N30.ai recommends the following corrective actions in order of priority:

1. Require authentication for the help page

Gate `/help` behind the same authentication as the API, or disable it in the production and staging builds.

2. Restrict it to a trusted network

If the help page is needed by developers, expose it only on an internal address or behind an allow-list, never on the public API host.

3. Review the exposed surface

Use the published route list to confirm that every controller, including the setup and delete routes, is one the application intends to serve.

STRICTLY CONFIDENTIAL — This document contains sensitive security and technical information. Reproduction, distribution, or disclosure to unauthorized parties without explicit written consent is strictly prohibited.

FINDING 9

Cross-Origin Policy Set to Wildcard

RISK LEVEL

Medium

AFFECTED SECURITY CONTROL

Security Configuration

COMPLIANCE IMPACT

`OWASP A05:2025`

AFFECTED ENDPOINTS

1



4.8

RISK SCORE

Affected Hosts

`api.sample-client.example` `/api/states`

Vulnerability Description

Preflight and actual responses return `Access-Control-Allow-Origin: *` together with `Access-Control-Allow-Headers: authorization,market,userid`, for the application's own origin and for a foreign origin alike. `Access-Control-Allow-Credentials` is not set.

The boundary was measured before rating it. The API accepts only the bearer token: it sets no cookie and rejects the application session cookie on the API path. Because credentials are not allowed, a foreign page cannot ride a victim's session automatically; it must already hold the token. This is therefore a hardening gap rather than a live cross-site bypass, and it becomes high severity the moment cookie-based authentication is added to the API.

Remediation & Solution Strategy

N30.ai recommends the following corrective actions in order of priority:

1. Return a fixed list of allowed origins

Answer preflight requests only for the application's own origin, and echo an allowed origin rather than a wildcard.

2. Keep credentials off the wildcard

If the policy must stay permissive, ensure `Allow-Credentials` is never set alongside a wildcard origin, so a browser cannot attach a session.

3. Recheck when authentication changes

Treat the API's authentication method as a prerequisite for this finding: adding cookie authentication turns it into a cross-site request forgery exposure.

STRICTLY CONFIDENTIAL — This document contains sensitive security and technical information. Reproduction, distribution, or disclosure to unauthorized parties without explicit written consent is strictly prohibited.

FINDING 10

Verbose Errors Disclose Internal Detail

RISK LEVEL

Medium

AFFECTED SECURITY CONTROL

Security Configuration

COMPLIANCE IMPACT

`OWASP A05:2025`

AFFECTED ENDPOINTS

3



5.0

RISK SCORE

Affected Hosts

`api.sample-client.example`	`/api/glAccounts/search`	`/api/balancesheetschedules/{id}/summary`

Vulnerability Description

A malformed numeric bound returns the backend exception verbatim. `POST /api/glAccounts/search?pageSize=-1` answered 500 with a gRPC error naming an internal peer address and port, a source file path and a gRPC status code, together with the internal detail message.

Malformed identifiers leak the controller's method signature and assembly name (`SampleClientApi.Controllers.BalanceSheetScheduleController`). The route `/api/states/:x` returns the ASP.NET diagnostics page with a stack trace and version, because custom errors are disabled.

None of this grants access directly, but it hands an attacker the internal architecture: a private service address and port, the technology stack, and the names and locations of server-side classes. It shortens the work of every other finding in this report.

Remediation & Solution Strategy

N30.ai recommends the following corrective actions in order of priority:

1. Return a generic error to the caller

Answer with a correlation identifier and a plain message, and keep the exception detail, gRPC status and file paths in the server-side log only.

2. Enable custom errors

Turn custom errors on for every environment reachable from the internet, so the ASP.NET diagnostics page is never returned.

3. Do not echo framework detail

Suppress the model-binding detail that names controller methods and assemblies in error responses.

STRICTLY CONFIDENTIAL — This document contains sensitive security and technical information. Reproduction, distribution, or disclosure to unauthorized parties without explicit written consent is strictly prohibited.

FINDING 11

Stack and Version Disclosure on Every Response

RISK LEVEL

Low

AFFECTED SECURITY CONTROL

Security Configuration

COMPLIANCE IMPACT

`OWASP A05:2025`

AFFECTED ENDPOINTS

1



3.0

RISK SCORE

Affected Hosts

`api.sample-client.example`

Vulnerability Description

API responses carry `Server: Microsoft-IIS/10.0`, `X-AspNet-Version: 4.0.30319`, `X-Powered-By: ASP.NET` and `X-AspNetMvc-Version: 5.2` on every call.

The disclosure identifies the exact platform and framework revisions, which lets an attacker match the target against known flaws for that stack without probing. The request-tracing module is also present but correctly blocked: `Trace.axd` answers 403, so its presence confirms tracing is configured without exposing it.

Remediation & Solution Strategy

N30.ai recommends the following corrective actions in order of priority:

1. Remove the identifying headers

Strip `X-Powered-By`, `X-AspNet-Version` and `X-AspNetMvc-Version` at the web-server or application level, and reduce the `Server` header to a generic value.

2. Keep the stack current

Treat the disclosed versions as a prompt to confirm the platform and framework are on supported releases.

STRICTLY CONFIDENTIAL — This document contains sensitive security and technical information. Reproduction, distribution, or disclosure to unauthorized parties without explicit written consent is strictly prohibited.

FINDING 12

Unhandled Exception on a Reachable Parameter

RISK LEVEL

Low

AFFECTED SECURITY CONTROL

Insecure Design

COMPLIANCE IMPACT

`OWASP A04:2025`

AFFECTED ENDPOINTS

2



3.2



Affected Hosts

`api.sample-client.example`

`/api/vendors/import/template`

`/api/glAccounts/search`

Vulnerability Description

The route `GET /api/vendors/import/template?templateName=default` answered HTTP 500 for any value of the parameter: a valid route that throws regardless of input. The route `POST /api/glAccounts/search?pageSize=-1` also returned 500 on a negative bound.

Both are unhandled-exception paths on parameters a normal user can reach. Beyond the availability of the individual call, they are where the internal detail in F10 is emitted, so they matter for both robustness and disclosure.

Remediation & Solution Strategy

N30.ai recommends the following corrective actions in order of priority:

1. Validate before processing

Reject the parameter at the boundary (a positive range for the page size, a check for the template name) and return a typed 400 rather than reaching unhandled server code.

2. Handle the failure inside the endpoint

Where the underlying call can fail, catch it and return a controlled response instead of letting the exception surface as a 500.

3. Test the reachable routes for empty and invalid input

Add a check that every route answers a typed error, never a 500, for out-of-range and malformed input.

STRICTLY CONFIDENTIAL — This document contains sensitive security and technical information. Reproduction, distribution, or disclosure to unauthorized parties without explicit written consent is strictly prohibited.

FINDING 13

Unbounded Result Size and Negative Values Crash the Service

RISK LEVEL

Low

AFFECTED SECURITY CONTROL

Insecure Design

COMPLIANCE IMPACT

`OWASP A04:2025`

AFFECTED ENDPOINTS

1



3.5

RISK SCORE

Affected Hosts

`api.sample-client.example` `/api/glAccounts/search`

Vulnerability Description

The paging parameters of `POST /api/glAccounts/search` are not bounded. A page size of 2147483647 returned the whole result set in one response of 823,090 bytes, with no cap; a page size of 0 returned a valid empty page; `pageSize=-1` and `pageNum=-5` each returned 500.

An uncapped page size is a resource-consumption and enumeration aid: it lets one request pull every record at once, which makes bulk extraction of any list fast, and it lets a caller impose a large response on the server cheaply. The negative-value crash is the same unhandled-input path as F12.

Remediation & Solution Strategy

N30.ai recommends the following corrective actions in order of priority:

1. Cap and validate the page size

Accept only a page size within a defined maximum and a page number of one or greater, and return a typed 400 for anything outside those bounds.

2. Apply the cap server-side

Clamp the value in the service layer, not only in the client, so the limit holds however the endpoint is called.

3. Rate-limit bulk reads

Treat repeated maximum-size pages as bulk extraction and alert on the pattern.

STRICTLY CONFIDENTIAL — This document contains sensitive security and technical information. Reproduction, distribution, or disclosure to unauthorized parties without explicit written consent is strictly prohibited.

FINDING 14

Authorization Failures Are Not Centrally Logged

RISK LEVEL

Low

AFFECTED SECURITY CONTROL

Logging and Monitoring

COMPLIANCE IMPACT

`OWASP A09:2025`

AFFECTED ENDPOINTS

1



3.0

RISK SCORE

Affected Hosts

`api.sample-client.example` `/api/{module}/audittrails`

Vulnerability Description

Audit-trail endpoints exist for several modules and carry a date, user and message, which is good. What was not observed is any record of an authorization decision: a denied read or an authentication failure. The evidence is suggestive rather than conclusive, so this finding is rated LIKELY.

The integrity of what is logged is also in question. F01 shows the audit trail recording the true session user while the document record stores a forged author, so the two can disagree and the trail no longer proves who authored a record on its own.

Without records of denied reads, the access-control findings in this report are silent: the out-of-scope reads in F02, F03 and F04 would not appear in any log an operator would review.

Remediation & Solution Strategy

N30.ai recommends the following corrective actions in order of priority:

1. Log authorization decisions, not only changes

Record denied reads and authentication failures alongside create and update events, with the caller, the object requested and the decision.

2. Make the record authoritative

Since authorship is stored from the request (F01), correct that first; then the audit trail and the record will agree.

3. Alert on the patterns that matter

Alert on repeated denied reads and on out-of-scope lookups, which are the observable signature of the access-control findings.

STRICTLY CONFIDENTIAL — This document contains sensitive security and technical information. Reproduction, distribution, or disclosure to unauthorized parties without explicit written consent is strictly prohibited.

VULNERABILITIES **OVERVIEW**

Every finding in this report, in one table, grouped by severity. Exploitability states how readily each can be used against the application, and the identifier in the second column is the one used on the finding's own pages.

SEV	VULNERABILITY	DESCRIPTION	EXPLOITABILITY	AFFECTED CONTROL
H	**Resident Personal Data Readable by Identifier** F04	Resident names, addresses and tax identifiers are returned for any identifier, whatever the caller's scope.	**Trivial** A SIGNED-IN ACCOUNT	**Access Control** OWASP A01:2025
H	**Out-of-Scope Buildings Readable by Identifier** F02	The building detail route returns records tied to entities outside the caller's scope.	**Trivial** A SIGNED-IN ACCOUNT	**Access Control** OWASP A01:2025
H	**Entity and Super-Entity Records Readable Globally** F03	The entity and super-entity detail routes return records with no path-level scope check.	**Trivial** A SIGNED-IN ACCOUNT	**Access Control** OWASP A01:2025
H	**Identity Asserted by a Client-Supplied Header** F05	The permission menu follows a userid header the client sends; the token carries no user claim.	**Easy** A SIGNED-IN ACCOUNT	**Authentication** OWASP A07:2025
H	**Authorship Forgery on Document Upload** F01	A document's author is stored from the request body, so a record can be attributed to another user.	**Easy** A SIGNED-IN ACCOUNT	**Data Integrity** OWASP A08:2025
M	**Staff Directory Disclosed to Any Caller** F07	The full staff list, with names, usernames and work emails, is returned to any signed-in caller.	**Trivial** A SIGNED-IN ACCOUNT	**Data Protection** OWASP A01:2025
M	**Interactive API Documentation Exposed** F08	The API help page lists every route and verb with no authentication required.	**Trivial** NO ACCOUNT NEEDED	**Security Configuration** OWASP A05:2025
M	**Per-Control Permission Flags Readable for Any User** F06	Any caller can read any other user's control-level permissions by naming them in the path.	**Easy** A SIGNED-IN ACCOUNT	**Access Control** OWASP A01:2025
M	**Verbose Errors Disclose Internal Detail** F10	Errors return an internal gRPC address, file paths and controller signatures.	**Easy** A MALFORMED REQUEST	**Security Configuration** OWASP A05:2025
M	**Cross-Origin Policy Set to Wildcard** F09	The API answers cross-origin requests with a wildcard origin together with the authorization header.	**Moderate** DEPENDS ON THE BROWSER	**Security Configuration** OWASP A05:2025
L	**Unbounded Result Size on a Search Endpoint** F13	The page size is not capped, so one request returns the whole result set.	**Limited** A SIGNED-IN ACCOUNT	**Insecure Design** OWASP A04:2025
L	**Unhandled Exception on a Reachable Parameter** F12	At least one valid route throws a server error for any input value.	**Limited** A SIGNED-IN ACCOUNT	**Insecure Design** OWASP A04:2025
L	**Stack and Version Disclosure on Every Response** F11	Server and framework versions are advertised in the response headers.	**Limited** A SINGLE REQUEST	**Security Configuration** OWASP A05:2025
L	**Authorization Failures Are Not Centrally Logged** F14	No records of denied reads or authentication failures were observed; rated LIKELY.	**Context** SETS THE VISIBILITY	**Logging and Monitoring** OWASP A09:2025

STRICTLY CONFIDENTIAL — This document contains sensitive security and technical information. Reproduction, distribution, or disclosure to unauthorized parties without explicit written consent is strictly prohibited.

EXPLOITABILITY **OVERVIEW**

How each finding would actually be used: the vector an attacker needs, what has to be true first, and what N30.ai executed while confirming it. Findings marked Confirmed were carried out non-destructively, and the one write artefact was removed.

LVL	VULNERABILITY	ATTACK VECTOR	EXPLOITATION REQUIREMENTS & DETAILS	EXECUTION STATUS
H	**Resident Personal Data Readable by Identifier** F04	**Network** A SIGNED-IN REQUEST	A single request naming a resident identifier or a surname returns a profile. No prior knowledge or elevated role is needed.	**Confirmed** READ PERFORMED WITH A LIVE TOKEN
H	**Out-of-Scope Buildings Readable by Identifier** F02	**Network** A SIGNED-IN REQUEST	Changing the identifier in the address returns a building outside the caller's scope. 240 of 259 identifiers answered.	**Confirmed** SWEEP PERFORMED, READ ONLY
H	**Entity and Super-Entity Records Readable Globally** F03	**Network** A SIGNED-IN REQUEST	The detail routes answer for identifiers outside the scope selector. 113 of 119 super-entity identifiers answered.	**Confirmed** SWEEP PERFORMED, READ ONLY
H	**Identity Asserted by a Client-Supplied Header** F05	**Network** A SIGNED-IN REQUEST	One token, a changed userid header, a different permission map. The boundary was measured: on 16 data endpoints the header changed nothing.	**Confirmed** MENU SPOOFING DEMONSTRATED
H	**Authorship Forgery on Document Upload** F01	**Network** A SIGNED-IN REQUEST	An upload carrying createdBy: 1745 was stored with that author, against the authenticated user 3318. The artefact was removed after the test.	**Confirmed** WRITE PERFORMED, THEN UNDONE
M	**Staff Directory Disclosed to Any Caller** F07	**Network** A SIGNED-IN REQUEST	One request returned 3,377 staff records with names, usernames and work emails, unpaged.	**Confirmed** DIRECTORY PULLED ONCE
M	**Interactive API Documentation Exposed** F08	**Network** AN ANONYMOUS REQUEST	The help page needs no token at all, so a bad or absent token still returns the full route map.	**Confirmed** FETCHED WITHOUT A TOKEN
M	**Per-Control Permission Flags Readable for Any User** F06	**Network** A SIGNED-IN REQUEST	The path user's control flags are returned whatever the caller's identity; several users were read to confirm.	**Confirmed** READ FOR SEVERAL USERS
M	**Verbose Errors Disclose Internal Detail** F10	**Network** A MALFORMED REQUEST	Any malformed bound or identifier returns the detail. It grants nothing directly; it shortens the work of the other findings.	**Confirmed** ERRORS REPRODUCED ON THE API
M	**Cross-Origin Policy Set to Wildcard** F09	**Network** A FOREIGN ORIGIN	Preflight from a foreign origin returns a wildcard origin. Credentials are not allowed and the API is bearer-only, so a foreign page must already hold the token.	**Confirmed** PREFLIGHT AND ACTUAL RESPONSES
L	**Unbounded Result Size on a Search Endpoint** F13	**Network** A SIGNED-IN REQUEST	A maximum page size returned the whole 823 KB result set in one response; negatives crash the call.	**Confirmed** BOUNDS TESTED, READ ONLY
L	**Unhandled Exception on a Reachable Parameter** F12	**Network** A SIGNED-IN REQUEST	A valid route throws for any parameter value and returns a server error.	**Confirmed** TWO ROUTES REPRODUCED
L	**Stack and Version Disclosure on Every Response** F11	**Network** A SINGLE REQUEST	The headers name the web server and the framework versions on every call. No flaw was proven against them.	**Confirmed** HEADERS OBSERVED
L	**Authorization Failures Are Not Centrally Logged** F14	**Local** AN OPERATOR'S REVIEW	No records of denied reads were observed; a per-module trail may exist for modules not sampled. This finding is rated LIKELY.	**Not attempted** LOG REVIEW, NOT PROVEN

STRICTLY CONFIDENTIAL — This document contains sensitive security and technical information. Reproduction, distribution, or disclosure to unauthorized parties without explicit written consent is strictly prohibited.

CORPORATE PROFILE

N30.ai

N30.ai is the cybersecurity consulting practice of Manteiga Foundation, a Florida public benefit company at the meeting of technology and art. We answer the four questions boards ask: could someone get in today (penetration testing of applications, APIs and networks, ranked by what to fix first, retest included); is our cloud set up right (AWS, Azure and Google Cloud identity, posture and data exposure, returned as a remediation plan); will we pass the audit (controls mapped, gaps found, evidence ready before the auditor asks); and are we ready for quantum (the post-quantum transition, host by host). Trusted where security is regulated: banking and financial services, the public sector and healthcare. The instruments behind our engagements are open source, under the MIT license.

OFFICE

N30.ai

2222 Ponce de Leon, 7th floor
Coral Gables, Florida

CORE FOCUS	GENERAL INQUIRIES	EMAIL	WEBSITE
Cybersecurity consulting	**+1 305 709 2909**	**hello@n30.ai**	**n30.ai**

STRICTLY CONFIDENTIAL — This document contains sensitive security and technical information. Reproduction, distribution, or disclosure to unauthorized parties without explicit written consent is strictly prohibited.