AllSecOps

OFFENSIVE SECURITY. CLEAR DIRECTION.

PUBLIC SAMPLE

Proof. Then progress.

Penetration testing assessment report

ASTER LABS

Web & API / Cloud / Active Directory / AI / Network

DOCUMENT

ASO-SAMPLE-2026-01

EDITION

28 October 2026 / v1.0

Illustrative only. Aster Labs, its systems, findings and evidence are fictional. No security assessment was performed. This sample demonstrates the format and depth of an AllSecOps deliverable.

info@allsecops.com

01 / LEADERSHIP BRIEF

Trust boundaries need work.

In this fictional assessment, five weaknesses span application, identity and infrastructure controls.

0	3	2	0
CRITICAL	HIGH	MEDIUM	LOW / INFO

What this means for the business

The fictional environment separates users at the sign-in screen but does not consistently preserve that separation behind it. A customer account can read another tenant's report. A compromised build identity can read material outside its intended scope. Internal reachability compounds excessive directory permissions.

The decision to make now

Assign one accountable owner to each trust boundary and prioritize server-side authorization, cloud privilege reduction and directory delegation cleanup. These changes address the causes of exposure; additional perimeter controls alone would leave the underlying paths intact.

> **Recommended first milestone: 7 days**
>
> Contain the three High findings, agree on regression tests and schedule independent verification. All five findings remain Open in this sample; no remediation or retest is claimed.

READ THIS REPORT BY ROLE

Leadership: pages 2, 5 and 11. Engineering: pages 3-10 and 12.
Procurement and assurance: pages 3, 4 and 13.

02 / SCOPE & AUTHORIZATION

The boundaries of the test.

A fictional, ten-business-day, gray-box engagement using seeded data and explicitly approved identities.

SURFACE	ILLUSTRATIVE TARGET	ACCESS MODEL
Web & API	portal.aster.example api.aster.example	2 users / 2 tenants
Cloud	aster-build / archive fictional AWS account	Read-only config + seeded canary
Active Directory	corp.aster.example LAB-APP-01	Standard domain user
AI application	assist.aster.example seeded retrieval corpus	User + test document
Internal network	192.0.2.0/24 198.51.100.0/24	Approved lab probe

Rules of engagement

Written authorization, target ownership and emergency contacts would be confirmed before a real engagement. In this scenario, testing is limited to lab accounts, benign markers and non-destructive validation. Stop conditions include service instability or contact with unapproved data.

Explicit exclusions

Physical penetration testing, social engineering, denial-of-service, persistence, production customer data and unapproved third-party services are excluded. Wireless, mobile, source code and Kubernetes testing are separate scopes, not covered by this sample.

Coverage is time-bound and sampled. The reserved .example names and documentation IP ranges above are illustrative and must not be treated as live targets.

03 / APPROACH & RISK MODEL

Evidence before severity.

A manual, hypothesis-led workflow supported by configuration review and narrowly scoped validation.

01

Map

Understand assets, identities, data flows and agreed threat scenarios.

02

Challenge

Test boundaries with paired identities, seeded objects and benign canaries.

03

Validate

Record the expected control, observed result and minimum proof of impact.

04

Prioritize

Tie likelihood and impact to an owner, corrective action and retest criterion.

QUALITATIVE SEVERITY / NO CVSS SCORE ASSIGNED

Critical	Readily achievable, broad compromise of a core business boundary.
High	Credible path to sensitive access or elevated control; limited prerequisites.
Medium	Meaningful exposure with constraints or an additional prerequisite.
Low / Info	Limited direct impact, hardening opportunity or contextual observation.

How to interpret the ratings

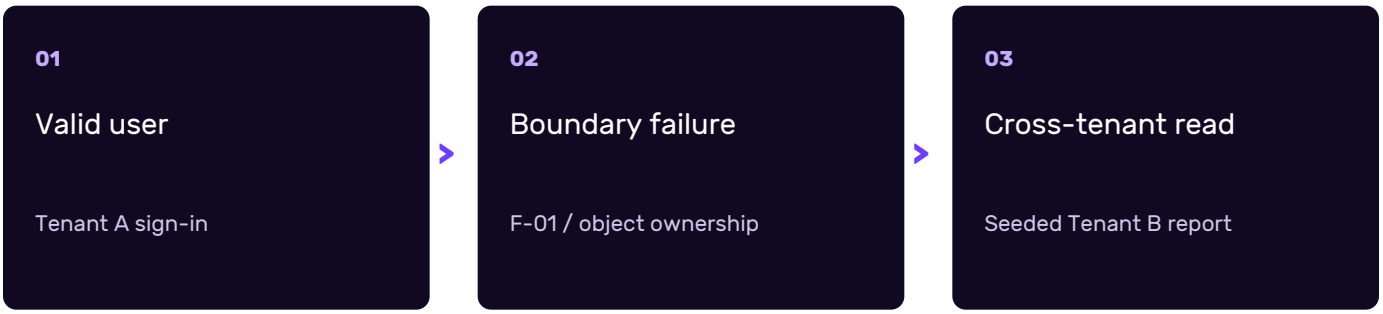
Ratings describe this fictional scenario. A real assessment would document business context, proof, compensating controls and any agreed scoring model.

04 / ATTACK NARRATIVE

Small gaps. Connected impact.

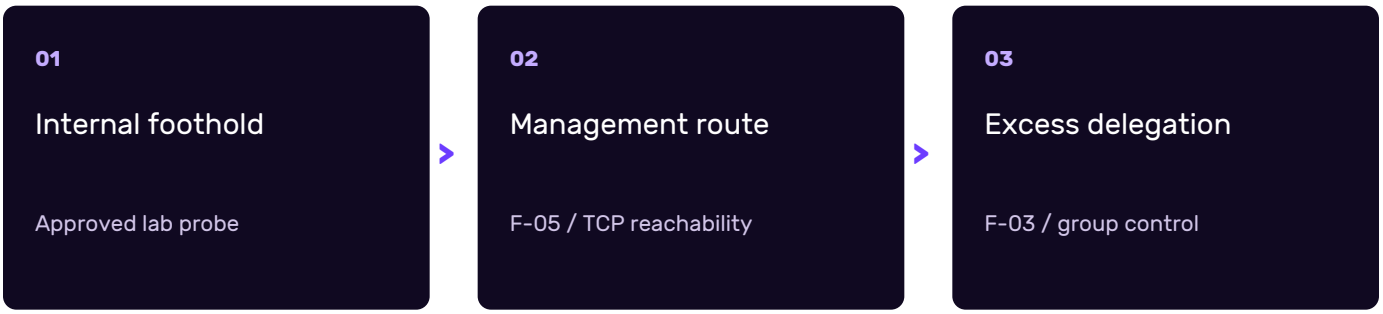
Two independent paths show why trust boundaries matter. No end-to-end domain compromise is claimed.

PATH A / CUSTOMER DATA



The sample API response includes a report belonging to Tenant B while the session belongs to Tenant A. Evidence supports one cross-tenant read; bulk extraction and write access were not tested.

PATH B / INTERNAL CONTROL



F-05 provides network reachability only. F-03 separately assumes a supplied domain account and shows control of a lab application-operator group. The diagram explains how these conditions could compound; it does not assert that reachability provides credentials.

Other independent exposures

F-02 tests cloud data access. F-04 tests an AI tool approval boundary. Neither is presented as a prerequisite for the paths above.

F-01 / WEB & API

Cross-tenant report access

HIGH

OPEN / NOT RETESTED

Target: api.aster.example / v1/reports
Owner: Application engineering

Control failure

The report lookup checks that a user is signed in but does not verify that the requested report belongs to the user's tenant. In this fictional case, a user in Tenant A receives a seeded report owned by Tenant B.

ILLUSTRATIVE EVIDENCE / NO LIVE CAPTURE

```
Fixture E-01 / synthetic HTTP observation
Session: tenant=A, role=member
Resource: seeded-report-B-017 / owner=tenant-B
Expected: 403 or non-disclosing 404
Observed: 200 / tenant-B marker returned
```

Business impact

A customer could read another customer's business reports if identifiers become known. The sample proves one seeded-object read; identifier enumeration, bulk access and data modification are untested.

Recommended correction

Apply authorization at the data access layer using the authenticated tenant context. Require both tenant and object ownership checks for reads, exports and shared-link flows. Do not rely on unguessable identifiers as the access control.

Verification criterion

Tenant A must receive a non-disclosing denial for Tenant B objects in detail, export and list routes. Confirm Tenant B still has access and add automated cross-tenant regression tests.

High impact / low prerequisite: an ordinary customer session can cross a tenant boundary.

F-02 / CLOUD SECURITY

Build identity exceeds its role

HIGH

OPEN / NOT RETESTED

Target: aster-build / fictional cloud archive
Owner: Cloud platform engineering

Control failure

The fictional CI identity has object-read permissions across an archive outside the build pipeline's purpose. A benign canary in the restricted archive is readable with this identity. No cloud account takeover is asserted.

ILLUSTRATIVE EVIDENCE / NO LIVE CAPTURE

```
Fixture E-02 / synthetic access decision
Principal: aster-build
Object: archive/restricted/canary.txt
Required business scope: build artifacts only
Observed: read allowed / CANARY-ASO-02
```

Business impact

Compromise of a build job could expose restricted archives beyond the software supply chain. This sample demonstrates one read using a supplied identity; it does not establish credential theft, privilege escalation or archive modification.

Recommended correction

Replace broad object permissions with named artifact paths and actions required by the build. Separate release and archive identities. Add a permissions boundary where appropriate and monitor denied and successful reads to sensitive storage.

Verification criterion

The build identity must be denied access to the restricted canary while normal artifact publishing succeeds. Verify effective permissions, including resource policies and inherited grants.

High impact / constrained prerequisite: the attacker must first control the build identity.

F-03 / ACTIVE DIRECTORY

Delegation crosses a privilege tier

HIGH

OPEN / NOT RETESTED

Target: corp.aster.example / Lab-App-Operators
Owner: Identity & Windows engineering

Control failure

An overly broad directory permission lets the fictional Helpdesk-L1 group modify membership of Lab-App-Operators. That destination group is mapped to local administration on LAB-APP-01, a seeded application server.

ILLUSTRATIVE EVIDENCE / NO LIVE CAPTURE

```
Fixture E-03 / synthetic lab transcript
Source identity: helpdesk.user / standard user
Delegated action: modify target group membership
Target: Lab-App-Operators -> LAB-APP-01 admin
Observed: seeded test user added; then removed
```

Business impact

A compromised help-desk account could obtain administrative control of the lab application tier. The scope does not support a claim of domain administrator access, credential dumping or compromise of other hosts.

Recommended correction

Remove unintended group membership delegation and review nested permissions along the same path. Separate support and administrative identities. Restrict server administration through role-specific groups and alert on sensitive membership changes.

Verification criterion

A standard help-desk identity must fail to alter the target group. Authorized operations must still succeed. Confirm nested group paths, endpoint role mappings and event capture after the change.

High impact / authenticated prerequisite: a standard help-desk user can acquire application administration.

F-04 / AI APPLICATION SECURITY

Retrieved content bypasses approval

MEDIUM

OPEN / NOT RETESTED

Target: assist.aster.example / document assistant
Owner: AI application engineering

Control failure

A retrieved document can influence the assistant to invoke a sandbox export tool without the user confirmation required by the application design. The synthetic test uses only a benign marker and an inert tool destination.

ILLUSTRATIVE EVIDENCE / NO LIVE CAPTURE

```
Fixture E-04 / synthetic agent trace
Input: test document with untrusted instructions
User intent: summarize the document
Expected: no export without explicit confirmation
Observed: export tool invoked / marker ASO-AI-04
```

Business impact

An attacker able to influence retrieved content could trigger unintended actions within the assistant's permissions. This example shows a sandbox action only; no production data loss, external transfer or generalized model compromise is established.

Recommended correction

Enforce tool authorization and explicit confirmation outside the language model. Constrain tools to approved destinations and task-scoped data. Treat retrieved text as untrusted and use allowlisted schemas with per-request permission checks.

Verification criterion

Replay the seeded document and variations. The application must block the export until a user approves the specific action. Confirm that normal summarization remains available and the decision is recorded.

Moderate impact / conditional likelihood: seeded content and an enabled tool workflow are required.

F-05 / NETWORK SECURITY

User segment reaches management

MEDIUM

OPEN / NOT RETESTED

Target: 192.0.2.0/24 -> 198.51.100.24
Owner: Network & infrastructure engineering

Control failure

The fictional user network can reach a management service on the application server subnet. A firewall exception permits the connection despite a documented management-only design. Authentication was not bypassed.

ILLUSTRATIVE EVIDENCE / NO LIVE CAPTURE

```
Fixture E-05 / synthetic connectivity result
Source: approved user-segment lab probe
Destination: 198.51.100.24 / management listener
Expected: blocked outside management segment
Observed: TCP handshake / authentication prompt
```

Business impact

An internal foothold has a larger management attack surface and may attempt credential misuse or exploit a future service weakness. The evidence supports reachability only, not successful logon or host compromise.

Recommended correction

Limit management traffic to designated administration hosts and approved identity-aware access paths. Remove broad exceptions, document justified dependencies and monitor policy changes. Keep administrative authentication controls in place.

Verification criterion

Repeat probes from user, application and administration segments. Unapproved paths must fail; the approved administration path must work. Verify both enforcement and corresponding network telemetry.

Moderate impact / additional control required: reachability exists, but authentication still applies.

05 / REMEDIATION ROADMAP

Turn evidence into ownership.

Suggested planning windows begin after report acceptance. These are sample targets, not agreed service levels.

01

WITHIN 7 DAYS

Contain the high-impact paths

Application, cloud and identity owners pair temporary restrictions with permanent corrective work.

F-01 tenant enforcement / F-02 scoped storage access / F-03 delegation cleanup

02

WITHIN 30 DAYS

Harden the control boundaries

AI and network teams remove unintended action paths and validate normal business workflows.

F-04 external tool authorization / F-05 management segmentation

03

WITHIN 60 DAYS

Make the fixes durable

Security engineering turns the retest criteria into regression coverage and review triggers.

Cross-tenant test suite / IAM review gates / delegation baselines / tool approvals

Success measure: Every finding has an owner, closure evidence and a verified test result.

06 / RETEST & HANDOVER

Closure requires proof.

A change ticket is a starting point. Verified behavior determines whether a finding can be closed.

FINDING	PRIORITY	CURRENT STATE	RETEST RESULT
F-01	**HIGH**	Open (illustrative)	Not performed
F-02	**HIGH**	Open (illustrative)	Not performed
F-03	**HIGH**	Open (illustrative)	Not performed
F-04	**MEDIUM**	Open (illustrative)	Not performed
F-05	**MEDIUM**	Open (illustrative)	Not performed

Evidence required for a real retest

Record the deployed build or configuration revision, remediation date, test identity, target, expected behavior and actual result. Repeat the original condition and adjacent variants; retain enough sanitized evidence for an independent reviewer to reproduce the conclusion.

Closure vocabulary

Verified resolved: original and adjacent checks pass.
Partially resolved: some exposure remains.
Not resolved: the original condition is reproducible.
Not retested: no verification was completed. Risk acceptance is a business decision, not a successful technical retest.

Handover package

Provide the final report, findings register, sanitized evidence references and retest addendum. Agree secure delivery and retention with the client before any real assessment data is collected.

07 / REFERENCES & REPORT NOTES

Built for decisions.

This original sample uses public reports as structural references, not as a source of client findings or copied evidence.

01 OffSec / MegaCorp One sample report

Useful reference for a leadership summary followed by an understandable attack narrative and actionable recommendations.

offsec.com / sample penetration testing report

02 Cure53 / Project 11 Web App UI, API & Infra

Useful reference for explicit work packages, stable finding identifiers, concise technical descriptions and transparent fix notes.

cure53.de / Project 11 report

03 Trail of Bits / DragonFly2 security assessment

Useful reference for coverage limitations, project goals and a separate fix-review section with clear result categories.

ostif.org / DragonFly2 assessment

Research and originality

References reviewed on 28 October 2026. They are examples of strong reporting practice, not an objective ranking or an endorsement of AllSecOps. All scenario text, diagrams, evidence fixtures and recommendations in this sample are original. No source logos or screenshots are reproduced.

Your next assessment should make the next step clear.

AllSecOps / info@allsecops.com