WhatsExposed

Acme

Continuous Exposure Assessment – Initial Assessment Report

Author: WhatsExposed Testing Team
Version: v1.0
Date 11/10/2026
Client Confidential

Private & Confidential

WhatsExposed conducted this assessment on the assets that existed as of 1 October 2026. Information security threats are continually changing, with new vulnerabilities discovered on a daily basis, and no application can ever be 100% secure no matter how much security testing is conducted.

This report cannot and does not protect against personal or business loss as a result of use of the applications or systems described. WhatsExposed offers no warranties, representations or legal certifications concerning the applications or systems it tests. All software includes defects: nothing in this document is intended to represent or warrant that security testing was complete and without error, nor does this document represent or warrant that the application tested is suitable to task, free of other defects than reported, fully compliant with any industry standards, or fully compatible with any operating system, hardware, or other application. By using this information, you agree that WhatsExposed shall be held harmless in any event.

Copyright @2026 WhatsExposed. All Rights Reserved.

This document is strictly private and confidential.

This report (including any enclosures and attachments) has been prepared for the exclusive use and benefit of the addressee(s) and solely for the purpose for which it is provided. Unless WhatsExposed provide express prior written consent, no part of this report should be reproduced, distributed, or communicated to any other company or third party.

The information included in this report has been prepared and included for the purpose of this document only and cannot be used as a precedent or in any other situation outside this report. We do not accept any liability if this report is used for an alternative purpose from which it is intended, nor to any third party in respect of this report.

WhatsExposed acknowledge all third-party trademarks in this document.

Table of Contents

Private & Confidential 2

Table of Contents 4

Executive Summary 5

Key Findings 6

Exposure Chains 9

Detailed Findings 11

1. Customer Portal Credentials in Stealer Logs 11
2. Insecure Transport 13
3. Registered Lookalike Domains 14
4. Publicly Reachable Pre-Production Environments 16
5. Staff Account Password Hash in Breach Data 18
6. Staff Mailboxes Indexed in Leak Sources 20
7. Catch-All DNS and Abandoned Cloud Endpoint 22
8. Mail Transport Hardening Not Published 24
9. Vulnerable FTP Service on Legacy Hosting Address 26
10. Microsoft 365 Tenant and Realm Discovery 28
11. Salesforce Aura API Exposed 30
12. Misconfigured Cookie Attributes 32
13. Content Security Policy (CSP) Header Misconfiguration 34
14. Customer Relationships Disclosed in CSP Header 36
15. Internet-Facing Sign-In Surfaces 38
16. Attack Surface Composition 40

Appendix A – Our Approach 42

Appendix B – Definitions 43

Appendix C – Asset Inventory 46

Appendix D – Supporting Scan Data 47

Executive Summary

WhatsExposed were commissioned by Acme to conduct a continuous exposure assessment against their internet-facing estate. The initial assessment began on 1st October 2026.

The purpose of the assessment was to identify the assets Acme exposes to the internet and the security misconfigurations, vulnerabilities and leaked information associated with them. This report contains the results found by WhatsExposed at the time of the assessment and is the baseline that subsequent reports measure change against.

Our approach simulates a threat actor's tactics, techniques and procedures to identify vulnerabilities associated with a system. Our proven assessment methodology follows industry and vendor best practices including, but not limited to, OWASP, OSSTMM and CREST benchmarks to assess and quantify risks associated with any patching deficiencies, security misconfigurations or insecure development practices.

Scope, Objectives and Limitations

Internet-facing estate:

- acme-example.com and all discovered subdomains
- 86 subdomains, 132 web applications and 130 addresses discovered during the assessment (Appendix C)
- 38 mailbox addresses and associated identity exposure

The assessment was conducted to meet the following objectives:

Passive Discovery	Identify the domains, hostnames, addresses, web applications and identities that make up the internet-facing estate from public and commercial data sources.
Uncredentialled Assessment	Assess the in-scope assets with limited or no knowledge, simulating the actions of a threat actor having to use publicly available information in an attempt to circumvent security controls and gain unauthorised access to resources.

Key Findings

During the Continuous Exposure Assessment, five Medium, three Low and eight Informational risk findings were identified, classifying the overall risk rating as Medium.



Please see Appendix B for the Risk Rating definitions.

Acme's own estate is small and mostly well kept. Sender authentication is at the strictest setting, the EU client portal sends a strict transport policy and a content security policy, the core applications sit behind managed AWS services with current certificates, and no confirmed software vulnerability was found on a system Acme is known to run. The exposure in this report comes from three other places: credentials exposed outside Acme, the trust placed in the Acme name, and things that were retired but never removed.

The assessment team found customer logins to the client portal in infostealer data: 4 captured logins with clear text passwords and a password reset link, covering 2 accounts at 2 organisations outside Acme. Separately, a breach of an unrelated website holds a password hash for one Acme staff mailbox.

The team also found an old shared hosting server still linked to the domain through historical DNS, running ProFTPD 1.3.5e with 12 published vulnerabilities matched by version.

Finally, 12 of 17 registered lookalike domains have mail configured, 5 pre-production copies of the client and web submission applications are open to the internet, and 42 hostnames fall through a catch-all DNS record to the website host.

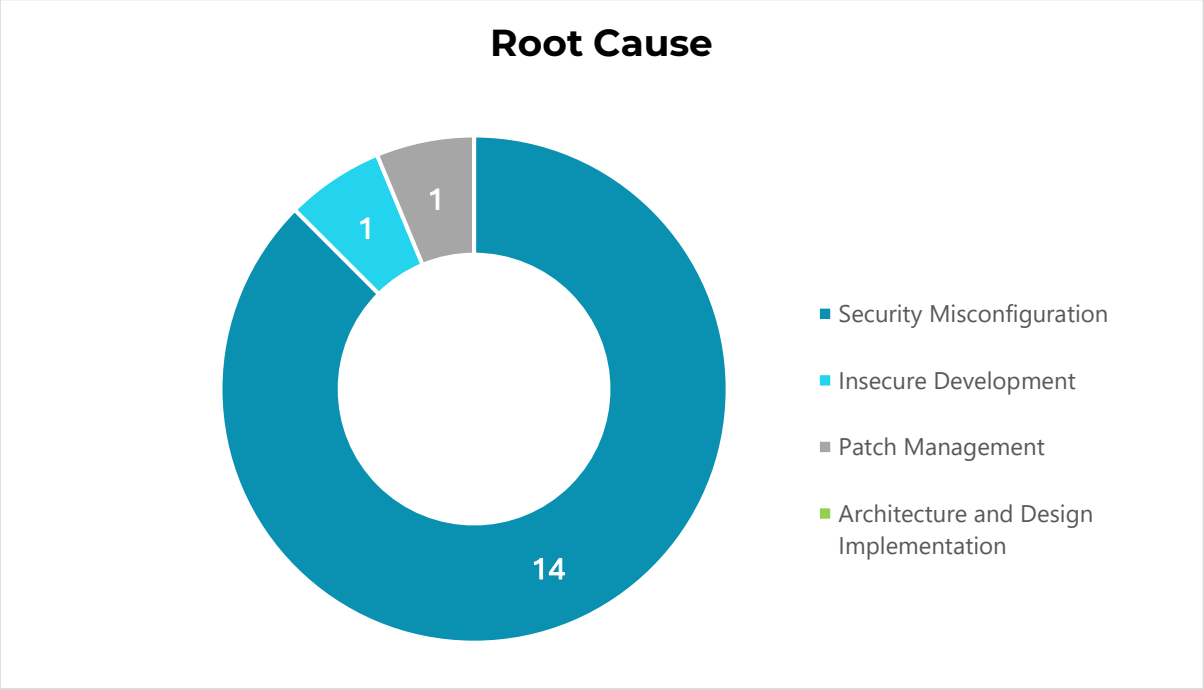
The table below summarises the findings identified. Further details associated with each of these findings is presented in the main body of the report and is referenced.

Ref	Finding	Risk Rating	Affected Assets
1	Customer Portal Credentials in Stealer Logs	**MEDIUM**	• client.acme-example.com • acme-example.com

Ref	Finding	Risk Rating	Affected Assets
2	Insecure Transport	**MEDIUM**	• 192.0.2.1 • 192.0.2.2 • web.client.acme-example.com • web.preprod-client.acme-example.com • web.us.client.acme-example.com • and 1 more
3	Registered Lookalike Domains	**MEDIUM**	• acme-example.com
4	Publicly Reachable Pre-Production Environments	**MEDIUM**	• web.preprod-client.acme-example.com • web.us.preprod-client.acme-example.com • submit.preprod-client.acme-example.com • submit.staging-client.acme-example.com • submit.us.preprod-client.acme-example.com • and 15 more
5	Staff Account Password Hash in Breach Data	**MEDIUM**	• acme-example.com
6	Staff Mailboxes Indexed in Leak Sources	**LOW**	• acme-example.com
7	Catch-All DNS and Abandoned Cloud Endpoint	**LOW**	• static-assets.client.acme-example.com
8	Mail Transport Hardening Not Published	**LOW**	• acme-example.com
9	Vulnerable FTP Service on Legacy Hosting Address	**INFO**	• 192.0.2.3 • 192.0.2.4
10	Microsoft 365 Tenant and Realm Discovery	**INFO**	• https://acme.auth0.com • https://acme.okta.com • https://acme.onelogin.com • https://login.microsoftonline.com • acme-example.com
11	Salesforce Aura API Exposed	**INFO**	• content.acme-example.com
12	Misconfigured Cookie Attributes	**INFO**	• blog.acme-example.com • client.acme-example.com • content.acme-example.com • app.us.client.acme-example.com • acme-example.com • and 4 more
13	Content Security Policy (CSP) Header Misconfiguration	**INFO**	• 192.0.2.5 • 192.0.2.6 • 192.0.2.7 • 192.0.2.8 • 192.0.2.9 • and 49 more
14	Customer Relationships Disclosed in CSP Header	**INFO**	• submit.client.acme-example.com • submit.us.client.acme-example.com • submit.staging-client.acme-example.com • submit.preprod-client.acme-example.com • submit.us.preprod-client.acme-example.com

Ref	Finding	Risk Rating	Affected Assets
15	Internet-Facing Sign-In Surfaces	INFO	• client.acme-example.com • us.client.acme-example.com • app.us.client.acme-example.com • web.client.acme-example.com • api-auth.acme-example.com • and 1 more
16	Attack Surface Composition	INFO	• 192.0.2.10 • 192.0.2.11 • 192.0.2.12 • 192.0.2.13 • 192.0.2.14 • and 8 more

The findings identified can be categorised into the following root causes:



Please see Appendix B for the Root Cause definitions.

Regulatory Context

Acme processes personal data on behalf of its business customers. The web submission policy (Finding 14) shows those customers include regulated financial services firms. Under Article 33(2) of the GDPR, a processor must notify the controller without undue delay after becoming aware of a personal data breach. That is why the log review in Finding 1 matters: it establishes whether there is anything to notify. Customers in scope of the Digital Operational Resilience Act (Regulation (EU) 2022/2554) will also treat Acme as an ICT third-party service provider and can be expected to ask about the controls in this report, MFA on the portal and transport security in particular.

This report does not constitute a compliance assessment and should not be read as one. The regulatory points are noted where a technical finding maps onto a specific obligation.

Exposure Chains

The findings in this report do not act independently. The analysis below sets out 4 paths where one finding supplies what another needs. Each chain is rated by its outcome and not by the rating of its individual components.

Each chain records a breaker: the single change that costs least and closes the path.

Ref	Exposure Chain	Findings Involved
A	Lookalike mail to staff credential capture to tenant access	• Registered Lookalike Domains • Staff Mailboxes Indexed in Leak Sources • Microsoft 365 Tenant and Realm Discovery • Internet-Facing Sign-In Surfaces
B	Pre-production copy as the easier way in to the same application	• Publicly Reachable Pre-Production Environments • Customer Portal Credentials in Stealer Logs
C	Unclaimed hostname under acme-example.com trusted by the submission flow	• Catch-All DNS and Abandoned Cloud Endpoint • Customer Relationships Disclosed in CSP Header • Content Security Policy (CSP) Header Misconfiguration
D	Breached staff password reused against the Microsoft 365 tenant	• Staff Account Password Hash in Breach Data • Microsoft 365 Tenant and Realm Discovery • Internet-Facing Sign-In Surfaces

Chain A: Lookalike mail to staff credential capture to tenant access

Path

DMARC at reject means mail that pretends to come from acme-example.com is refused, so an attacker uses one of the 12 lookalike domains already set up for mail. The 38 indexed mailbox addresses supply the recipients, and the two naming conventions supply the rest of the staff. The tenant discovery confirms Microsoft 365 is the login page to copy and lets each address be checked first. A staff member enters their password on the copy.

Outcome

A staff mailbox and whatever that account can reach in Microsoft 365. From a staff mailbox an attacker can also send convincing mail to customers from the real domain.

Breaker

Phishing-resistant MFA through conditional access for all staff removes the value of a captured password.

Chain B: Pre-production copy as the easier way in to the same application

Path

The pre-production and staging copies of the client and web submission applications are reachable by anyone. The leak index already lists a staging hostname. An attacker who wants the production application studies and attacks the copy first, where changes land earlier, monitoring is lighter and test accounts tend to have weak or shared passwords. What the copies hold was not tested.

Outcome

Knowledge of the application's next release and its weaknesses, possible access to test data that may include real documents, and credentials or keys that are sometimes shared with production.

Breaker

Restricting pre-production to known source addresses closes this for every environment at once.

Chain C: Unclaimed hostname under acme-example.com trusted by the submission flow

Path

The catch-all record means hostnames nobody has set up still resolve to a third-party platform, and one CloudFront endpoint appears to be abandoned. If either can be claimed by someone else, that person serves their own page under a real acme-example.com name with a valid certificate. The web submission application allows any acme-example.com subdomain to embed it, and has no script policy, so the page can frame the genuine submission flow and draw over it. Whether a takeover is possible was not tested.

Outcome

A convincing attack page on Acme's own domain that wraps the real application.

Breaker

Removing the wildcard record and the abandoned endpoint removes the foothold.

Chain D: Breached staff password reused against the Microsoft 365 tenant

Path

An attacker takes the bcrypt hash for the staff mailbox from the breach data and tries to recover the password. The Microsoft 365 tenant for acme-example.com is discoverable, so the attacker knows where to try it. If the password matches the staff member's Acme password and MFA is not enforced on that account, the sign-in succeeds. Both conditions are untested, and bcrypt makes the first step slow unless the password is weak.

Outcome

Access to one staff mailbox and the files behind it, and a trusted acme-example.com sender for further phishing.

Breaker

Reset the password on the affected account and confirm MFA is enforced on it.

Detailed Findings

1. Customer Portal Credentials in Stealer Logs

MEDIUM	Risk	Asset	https://client.acme-example.com
MEDIUM	Impact	Root Cause	Security Misconfiguration
MEDIUM	Likelihood	Affected Services	HTTPS, client portal sign-in
8.2	CVSS:4.0 Score	CVSS:4.0 Vector	AV:N/AC:L/PR:N/UI:N/S:U/C:H/I:L/A:N

Description

Infostealer malware captures the usernames, passwords and session cookies saved in a victim's browser and sells them on criminal marketplaces. Dark web monitoring returned 4 captured logins to the Acme client portal, covering 2 accounts at 2 organisations outside Acme. All passwords were captured in clear text, and two of the records carry the same password reset link.

All four logins are for client.acme-example.com. Each account appears in two separate log records, which suggests the same machine was captured more than once or the logs were resold.

None of these are Acme staff accounts, and nothing was leaked by Acme. The logins were taken by malware on the account holders' own computers. The risk still lands on Acme, because the door these credentials open is the Acme portal.

Impact

The client portal is where business customers work with their end users' submissions. Anyone signing in with one of these accounts would likely see the personal data of that customer's end users. A sign-in with a stolen password looks the same as a normal one, so it would not stand out in the logs.

Stealer logs usually include session cookies too. Where they do, a second factor does not stop access until the session is revoked.

As a processor, Acme would need to tell the affected customer without undue delay if any of these accounts has been misused.

Likelihood

The credentials are already in criminal hands and need no technical skill to use. All the passwords follow a simple word-plus-numbers pattern, which shows the portal accepts weak passwords. No credential was tested, so it is not known whether the passwords still work or whether these accounts require a second factor. The finding is rated on the basis that at least one is still live.

Recommendation

It is recommended that Acme take the following steps, in this order.

- Identify the two accounts from the portal user list. Force a password reset on each and revoke all active sessions and API tokens.
- Invalidate outstanding password reset links and confirm reset links are single use and expire quickly.
- Review portal logs for these accounts for new devices, new countries, and bulk viewing or export of submissions.

Strategic:

- Make multi-factor authentication mandatory for every portal user.
- Alert on sign-ins from a new device.
- Block passwords that appear in known breach lists or follow simple patterns.

Technical Details

Bucket
leaks.logs
Host
client.acme-example.com
Source
https://client.acme-example.com/password/reset/[token withheld]
System ids
[log ref 1], [log ref 2], [log ref 3], [log ref 4]

MATCHED LINES

[log ref 1] · lines 1742-1746

```
URL: https://client.acme-example.com/password/reset/[token withheld]
Username: s******@[customer A domain withheld]
Password: [withheld]
Application: Microsoft_[Edge]_Default
===============
```

[log ref 2] · lines 1741-1745

```
URL: https://client.acme-example.com/password/reset/[token withheld]
Username: s******@[customer A domain withheld]
Password: [withheld]
Application: Microsoft_[Edge]_Default
===============
```

[log ref 3] · lines 10-14

```
URL: https://client.acme-example.com/
Username: s******@[customer B domain withheld]
Password: [withheld]
Application: Google_[Chrome]_Default
===============
```

[log ref 4] · lines 10-14

```
URL: https://client.acme-example.com/
Username: s******@[customer B domain withheld]
Password: [withheld]
Application: Google_[Chrome]_Default
===============
```

2. Insecure Transport

MEDIUM	Risk	Asset	6 assets, listed in Technical Details
MEDIUM	Impact	Root Cause	Security Misconfiguration
MEDIUM	Likelihood	Affected Services	HTTPS (TCP 443)
6.3	CVSS:4.0 Score	CVSS:4.0 Vector	AV:N/AC:H/AT:P/PR:N/UI:N/VC:L/VI:L/VA:N/SC:N/SI:N/SA:N

Description

Transport Layer Security protects data in transit between a user and the application. Versions 1.0 and 1.1 of the protocol have been withdrawn, and the 3DES cipher is considered weak.

web.client.acme-example.com, web.us.client.acme-example.com, web.preprod-client.acme-example.com and web.us.preprod-client.acme-example.com all negotiate TLS 1.0 and TLS 1.1 and accept the DES-CBC3-SHA cipher. These are the client web application in the EU and the US and its two pre-production copies, all served through Amazon CloudFront in front of S3. The client portal and the web submission application accept TLS 1.2 and 1.3 only, so this is a per-distribution setting that was left at an older value.

The same weakness was reported on 192.0.2.1 and 192.0.2.2, two HubSpot addresses that serve blog.acme-example.com and thank-you.acme-example.com. Requested by name, both hostnames negotiated TLS 1.2 and 1.3 only. The older protocols are the platform's default on its shared addresses and are not a setting Acme controls.

Impact

TLS 1.0 and 1.1 are withdrawn protocols and 3DES is open to the Sweet32 attack on long-lived connections. An attacker positioned on the network path could downgrade a connection and attempt to recover session data. For a company that sells into regulated firms, the more immediate effect is that the setting is flagged in customer security questionnaires and external scans.

Likelihood

An attacker needs a position on the network path and a client willing to negotiate down, which limits interception in practice. The weakness is visible to any external scan, so it is likely to be raised by a customer or auditor.

Recommendation

It is recommended that Acme set the CloudFront security policy on these four distributions to TLSv1.2_2021 or later, to match the distributions that are already correct. No action is needed on the HubSpot addresses.

Strategic Considerations: record the required policy in the infrastructure templates so a new distribution cannot be created with the older default.

Technical Details

Technical evidence withheld from this sample. In the client report this section lists, for each host, the deprecated protocol versions and weak cipher suites accepted, with the commands to reproduce the check.

3. Registered Lookalike Domains

MEDIUM	Risk	Asset	acme-example.com
MEDIUM	Impact	Root Cause	Security Misconfiguration
MEDIUM	Likelihood	Affected Services	DNS, SMTP
5.1	CVSS:4.0 Score	CVSS:4.0 Vector	AV:N/AC:L/AT:N/PR:N/UI:A/VC:N/VI:N/VA:N/SC:L/SI:L/SA:N

Description

Lookalike domains differ from a genuine domain by one character or by top-level domain, and can be used to impersonate an organisation in email and on the web. Registration alone does not show intent. A lookalike may belong to an unrelated business with a similar name, or be held defensively by the organisation it resembles. Domain permutation monitoring raised 17 lookalikes of acme-example.com, and the assessment identified 16 of them as registered. The seventeenth, acme-example.ph, is not registered: it resolves only because the .ph registry answers DNS queries for every unregistered name. Twelve of the 16 publish mail exchanger (MX) records, which tell other mail servers where to deliver email for a domain, on providers including Google, Microsoft, Rackspace and GoDaddy.

Registration records show that five of the 16 were registered between 1998 and 2012, and three since May 2025, none of those three with mail configured. acme-example.org, acme-example.info, acme-example.me and acme-example.us share name servers, a mail provider and DMARC settings, which indicates a single holder. DMARC is the DNS record that tells receiving mail servers what to do with mail that fails sender checks. The holder of these four could not be identified from public records.

Impact

A lookalike with mail configured lets its holder send and receive email under a name resembling Acme's. The holder sets that domain's SPF record, the list of servers allowed to send its mail, and its DMARC record, so its messages can pass the sender checks receiving servers apply. Acme's own DMARC policy is at reject, which blocks direct spoofing of acme-example.com and leaves a lookalike as the route for anyone impersonating Acme.

Staff, customers and the people asked to submit personal documents could all be approached this way. A request for an personal document that appears to come from Acme is one such a recipient may expect, and anything sent in reply would reach the lookalike's holder. The four-domain cluster carries Acme's exact name and all four have mail exchangers, so its unidentified holder is the main source of the risk rated here.

Likelihood

An impersonator needs only a registered lookalike with mail configured and a recipient who acts on the message. No exploit and no access to Acme systems is required. Twelve such domains are set up for mail today, although no use of any of them against Acme was observed during the assessment. Medium applies until the holder of acme-example.org, acme-example.info, acme-example.me and acme-example.us is known. If Acme holds those four, the finding falls to Informational. The likelihood is assessed as Medium.

Recommendation

It is recommended that Acme establish who holds each lookalike before acting against any of them, in this order:

- Confirm whether Acme, or a brand protection agent acting for it, holds **acme-example.org**, **acme-example.info**, **acme-example.me** and **acme-example.us**, by checking the registrar and DNS

hosting accounts Acme controls. Then mark every other registered domain in the list as Acme's own, a legitimate third party or unexplained. This check alone decides whether the finding stands at Medium.

- Add each unexplained domain with a mail exchanger to the blocked sender list on the mail gateway and the web proxy block list. Watch the same domains for content changes and for new certificates in the certificate transparency logs, the public record of every certificate issued, so that a change in use is seen early.
- Register close variants that are still free, such as **acme-example.ph**, and request a takedown from the registrar or hosting provider only where a domain shows Acme branding.

Technical Details

Technical evidence withheld from this sample. In the client report this section lists each lookalike domain with the technique that produced it, whether it resolves in DNS and its mail exchanger.

4. Publicly Reachable Pre-Production Environments

MEDIUM	Risk	Asset	20 pre-production hostnames
MEDIUM	Impact	Root Cause	Security Misconfiguration
MEDIUM	Likelihood	Affected Services	HTTPS (TCP 443)
6.3	CVSS:4.0 Score	CVSS:4.0 Vector	AV:N/AC:L/AT:P/PR:N/UI:N/VC:L/VI:L/VA:N/SC:N/SI:N/SA:N

Description

Pre-production environments are the development, test and staging copies of a service, used to build and check changes before they are released. They run code that has not yet been released and are not always maintained to the same standard as production, so they are normally reachable only by the people working on the service. The usual control is to limit who can reach them by network location.

The assessment identified 45 hostnames carrying a development, test, staging, UAT (user acceptance testing), sandbox or pre-production name. Of these, 25 fall through to the catch-all record, reach the website host and are covered in the catch-all DNS finding. The other 20 resolve to their own addresses, and 5 of them returned a working application. These are pre-production copies of the client web application at web.preprod-client.acme-example.com and web.us.preprod-client.acme-example.com, and of the web submission application, through which end users submit personal documents, at submit.preprod-client.acme-example.com, submit.staging-client.acme-example.com and submit.us.preprod-client.acme-example.com. The remaining 15 include development and staging project and bug tracking hosts, pre-production automation hosts and hsm.preprod-us.client.acme-example.com. That name discloses the use of a hardware security module, a dedicated device for generating and holding cryptographic keys, in the US pre-production environment.

Impact

Five pre-production applications are reachable by anyone on the internet, and three of them are copies of the application end users submit personal documents through. A pre-production copy can run unreleased code and may be patched and monitored less closely than production. Where it shares code with the live service, a weakness found in the copy can point to the same weakness in production.

For a provider handling customer documents the concern is that test submissions may have been made with real personal documents, or that an environment holds production credentials. If so, a weakness in a pre-production application could expose end users' personal documents, which are personal data subject to the security of processing obligation in GDPR Article 32, or give a route into the live service.

Likelihood

Reaching the applications requires nothing beyond an internet connection. No source address restriction was observed, and the hostnames are discoverable from passive DNS and public certificate records. A breach and leak index monitored for the assessment already lists a staging hostname. Turning that access into impact is harder: it needs an exploitable weakness in a pre-production application or real data held in one, and neither was established, because the contents of the applications were not tested. The likelihood is assessed as Medium.

Recommendation

It is recommended that Acme decide for each pre-production hostname whether it needs to be reachable from the internet, starting from the position that it does not, and then:

- Restrict pre-production applications by source address, behind a VPN or an allow list of Acme addresses. For applications served through Amazon CloudFront, this is an AWS WAF IP set rule on

the distribution's web ACL that admits only those addresses, or CloudFront signed access, which serves content only to requests carrying a signed URL or signed cookie.

- Confirm that no pre-production environment holds real personal documents or production credentials. Remove any that are found, rotate the affected credentials and use synthetic test documents in their place.
- Remove the DNS records for pre-production environments that are no longer in use, so that retired names do not point at addresses that could later be reassigned to another party.

.

Technical Details

Technical evidence withheld from this sample. In the client report this section lists each pre-production hostname with the addresses it resolves to, the hosting provider and the web response returned.

5. Staff Account Password Hash in Breach Data

MEDIUM	Risk	Asset	acme-example.com (1 staff mailbox)
HIGH	Impact	Root Cause	Security Misconfiguration
LOW	Likelihood	Affected Services	Microsoft 365 sign-in (HTTPS)
5.9	CVSS:4.0 Score	CVSS:4.0 Vector	AV:N/AC:H/AT:P/PR:N/UI:N/VC:L/VI:L/VA:N/SC:L/SI:N/SA:N

Description

When a website is breached, its account records are copied and later traded and indexed. A record usually holds the email address used to register and a hash of the password. A hash is a one-way transformation: it cannot be reversed directly, but an attacker who holds it can hash candidate passwords until one produces the same value. A member of staff who registers on an outside website with a work address places that address, and any password also used at work, into this data. bcrypt is a password hashing algorithm designed to be deliberately slow, so that every guess costs the attacker computing time.

Dark web monitoring, drawn from a third-party breach index, surfaced one breach record for an acme-example.com staff mailbox. The record comes from a breach of an unrelated website, not of Acme, and holds the password used on that website as a bcrypt hash. The record carries no breach date, so when the password was set or exposed is not known.

Impact

The impact described here is potential and was not demonstrated. If the password were recovered from the hash, and the same password is or was used on the staff member's Acme Microsoft 365 account, an attacker could sign in to Microsoft 365 as that member of staff from anywhere on the internet. That would give access to their mail and files for as long as the password stayed unchanged and the sign-in went unnoticed.

The account would also give the attacker a genuine acme-example.com sender address. Phishing sent from a real Acme mailbox does not depend on a lookalike domain, so customers and colleagues have little reason to doubt it, and a request from a known contact is more likely to be acted on. Replies would arrive in the compromised mailbox, where the attacker could read and answer them under the staff member's name.

Likelihood

An attacker would need to recover the password from the hash and find it in use on the staff member's Acme Microsoft 365 account. bcrypt is costly to crack unless the password is short or common, and a password set for another website may never have matched the Acme one. The monitoring platform rated this record High; the assessment rates it lower for these reasons. The rating would rise if multi-factor authentication (MFA), a second proof of identity beyond the password, is not enforced on this account. The likelihood is assessed as Low.

Recommendation

It is recommended that Acme take the following steps, in order of priority:

- Reset the password on the affected mailbox and confirm that MFA is enforced on it. In Microsoft Entra ID, check that a Conditional Access policy or security defaults require MFA for the account and that no exclusion covers it, because MFA is what stops a recovered password from being enough on its own.

- Review the Microsoft Entra ID sign-in logs for the account for successful sign-ins from unfamiliar locations or devices, as far back as retention allows, since the record carries no breach date. If any are found, revoke the account's active sessions and refresh tokens and treat the mailbox as compromised.
- Remind staff not to reuse work passwords on other websites. A password manager makes a unique password for every site practical, which removes the link between a breach elsewhere and an Acme account.

Technical Details

Technical evidence withheld from this sample. In the client report this section shows the matched breach record: the staff address, the hash type, a truncated hash and the breached source site.

6. Staff Mailboxes Indexed in Leak Sources

LOW	Risk	Asset	acme-example.com
MEDIUM	Impact	Root Cause	Security Misconfiguration
LOW	Likelihood	Affected Services	Email, Microsoft 365
3.7	CVSS:4.0 Score	CVSS:4.0 Vector	AV:N/AC:H/PR:N/UI:N/S:U/C:L/I:N/A:N

Description

A leak index is a third-party service that collects data from breaches and leaks and makes it searchable by selector, such as an email address or a domain. An entry shows that a value appeared in at least one indexed data set. It does not show that any account was accessed. The assessment identified 38 acme-example.com mailbox addresses and 8 acme-example.com hostnames in the breach and leak data held by the index. The addresses follow two naming conventions: the first name alone, and the first name and surname joined by a full stop. The list also contains shared role mailboxes for general enquiries, sales, support and recruitment.

The index records only that each address appears in leak data. It carries no password, hash or session data for any of the 38 addresses. One address also appears in a breach record that holds a password hash, which is reported in the separate breach-data finding. No acme-example.com account appears in the infostealer log data reviewed for this assessment. Infostealer logs hold the credentials and session cookies that malware harvests from infected devices.

Impact

The 38 addresses form a ready target list for phishing and for password spraying, in which one common password is tried against many accounts to stay under lockout limits. Because both naming conventions are known, further staff addresses can be worked out from a name alone, so the exposure is not limited to the 38 listed. The shared mailboxes for sales, support and recruitment receive messages from outside parties by design, so a lure written as a customer or candidate enquiry looks like normal traffic. The 8 hostnames confirm which systems exist, including app-dev.acme-example.com and the US portal at us.client.acme-example.com, without any scanning.

The finding is rated Low because no credential is attached to any address. The data tells an attacker whom to target, not how to sign in. A mailbox is at risk only if a password is obtained by other means and the sign-in controls do not stop its use.

Likelihood

The addresses and hostnames are already indexed and searchable, so obtaining them requires no access to Acme systems and no technical skill. A phishing message or a spraying attempt against the list needs little preparation. Turning an address into access still requires a step the leak data does not supply: a valid password, and a sign-in route that does not demand a second factor. The outcome therefore depends on the mail and sign-in controls on the Microsoft 365 tenant.

Recommendation

It is recommended that Acme:

- Confirm that multi-factor authentication is enforced for every mailbox on the list, including the shared mailboxes, and block legacy authentication across the Microsoft 365 tenant with a Conditional Access policy. Legacy authentication accepts a password without a second factor, so it bypasses multi-factor authentication. For a shared mailbox, which staff open through delegated access, blocking sign-in on its underlying account closes the same route.

- Review the 38 addresses against the current staff list and disable or remove the mailboxes of anyone who has left. A mailbox nobody reads will not report a phishing message or an unexpected sign-in prompt, and misuse of it can go unnoticed.
- Use the list as the audience for the next phishing awareness exercise. These are the addresses already within an attacker's reach, so an exercise aimed at them tests the people most exposed.

.

Technical Details

Technical evidence withheld from this sample. In the client report this section shows the leak index query, the selectors used and the 38 indexed addresses.

7. Catch-All DNS and Abandoned Cloud Endpoint

LOW	Risk	Asset	42 hostnames and static-assets.client.acme-example.com
MEDIUM	Impact	Root Cause	Security Misconfiguration
LOW	Likelihood	Affected Services	DNS, HTTPS (TCP 443)
3.1	CVSS:4.0 Score	CVSS:4.0 Vector	AV:N/AC:H/PR:N/UI:R/S:U/C:N/I:L/A:N

Description

A wildcard, or catch-all, DNS record answers for every name under a domain that has no record of its own, so a name that was never created, or has since been retired, still resolves. Apart from www.acme-example.com itself, the assessment identified 42 acme-example.com hostnames that resolve only to the two WP Engine addresses serving the main website, 38 of them as aliases of the same wpeproxy.com target. Of the 42, 41 return an identical Cloudflare HTTP 409 error and no content. They include names unlikely to have been created on purpose, such as production.acme-example.com and www.submit.us.preprod-client.acme-example.com. The pattern is consistent with a wildcard record for the zone. One of these names, admin.client.acme-example.com, has no application behind it but appears in the leak data reviewed during this assessment, which suggests a system that has since been retired and now falls through to the catch-all.

Separately, static-assets.client.acme-example.com points at an Amazon CloudFront distribution, a content delivery service, with an Amazon S3 storage bucket as its origin. It returns an S3 HTTP 404 with the error code NoSuchKey and presents a certificate that expired on 13 April 2025, which is consistent with a distribution that is no longer maintained.

Impact

No hostname takeover was identified, and none was attempted. The issue observed is DNS hygiene. Under a catch-all, retired hostnames never stop resolving, so a name such as admin.client.acme-example.com still answers after the system behind it appears to have been retired. This matters because the web submission application allows any acme-example.com subdomain to embed it: its frame-ancestors directive, which tells browsers which sites may display it inside their own pages, includes *.acme-example.com. A hostname that came under another party's control would therefore be trusted by the submission flow. No route to such control was found.

The static-assets endpoint appears unmaintained. Its bucket still exists, so the takeover pattern in which another party re-creates a deleted bucket does not apply. The catch-all also clutters the asset inventory: 41 of the 132 web applications in the assessment data are the same error page, which makes genuine exposure harder to spot.

Likelihood

Abuse would require bringing an acme-example.com hostname under another party's control, serving content on it with a valid certificate, and then bringing a user to it. No route to do so was found. An unclaimed name currently fails the HTTPS handshake because the hosting platform holds no certificate for it, and WP Engine is listed as not vulnerable in can-i-take-over-xyz, the public reference list of services open to subdomain takeover. The S3 origin behind static-assets.client.acme-example.com still exists. Nothing was claimed or registered to test either point. The likelihood is assessed as Low.

Recommendation

It is recommended that Acme remove the catch-all, retire the abandoned endpoint and limit which hostnames may embed the web submission application:

- Confirm whether a wildcard record exists for **acme-example.com** and whether it is needed. If it is not, remove it and create explicit records for the hostnames in use, so that retired names such as **admin.client.acme-example.com** stop resolving. Confirm with the hosting provider that a hostname under **acme-example.com** cannot be attached to another customer's account.
- Delete the CloudFront distribution behind **static-assets.client.acme-example.com** and its DNS record if it is unused. If it is still needed, renew its certificate.
- Replace ***.acme-example.com** in the **frame-ancestors** directive of the web submission application with the named hostnames that need to embed it. This removes the trust currently extended to every **acme-example.com** subdomain, including retired and unknown ones.

Technical Details

Technical evidence withheld from this sample. In the client report this section shows the catch-all DNS counts with sample hostnames, and the certificate and web response returned by the abandoned cloud endpoint.

8. Mail Transport Hardening Not Published

LOW	Risk	Asset	acme-example.com
LOW	Impact	Root Cause	Security Misconfiguration
LOW	Likelihood	Affected Services	SMTP, DNS
2.3	CVSS:4.0 Score	CVSS:4.0 Vector	AV:A/AC:H/AT:P/PR:N/UI:N/VC:L/VI:L/VA:N/SC:N/SI:N/SA:N

Description

Mail protection for a domain has two parts. Sender authentication lets a receiving server check that a message claiming to come from the domain is genuine: SPF lists the servers allowed to send, DKIM adds a cryptographic signature, and DMARC tells receivers what to do when neither check passes. Transport hardening protects mail on its way to the domain's own mail servers, and protects the DNS answers that tell senders where those servers are. The sender authentication for acme-example.com is in good order. The DMARC policy is p=reject, and the SPF record ends in -all with 9 of the 10 permitted DNS lookups used.

The transport side is not in place. No MTA-STS policy is published. MTA-STS (Mail Transfer Agent Strict Transport Security) is a policy a domain publishes to tell sending mail servers to deliver its mail only over TLS encryption with a valid certificate. No TLS-RPT record is published, so failed or downgraded TLS deliveries are not reported to Acme. The DNS zone is not signed with DNSSEC, so a resolver cannot tell a genuine answer for the domain's records from a forged one. DKIM keys sit under selector names that cannot be listed, and of seven common names checked, only selector1 returned a key. That key is 1024-bit RSA.

Impact

Without MTA-STS, a server sending mail to Acme can be made to deliver it unencrypted by an attacker on the network path who strips the offer of encryption from the connection, and the content of that mail could then be read in transit. Without TLS-RPT, sending servers have no address to report the failure to, so Acme would not be told it had happened. TLS-RPT enforces nothing on its own. Without DNSSEC, the same kind of attacker can forge the domain's MX, SPF and DMARC answers, which could direct inbound mail to a server of the attacker's choosing.

The DKIM key on **selector1** is below the 2048 bits recommended for DKIM signers in RFC 8301. A valid DKIM signature for the domain is enough for mail to pass DMARC, so anyone able to recover that private key could sign mail as Acme and the **p=reject** policy would not stop it.

Likelihood

The missing MTA-STS policy and the unsigned zone can each be used only from a position on the network path between mail servers, or on the path of the DNS lookups a sending server makes. An ordinary internet user does not hold that position, and gaining it generally means controlling network infrastructure that carries the traffic. The DKIM key needs no network position, but signing with it would first require the private key to be recovered from the published public key. The likelihood is assessed as Low.

Recommendation

It is recommended that Acme make the following changes.

- Publish a TLS-RPT record at _smtp._tls.acme-example.com with a reporting address, and an MTA-STS policy in testing mode. MTA-STS needs both a _mta-sts.acme-example.com TXT record and a policy file served at https://mta-sts.acme-example.com/.well-known/mta-sts.txt that lists the two

inbound mail gateway hosts. Review the reports for a few weeks, then change the policy mode to enforce.

- Rotate the **selector1** DKIM key to 2048 bits in the service that signs with it. Publish the new key under a new selector, move signing to it, and remove the 1024-bit key once mail signed with it has been delivered.
- Enable DNSSEC signing for the zone in **Amazon Route 53** and add the resulting DS record at the domain registrar to complete the chain of trust. This is the lowest priority of the three unless a customer or regulatory requirement applies. Separately, the SPF record uses 9 of its 10 permitted lookups, so the count should be checked before another sending service is added.

Technical Details

Technical evidence withheld from this sample. In the client report this section shows the mail records read from public DNS and the transport hardening records that were not found.

9. Vulnerable FTP Service on Legacy Hosting Address

INFO	Risk	Asset	192.0.2.4, 192.0.2.3
HIGH	Impact	Root Cause	Patch Management
INFO	Likelihood	Affected Services	FTP (TCP 21), HTTP, HTTPS
N/A	CVSS:4.0 Score	CVSS:4.0 Vector	AV:N/AC:L/AT:N/PR:L/UI:N/VC:H/VI:H/VA:H/SC:N/SI:N/SA:N

Description

acme-example.com currently resolves to 192.0.2.15 and 192.0.2.16 on WP Engine. Passive DNS history links the domain to six further addresses that it does not resolve to today. Two of them still answer and carry findings.

192.0.2.4 is on a shared hosting provider. It answers on TCP 21, 80 and 443. Port 21 runs ProFTPD 1.3.5e, against which the platform matched 12 CVEs, 10 of them at CVSS 7.0 or above and the highest at 8.7. Ports 80 and 443 serve a hosting control panel default page behind a self-signed certificate for odin.com with a 1024-bit key. The FTP service presents a certificate in the hosting provider's own name, which points to a shared hosting server.

192.0.2.3 is on Amazon Web Services in London. It serves nginx on TCP 80 and 443. The certificate check found a self-signed certificate issued to the address itself that expired on 27 August 2020, and the service accepts two weak CCM8 cipher suites. Internet exposure data from September 2026 shows the same HTTPS port presenting a current certificate for an unrelated third-party domain, which suggests the address has since been reassigned to another AWS customer.

Impact

If either address still belongs to Acme, it is an unmanaged system outside the estate the team actively looks after. A forgotten hosting account commonly holds an old copy of the website, its database and the credentials that went with it. The most serious of the ProFTPD issues need an authenticated FTP user, so the practical route is a reused or leaked password followed by one of the code execution flaws.

If the addresses no longer belong to Acme, the findings are somebody else's and should come out of scope. The question then is only whether any Acme content or credentials were left behind when the hosting was given up.

Likelihood

Undetermined pending ownership confirmation. The link to acme-example.com is historical DNS only, every CVE is unverified, and a shared hosting server is patched by the provider on its own schedule. No exploit prevalence data was available for these CVEs. For 192.0.2.3 the third-party certificate makes it likely the address is no longer Acme's. The rating rests on 192.0.2.4 and will be revised once ownership is confirmed.

Recommendation

It is recommended that Acme confirm whether it still holds 192.0.2.4, 192.0.2.3 and the four other historical addresses: 192.0.2.17, 192.0.2.18, 192.0.2.19 and 203.0.113.1.

If Acme still holds the shared hosting account, retrieve anything needed from it, close the account and confirm FTP access is gone. If the AWS address is Acme's, terminate the instance or bring it under management. Any address confirmed as not Acme's will be removed from the next assessment.

Either way, rotate any password that was ever used for FTP or the hosting control panel on the old provider.

Technical Details

Technical evidence withheld from this sample. In the client report this section gives, for each address, the hosting provider, open ports, service banner and the 12 CVEs matched by version.

10. Microsoft 365 Tenant and Realm Discovery

INFO	Risk	Asset	acme-example.com tenant
LOW	Impact	Root Cause	Security Misconfiguration
INFO	Likelihood	Affected Services	Microsoft Entra ID
N/A	CVSS:4.0 Score	CVSS:4.0 Vector	AV:N/AC:L/PR:N/UI:N/S:U/C:N/I:N/A:N

Description

Every Microsoft 365 customer has a tenant in Microsoft Entra ID, the identity service that holds its user accounts and handles sign-in. Microsoft publishes two answers about every tenant. The user realm endpoint tells the sign-in page where to send a user, based on the domain of the address typed in. The **OpenID Connect** discovery document tells applications where to sign users in and which keys sign their tokens. Both answer any unauthenticated request, and neither can be switched off.

The assessment identified that acme-example.com is a managed Entra ID tenant, which means users sign in at the Microsoft sign-in page rather than being redirected to a separate federation service. Both answers return the tenant's sign-in configuration to an unauthenticated request. The SPF record for the domain, which lists the services allowed to send mail as acme-example.com, includes Microsoft, HubSpot, Salesforce, SendGrid and Pardot. The platform also reported reachable sign-in pages for the name acme at acme.auth0.com, acme.okta.com and acme.onelogin.com. Each of those platforms answers a made-up tenant name in exactly the same way, so there is no evidence that Acme has a tenant on any of them. Whether multi-factor authentication (MFA), a second check in addition to the password, is enforced cannot be seen from outside and was not assessed.

Impact

These answers are the same for every Microsoft 365 customer, so they do not reflect a choice Acme has made. They matter only as the place a captured password would be used. An attacker who obtains an Acme password knows to present it at the Microsoft sign-in page, which is reachable from anywhere on the internet. The same knowledge tells an attacker which sign-in page to copy when building a phishing page to capture that password.

What a captured password is then worth depends on whether the tenant requires MFA at sign-in, and that is not visible from outside. A Microsoft 365 account normally opens a member of staff's mailbox and files, so if a password alone is accepted, the exposure is the content of that account. The sign-in pages reported on the three other platforms add no exposure, as nothing indicates Acme uses them.

Likelihood

Reading the realm and discovery answers takes one unauthenticated request and no particular skill, so obtaining the information is not a barrier. The information does not provide a route in by itself. To make use of it, an attacker must first capture a valid password for an Acme user and then pass whatever sign-in controls the tenant applies, and this finding provides neither. The likelihood is assessed as Informational.

Recommendation

It is recommended that Acme harden its Microsoft sign-in:

- Enforce MFA for every user through conditional access with no exclusions, and block legacy authentication. Conditional access is the Entra ID feature that sets the conditions an account must meet to sign in, and the **Require multifactor authentication for all users** and **Block legacy**

authentication templates cover both steps. Legacy protocols cannot prompt for MFA, so leaving them open undoes the first control.

- Turn on a sign-in risk policy, such as the **Require multifactor authentication for risky sign-ins** template (Microsoft Entra ID P2 licence), and alert on password spray, where one common password is tried against many accounts. Add company branding to the sign-in page so that a generic copied page looks different from the genuine one and is easier for staff to spot.
- Confirm whether Acme uses Auth0, Okta or OneLogin. If it does, apply the same MFA requirement there. If it does not, no action is needed on those three.

Technical Details

Technical evidence withheld from this sample. In the client report this section shows the realm discovery request and the identity surfaces it confirmed, with a command to reproduce it.

11. Salesforce Aura API Exposed

INFO	Risk	Asset	content.acme-example.com
LOW	Impact	Root Cause	Insecure Development
INFO	Likelihood	Affected Services	HTTPS (TCP 443)
N/A	CVSS:4.0 Score	CVSS:4.0 Vector	AV:N/AC:H/AT:P/PR:N/UI:P/VC:L/VI:N/VA:N/SC:N/SI:N/SA:N

Description

Salesforce Experience Cloud is the Salesforce product for building public websites and portals on top of an organisation's Salesforce environment. Visitors who have not signed in are served as the site's guest user, a Salesforce user whose profile and sharing rules decide which objects and records an anonymous visitor can read. The Aura endpoint, /aura, is the interface the site's pages call to fetch and submit Salesforce data, and it answers unauthenticated requests on every Experience Cloud site by design. What the endpoint returns is therefore limited by the guest user's access, not by the endpoint itself.

The assessment identified that content.acme-example.com is a Salesforce Experience Cloud site and that its /aura endpoint answered an unauthenticated request. An endpoint that responds in this way is expected behaviour for the platform and is not a weakness in itself. The guest user profile for the site was not tested. No Aura query was issued, so the assessment did not establish which objects or records, if any, the guest user can read.

Impact

No impact was demonstrated. No Salesforce record was requested or returned, and the assessment holds no evidence either way on what the guest user of **content.acme-example.com** can read.

The realistic potential lies in the guest user configuration. Where guest access is too broad, the Aura API lets an anonymous visitor query Salesforce objects directly, with no account and no credential. Such an exposure would not depend on any flaw in Salesforce: it would follow from access granted to the guest user in Acme's own configuration, which is why the root cause is recorded as Security Misconfiguration. Over-broad guest access is a known cause of customer and lead data exposure from Salesforce sites (Krebs on Security, "Many Public Salesforce Sites are Leaking Private Data", April 2023). Customer and lead records usually hold personal data, and an unauthorised disclosure of them would be a personal data breach under the GDPR.

Likelihood

Reaching the **/aura** endpoint requires no account, no credential and no knowledge beyond the hostname, and the endpoint is present on every Experience Cloud site. That barrier is low, but reaching the endpoint achieves nothing unless the guest user can read data the public pages do not need. The guest user configuration was not tested and no Aura query was issued, so no route to data was observed. The likelihood is assessed as Informational.

Recommendation

It is recommended that Acme confirm the guest user of the **content.acme-example.com** site can read only what its public pages need:

- Review the guest user profile and the sharing rules for the site, and remove read access to any object or field the public pages do not need. This is the configuration that decides what the **/aura** endpoint returns to an anonymous visitor.

- Run Salesforce's guest user access report and check each object and permission it lists against what the public pages require. Access that cannot be tied to a page on the site should be removed.
- Confirm that the **Secure guest user record access** setting is turned on in the org's Sharing Settings. With it on, guest users can be given access to records only through guest user sharing rules, which keeps that access visible and reviewable.

Technical Details

Technical evidence withheld from this sample. In the client report this section shows the detection result, the matched URL and the response details, with a command to reproduce it.

12. Misconfigured Cookie Attributes

INFO	Risk	Asset	9 assets, listed in Technical Details
MEDIUM	Impact	Root Cause	Security Misconfiguration
INFO	Likelihood	Affected Services	HTTPS (TCP 443)
N/A	CVSS:4.0 Score	CVSS:4.0 Vector	AV:N/AC:H/AT:P/PR:N/UI:P/VC:L/VI:L/VA:N/SC:N/SI:N/SA:N

Description

Cookies carry session state between a browser and a web application, and the attributes set with each cookie tell the browser when it may send it. The **Secure** attribute instructs the browser to send a cookie only over an encrypted HTTPS connection, never over plain HTTP. **SameSite** controls whether a cookie accompanies requests started from another site, and **HttpOnly** keeps a cookie out of reach of page scripts.

The assessment identified that the EU and US client portals, **client.acme-example.com** and **app.us.client.acme-example.com**, set the **laravel_session** session cookie and the **XSRF-TOKEN** anti-forgery cookie without the **Secure** attribute. Both portals send a **Strict-Transport-Security** header, which tells a browser to refuse plain HTTP for that host. The header arrives in the same response that sets the cookies, so the gap is narrow. The other 7 assets return cookies set by third-party platforms: a Pardot tracking cookie on **try.acme-example.com**, Salesforce consent and locale cookies on **content.acme-example.com**, and the Cloudflare bot management cookie **__cf_bm** on the main website, blog, thank-you and trust centre pages, which Cloudflare sets with **SameSite=None** by design.

Impact

A cookie without **Secure** is sent over plain HTTP whenever the browser is directed to the host that way. An attacker on the same network as a portal user, such as shared or public Wi-Fi, could read **laravel_session** in transit and use it to take over that user's portal session, with the same access the user has to customer submissions.

Customer submissions hold personal data about the individuals who submit them, so access gained this way would amount to a personal data breach under GDPR. No session was intercepted or used during the assessment; the impact is assessed from the configuration observed. The third-party cookies are of low value to an attacker.

Likelihood

An attacker would need a position on the same network as a portal user, and would need that user's browser to make a plain HTTP request to a portal, for example by following an **http://** link. Both portals send **Strict-Transport-Security** with a one-year **max-age** and **includeSubDomains**, and the cookies are set in the same HTTPS response, so a browser holding a portal cookie normally also holds the instruction to convert such a request to HTTPS before it is sent. The exposure is limited to a browser that holds a portal cookie without the stored policy, for example after its site data has been partly cleared.

Recommendation

It is recommended that Acme:

- Set the session cookies to **Secure** in the Laravel session configuration for both the EU and US environments, using the **secure** option in **config/session.php** or the **SESSION_SECURE_COOKIE** environment variable it reads. Laravel applies the same option to **XSRF-TOKEN**, so one change per environment covers both cookies.

- Set the **same_site** option in the same file to **lax** or **strict**, so the browser does not send **laravel_session** on requests started from another site. Keep a less restrictive value only if a portal must be embedded in another site's pages.
- Leave the third-party cookies unchanged. Pardot, Salesforce and Cloudflare set them, and **SameSite=None** on **__cf_bm** is Cloudflare's intended behaviour.

Technical Details

Technical evidence withheld from this sample. In the client report this section lists each affected cookie by asset, with the attributes that were and were not set.

13. Content Security Policy (CSP) Header Misconfiguration

INFO	Risk	Asset	54 assets, listed in Technical Details
INFO	Impact	Root Cause	Security Misconfiguration
INFO	Likelihood	Affected Services	HTTPS (TCP 443)
0.0	CVSS:4.0 Score	CVSS:4.0 Vector	AV:N/AC:L/AT:N/PR:N/UI:N/VC:N/VI:N/VA:N/SC:N/SI:N/SA:N

Description

A Content Security Policy (CSP) is a response header that tells the browser which sources a page may load scripts and other content from. Subresource Integrity (SRI) is an attribute that lets the browser check that a script loaded from another origin has not been changed. The assessment identified policies that leave script loading unrestricted or weakly restricted, and scripts loaded without an integrity attribute, across 49 assets.

The web submission application publishes a policy of the same shape on five hostnames covering production, US, staging and pre-production. It contains only a frame-ancestors directive, which controls which sites may embed the page, so it places no limit on where scripts load from. The documentation site, docs.acme-example.com, allows inline scripts, eval and a wildcard source; it is hosted on GitBook, which sets the policy. The policies on acme-example.com, the trust centre at trust.acme-example.com and the Pardot pages behind try.acme-example.com are set by the platforms that host them, and those on acme-example.com and trust.acme-example.com permit 'unsafe-eval'. Scripts from other origins load without an integrity attribute on 47 hosts. Of these, 41 are the catch-all error page described in the catch-all DNS finding. The other six are docs.acme-example.com and the five web submission hostnames, with the production and pre-production instances loading a third-party error-tracking script from its vendor CDN.

Impact

The web submission application is where end users upload their documents. Its policy places no limit on where scripts load from or where a script running in the page can send data. If a third-party script the application loads were altered at its source, or an injection flaw allowed script to be added to the page, the browser would run it without restriction, and it could read what the user submits.

The gap is one of containment rather than exposure. The weaker policies on the documentation site, the corporate site and the trust centre are set by their hosting platforms, so what Acme can change there depends on each platform. The web submission policy is the one Acme sets, and it covers the application that receives personal documents.

Likelihood

Making use of the missing script policy needs a second flaw to combine with, such as a compromised third-party script or a content injection weakness in the web submission application. The assessment identified neither. The headers were read passively from public responses and no attack was attempted. Adding script directives would limit the damage such a flaw could do; their absence does not create an entry point. The likelihood is assessed as Informational.

Recommendation

It is recommended that Acme:

- Add **script-src**, **object-src 'none'** and **base-uri 'self'** to the web submission policy alongside the existing **frame-ancestors** directive, with **script-src** limited to the application's own origin and the third-party sources it needs. Publish the change first as **Content-Security-Policy-Report-Only**

with a reporting endpoint, so every source the application relies on is identified before the policy is enforced.

- Add Subresource Integrity to scripts loaded from CDNs, starting with the error-tracking script on the web submission application. An integrity attribute holding the file's SHA-384 hash, with crossorigin="anonymous", makes the browser refuse a copy that has been altered.
- After enforcement, review the violation reports and remove sources the application no longer uses. Where GitBook and the hosts of the corporate site and trust centre allow the policy to be configured, remove **'unsafe-inline'**, **'unsafe-eval'** and wildcard sources there too.

Technical Details

Technical evidence withheld from this sample. In the client report this section shows the policy returned by each host, with external origins withheld, and the results of the weak policy check.

14. Customer Relationships Disclosed in CSP Header

INFO	Risk	Asset	Web submission application, five hostnames
INFO	Impact	Root Cause	Security Misconfiguration
INFO	Likelihood	Affected Services	HTTPS (TCP 443)
0.0	CVSS:4.0 Score	CVSS:4.0 Vector	AV:N/AC:L/AT:N/PR:N/UI:N/VC:N/VI:N/VA:N/SC:N/SI:N/SA:N

Description

A Content Security Policy is a response header containing rules that the browser enforces for a page. Its **frame-ancestors** directive names the external sites permitted to display the page inside a frame on their own pages, and the browser refuses to frame the page anywhere else. A service that customers embed in their own websites therefore has to list each customer's site in this directive for the page to load there. The header is sent in full with every response, so its contents are visible to anyone who requests the page.

The assessment identified a **Content-Security-Policy** header of the same form on all five hostnames of the web submission application, across its production, staging and pre-production environments. The policy is one shared list, sent to every visitor. On the production and pre-production hostnames its **frame-ancestors** directive lists 280 external origins, meaning website addresses outside Acme's own domains, alongside Acme's wildcards **https://*.client.acme-example.com** and **https://*.acme-example.com**. The staging hostname lists 234. The external origins name Acme customers, including regulated financial services firms, together with their development and test domains.

Impact

Anyone can read which organisations use Acme. A customer list is commercially sensitive, and the header gives it to competitors and the public with no effort required. A customer that has not announced its use of Acme has the relationship disclosed regardless. The development and test domains go further, naming customers' own non-production systems, so Acme is publishing information about its customers' systems as well as its own.

The list is also useful to anyone planning an impersonation campaign. It names the brands whose end users expect to see an Acme page, so a fake page sent under one of those brands would match what those end users already expect. The header itself gives no access to Acme systems or data.

Likelihood

Reading the list requires no account, no specialist tool and no unusual request. The header is returned to every unauthenticated visitor, and a browser's developer tools display it. Putting the list to use is a separate step carried out entirely outside Acme's systems, such as building lookalike pages under a customer's brand. The rating reflects a disclosure of business information rather than a weakness in the platform itself. The likelihood is assessed as Informational.

Recommendation

It is recommended that Acme stop sending the full customer list to every visitor:

- Serve the policy per customer, so that **frame-ancestors** in each response names only the site embedding the page, plus Acme's own origins where they are needed, in place of one shared list. A visitor then sees at most the one customer whose page they already have open, and the framing restriction still applies.
- Remove the entries for customers and test domains no longer in use, and add removal of a customer's origins to the offboarding process so that the list does not build up again.

- Keep customer development and test domains out of the policy on the production hostnames, and permit them only on the staging and pre-production hostnames. This reduces what production discloses while the per-customer policy is built.

Technical Details

Technical evidence withheld from this sample. In the client report this section shows the frame-ancestors directive with customer origins withheld, a command to reproduce it, and the hosts affected.

15. Internet-Facing Sign-In Surfaces

INFO	Risk	Asset	Internet-facing sign-in surfaces
INFO	Impact	Root Cause	Security Misconfiguration
INFO	Likelihood	Affected Services	HTTPS (TCP 443)
0.0	CVSS:4.0 Score	CVSS:4.0 Vector	AV:N/AC:L/AT:N/PR:N/UI:N/VC:N/VI:N/VA:N/SC:N/SI:N/SA:N

Description

A sign-in surface is any page or service that accepts a username and password. Customer portals, cloud mail and content platforms are normally reachable from the internet, and that reachability is not a weakness by itself. What decides the risk is whether a password alone is enough to sign in, or whether multi-factor authentication (MFA), a second proof such as a code from a phone app, is also required.

The assessment identified seven Acme hostnames that accept a login over HTTPS on TCP port 443: three for the client portal in the EU and US, two for the client web application, an SDK authentication endpoint at **api-auth.acme-example.com**, and the WordPress site at **acme-example.com**. The Microsoft 365 tenant for **acme-example.com**, four third-party platforms (Salesforce, HubSpot, GitBook and SafeBase) and the pre-production and staging copies reported in another finding also accept a login. No sign-in was attempted, and whether MFA is enforced was not tested on any of them.

Impact

Each surface is a place where a valid password is enough to sign in unless a second factor is enforced. Other findings in this assessment report leaked customer credentials and a breached staff password hash. Whether either can be used depends on whether the surface the account belongs to requires MFA. That question was not tested, and its answer decides how much risk those findings carry.

A customer account on client.acme-example.com or us.client.acme-example.com gives access to whatever that customer holds in the service, and a staff account opens the organisation's Microsoft 365 services. This finding records the surfaces and carries no direct impact of its own, so the impact is rated Informational.

Likelihood

The surfaces are discoverable from public DNS and respond to anyone on the internet. Using one requires a valid password for an account on that surface, and then the absence of a second factor. Rate limiting and account lockout were not tested, so no claim is made that password guessing against any surface would succeed. Reachability of a sign-in page is normal for a customer-facing service, and this finding records an observation rather than a weakness. The likelihood is assessed as Informational.

Recommendation

It is recommended that Acme take the following steps for each surface listed in the Technical Details, and record the result against each entry:

- Confirm that MFA is mandatory for every client portal user and every staff account. For Microsoft 365, that means an Entra ID Conditional Access policy requiring MFA for all users, excluding only named emergency access accounts. This answer decides the risk carried by the leaked customer credentials and the breached staff hash.
- Alert on sign-ins from a new device or country, for portal users and staff alike. The Entra ID Protection detection for unfamiliar sign-in properties covers staff accounts. The client portals need an equivalent alert built from their own authentication logs.

- Restrict administrative sign-in surfaces to known source addresses where practical: the WordPress administration login on **acme-example.com**, the pre-production and staging copies, and administrator access to the third-party platforms where supported. Customer-facing portals cannot be restricted this way, which leaves MFA as their main control.

Technical Details

Technical evidence withheld from this sample. In the client report this section lists each sign-in surface with its hosting and the addresses it resolves to.

16. Attack Surface Composition

INFO	Risk	Asset	Organisation-wide
INFO	Impact	Root Cause	Security Misconfiguration
INFO	Likelihood	Affected Services	All
0.0	CVSS:4.0 Score	CVSS:4.0 Vector	AV:N/AC:L/AT:N/PR:N/UI:N/VC:N/VI:N/VA:N/SC:N/SI:N/SA:N

Description

An organisation's attack surface is every internet-facing service an outside party can reach, including those that suppliers run on its behalf. The asset inventory lists 387 assets: 86 subdomains, 132 web applications, 130 addresses (121 **IPv4** and 9 **IPv6**), 38 mailbox addresses and the domain itself. The web application count is inflated: 41 are the catch-all error page reported in a separate finding, 64 are requests made to bare content delivery network (CDN) and hosting addresses, most of them **CloudFront** edges returning **HTTP 403**, and 27 are named applications. Addresses that belong to CDN, load balancer and hosting providers are shown against the hostnames that resolve to them and are not listed as assets.

The core estate is small. The client portal and web submission applications run on Amazon Web Services in Ireland and the US, and the marketing website at www.acme-example.com runs on WordPress, hosted by WP Engine behind Cloudflare. The other services are third-party platforms: GitBook for documentation, SafeBase for the trust centre, and HubSpot, Salesforce and Pardot for marketing and customer relationship management (CRM). The website discloses the version numbers of its page builder and the names of four plugins. Plugins are add-ons to WordPress, and their names appear in the file paths each page loads. No vulnerability was matched against these versions in this assessment.

Impact

No weakness is reported in this finding. Most of Acme's internet-facing surface is operated by other companies, so its security rests on how each platform is configured and on who holds the administrator accounts. Control of one of those accounts would give control of the service it carries, such as the website content or the registration of **acme-example.com**, without any access to Acme's own applications. For a service provider, the website, trust centre and documentation are what customers and prospects read when they assess the service.

The plugin names and versions tell an outside party which components to check against published vulnerabilities, without sending any unusual request. While the components are kept current, this is reconnaissance value only. A plugin left on an older version becomes identifiable from outside as soon as a vulnerability in that version is published.

Likelihood

Everything in this finding was observed without authentication, from public DNS records and from responses any visitor receives. No weakness is reported, so there is nothing for an attacker to act on directly. The information shortens an attacker's preparation and becomes useful only in combination with a weakly protected platform account or an out-of-date plugin. Whether the third-party administrator accounts require multi-factor authentication (MFA) was not assessed. The likelihood is assessed as Informational.

Recommendation

It is recommended that Acme treat the third-party platforms as part of its estate, with a named owner, enforced MFA and current software for each.

- Record a named owner and an MFA requirement against each administrator account on **WP Engine**, **Cloudflare**, **HubSpot**, **Salesforce**, **GitBook**, **SafeBase** and the DNS registrar, and confirm in each platform's own settings that the requirement is enforced. These accounts control the services that sit outside Acme's own applications.
- Keep WordPress core and all plugins on automatic updates, and remove any plugin that is no longer used. Premium plugins are usually updated through the vendor's licence rather than the WordPress plugin directory, so confirm that each licence is active and that updates arrive.
- Review the provider roll-up in Technical Details whenever a supplier is added, replaced or retired, and update the owner and MFA records to match. Close any accounts left behind on a retired supplier at the same time.

Technical Details

Technical evidence withheld from this sample. In the client report this section shows the provider roll-up by category and the technology detections behind it.

Appendix A – Our Approach

Continuous Exposure Assessment

WhatsExposed worked with Acme personnel to deliver this engagement. WhatsExposed have developed a proven continuous exposure assessment methodology that follows industry and vendor best practices including, but not limited to, OWASP, OSSTMM, CREST and NIST to assess and quantify the risks of an organisation's internet-facing estate.

Phase 0 **Planning & Scoping**	WhatsExposed and Acme engaged in initial discussions to understand requirements and project scope. All the necessary pre-requisites were discussed and met prior to start date. Any limitations, risks or areas of concern were flagged to WhatsExposed, and appropriate measures taken to ensure successful project delivery.
Phase 1 **Passive Reconnaissance**	Our team began the assessment by carrying out passive information gathering against the in-scope organisation. Public DNS and certificate records, internet exposure data and dark web leak sources were used to identify the domains, hostnames, addresses, web applications and identities that make up the internet-facing estate.
Phase 2 **Active Reconnaissance**	An active scan of the discovered assets and their underlying infrastructure was conducted to identify exposed services, security misconfigurations and publicly known vulnerabilities associated with them.
Phase 3 **Correlation & Analysis**	The results were consolidated and correlated across assets. Related findings were chained to derive the exposure paths an attacker could follow. No exploitation was performed, and findings that rest on inference are identified as such in the report.
Phase 4 **Reporting**	The findings are categorised into root causes and risk assessed according to WhatsExposed's risk ratings – Critical, High, Medium, Low and Informational. WhatsExposed outlined detailed remediation procedures to mitigate the risks associated with the identified findings. We shall endeavour to discuss the findings with Acme to address any queries the team may have by means of a walkthrough of the report should you wish to do so.
Phase 5 **Follow-Up & Re-Assessment**	The assessment is repeated on a continuous basis. Each subsequent report sets out what has changed since the previous assessment. WhatsExposed shall assist Acme with validation of remediation procedures or provide further guidance with any security improvement programmes.

Appendix B – Definitions

Root Causes

Security Misconfiguration A security misconfiguration arises when security controls of an asset are not defined, implemented or implemented incorrectly. This implies that the configuration settings do not comply with industry standards which are critical to maintaining security and reduce business risk.

Insecure Development Vulnerabilities arising from poorly developed applications which do not follow secure software development practices. Poorly documented secure coding standards and a lack of secure code reviews results in insecure code being developed.

Patch Management Security vulnerabilities resulting from outdated, unsupported, or unpatched software which contain publicly known vulnerabilities. Insufficient patch management hinders the ability to correct software errors in a timely manner.

Risk Ratings

The Risk Ratings are based on the standard risk model: **Risk = Impact x Likelihood**

Impact	High	MEDIUM	HIGH	CRITICAL
	Medium	LOW	MEDIUM	HIGH
	Low	LOW	LOW	MEDIUM
	INFORMATIONAL	Low	Medium	High
		Likelihood		

Impact

The testing team calculates the scores by combining a variety of factors, this includes vulnerability rankings by well-known security organisations such as CVSS scores and a direct examination of the impact in the context of the tested environment. Each finding can have unique consequences but in general, the ratings are determined by inspection of how a vulnerability affects the following factors; confidentiality, integrity, availability as well as any reputational or financial cost. Impact is classified into the following four categories:

High	The vulnerability presents a high security threat to the data, host or system in scope that could affect the confidentiality, integrity or availability of large amounts of sensitive data. As a result, significant financial and/ or reputational damage along with non-compliance penalties are common effects of such threats.
Medium	The vulnerability presents a medium security threat to the data, host or system in scope that could be exploited under certain conditions. These are often limited in scope such as the disclosure of small amounts of confidential information or requiring user intervention to gain further access to information assets. Moderate financial and/ or reputational damage along with alleviated non-compliance penalties are common effects of such threats.
Low	The vulnerability presents a small effect on the security posture of the affected system. Examples of this include disclosure of information that an adversary could leverage to carry out more sophisticated attacks. Trivial consequences or penalties could be incurred.
Informational	The vulnerability does not pose a direct threat to the business. It is raised for informational purposes only which could help the organisation to improve the security posture of the system. It is possible that the vulnerability could be used in conjunction with another, leading to a more sophisticated attack scenario.

Likelihood

The testing team calculates likelihood based on the ease of exploiting a vulnerability. This accounts for numerous factors such as the threat actor's skill level, opportunity and motive, along with the vulnerability's ease of discovery and exploitation. Likelihood is classified into the following four categories:

High	Exploitation requires minimal requirements and an adversary with limited technical skills. Vulnerability discovery could be carried out using an automated tool. There are limited mitigating factors implemented by the organization that prevent the vulnerability from being exploited.
Medium	Exploitation requires some additional requirements which reduce the situations in which it can be exploited. An adversary with moderate technical knowledge, coupled with another attack vector such as social engineering and user interaction may be required to exploit this vulnerability. The organisation applied several mitigating factors which prevents the vulnerability from being easily exploitable.
Low	Exploitation requires a large amount of time and resources which significantly reduce the likelihood of success. An adversary with advanced technical knowledge and other complex attack vectors are required to exploit this vulnerability. The organisation applied adequate mitigating factors which causes the vulnerability to be difficult to identify and exploit.

Informational	Exploitation is unlikely in practice by a sophisticated adversary with unlimited amount of time and resources. Complex and chained attack tactics would be required to successfully exploit this vulnerability. The organisation applied mitigating factors which considerably reduce the likelihood of its exploitation.

Risk

Once the impact and likelihood are determined, the risk is calculated by comparing these two factors. The risks should be used in order to determine the order of priority to address each risk.

CRITICAL	Critical risks should be addressed as a matter or urgency. Restricting access to the vulnerable system should be considered until remediation has taken place.
HIGH	High risks should be addressed as soon as possible to ensure that all assets are adequately protected. Monitoring, reviewing system logs and responding to incidents should be carried out imminently until remediation actions occur.
MEDIUM	Medium risks often do not require urgent remediation; however, if left unchecked for a long period of time, more severe vulnerabilities may arise. As such these should be addressed in the near future.
LOW	Low risks should be treated as the lowest threat that a system can possess and as such should be addressed in the longer term.
INFORMATIONAL	Informational risks should be reviewed but not necessarily remediated. Often, these can be used as part of a defence in depth approach to obtain the highest security posture possible.

Appendix C – Asset Inventory

Withheld from this sample. In the client report this appendix lists every asset discovered: 130 addresses with provider, open ports and services; 87 hostnames with what they resolve to and whether they follow a pre-production naming pattern; and 132 web applications with status, server and technology.

Appendix D – Supporting Scan Data

Withheld from this sample. In the client report this appendix holds the supporting scan data: known vulnerabilities by host with CVSS scores, detection results by template, and the identity exposure lists. Passwords and hashes are never reproduced in a client report.