Release-readiness audit

How much bot traffic gets past a launch, sign-up or drop flow, measured on staging, with evidence the client can check.

CLIENT	Sample (lab target)
FLOW TESTED	Sign-up (demo shop, lab)
TEST WINDOW	20 November 2026, 20:38 to 20:39 UTC
ENGAGEMENT REF	not supplied
EVIDENCE BUNDLE	kessler-evidence.json, generated 20 November 2026, 20:40 UTC
BUNDLE CHECKS	Checksum PASS; signature PASS (see page 5)
PREPARED BY	Somer North Ltd

WHAT THIS REPORT IS

A short summary of one authorised, staging-only test of the flow above. The detailed engineering evidence is in the appendix named on page 5. Each figure below is counted from the signed bundle, and each interval is a 95% Wilson range for the count it sits beside.

What got through

Flow classified as OWASP Automated Threats OAT-019 Account Creation (OWASP Automated Threats to Web Applications)

> **No defence was in the path for these runs. This report shows what the flow accepted, not what a defence stopped.**

BOT TIERS, SIMPLEST FIRST

- 5 of 5 plain HTTP client (tier 0b) sessions completed the flow (95% range 57-100%).
- 5 of 5 Chrome-impersonating HTTP client (tier 0) sessions completed the flow (95% range 57-100%).
- 5 of 5 header-matched HTTP client (stealth) sessions completed the flow (95% range 57-100%).
- 5 of 5 real browser driven as a bot (browserbot) sessions completed the flow (95% range 57-100%).

THE REAL-BROWSER BASELINE (FALSE-POSITIVE VIEW)

- 5 of 5 real-browser baseline sessions completed the flow (95% range 57-100%).
- Real-browser baseline blocks were NOT RUN: no defence was in the path for these runs.

The real-browser baseline is the control for false positives. A block here would be a legitimate visitor turned away.

> ### HOW TO READ THE RANGES
>
> Each tier had 5 configured units. A 95% range is wide at that size, so read the range, not the single percentage.

Per-tier outcome

Each cell is a count of sessions out of the units offered. NOT RUN means the stage did not happen, or there was nothing to divide by.

TIER	OFFERED	REACHED THE TARGET	CHALLENGED AT THE GATE	COMPLETED THE FLOW
plain HTTP client (tier 0b)	5 / 5 (100%)	5 / 5 (100%)	*NOT RUN*	5 / 5 (100%)
Chrome-impersonating HTTP client (tier 0)	5 / 5 (100%)	5 / 5 (100%)	*NOT RUN*	5 / 5 (100%)
header-matched HTTP client (stealth)	5 / 5 (100%)	5 / 5 (100%)	*NOT RUN*	5 / 5 (100%)
real browser driven as a bot (browserbot)	5 / 5 (100%)	5 / 5 (100%)	*NOT RUN*	5 / 5 (100%)
real-browser baseline	5 / 5 (100%)	5 / 5 (100%)	*NOT RUN*	5 / 5 (100%)

DEFENCE IN THE PATH

No defence was in the path for these runs. Stages marked NOT RUN did not happen, so they are not a result.

ANALYST FINDINGS

Written by the analyst who ran the test. This section is not generated.

This sample is the undefended baseline. The lab flow had no bot defence in front of it, so every tier, from a plain script to a browser driven as a bot, completed sign-up 5 times out of 5. That is the expected result, and it is the reference point a real audit starts from.

In a client audit, the same tiers run against your staging flow with your defence switched on. This page then shows, tier by tier, where the defence stops bots and where it stops being effective, and whether it turns away real-browser sessions.

Recommendations in a client report are tied to a figure on this page, for example:

- if header-matched clients get through but plain scripts do not, the defence is relying on header checks that are cheap to copy;
- if real-browser sessions are blocked, the defence is costing real customers and the threshold needs tuning before the event.

How the test was run

ATTACK CLASS (OWASP AUTOMATED THREATS)

- **OAT-019 Account Creation.** Flow sign-up (inferred from the flow: sign-up: account creation).

Ids and names are the OWASP Automated Threats to Web Applications handbook v1.2 entries. A flow with no declared or inferred entry reads "unmapped". The mapping classifies the flow; it does not measure it.

TIERS

- **plain HTTP client (tier 0b).** A plain HTTP library with its own TLS stack and header set. The naive control.
- **Chrome-impersonating HTTP client (tier 0).** An HTTP client that impersonates Chrome's TLS and HTTP/2 stack. A competent scripted client.
- **header-matched HTTP client (stealth).** An HTTP client that sends a browser header set measured from the browser baseline.
- **real browser driven as a bot (browserbot).** A real browser engine driven by a script the way a bot drives one: few requests, no dwell.
- **real-browser baseline.** A real browser engine driven through the same task by a script. A baseline for false positives, not a panel of real visitors.

BUDGET AND SAMPLE

Each tier had a fixed unit budget set in its run manifest, and the test sent no more than that. The table gives the units configured and the units the manifest records as sent. 5 sealed run(s) were exported; 0 were excluded as incomplete.

TIER	CONFIGURED	SENT (MANIFEST)	ROWS CAPTURED	LAB GATE
0	5	5	10	none
0b	5	10	10	none
browserbot	5	10	10	none
real-browser baseline	5	10	230	none
stealth	5	10	10	none

AUTHORISATION

Engagement reference: **not supplied**. Target, hosts, window and stop contact are set out in the signed authorisation letter for this engagement. Nothing ran outside it. The test ran in Somer North's in-house test lab (Kessler), against a staging copy of the flow.

Limits, and how to check the evidence

The bundle is the source of truth. This report was rendered from it, and the two checks below travel with it.

BUNDLE FILE	`kessler-evidence.json`
CHECKSUM (SHA-256)	`8b8a01e7aa79da62d57196a9a57cca446a5da942cc11c2a5d510ac0eafe8637d`
KEY FINGERPRINT	`93d9c7fcc3e27070`
VERDICT	PASS on both checks

WHAT THE TWO CHECKS MEAN

- **Checksum** detects accidental damage: a truncated download, a bad copy. It does not prove who made the file, because anyone who edits it can recompute the digest.
- **Signature** (Ed25519) proves who produced the file *if you trust the key*. The key is not pinned in this report. Compare the fingerprint above with the one the operator gives you, by a channel other than this file.

WHAT THE FUNNEL COLUMNS MEAN

- **Offered:** units configured for the tier in the run manifest.
- **Reached the target:** sessions with at least one request the target answered.
- **Challenged at the gate:** sessions the defence issued a challenge to. NOT RUN where no defence was in the path.
- **Completed the flow:** sessions whose success response is in the capture (for sign-up, a 201 from the registration request).

CHALLENGE DETECTION

The rows record request headers, status codes and the lab gate value, but not response headers or bodies. Challenge detection therefore rests on the status codes and the gate value. A provider challenge header cannot be read from this bundle. The rows in analysis hold 0 responses with status 403, 429 or 503.

LIMITS

- Five units per tier: a small sample, not a statistical result.
- Signed with a sample key for this lab example, not an operator key; the key is not pinned.

STEPS

1. Open the evidence verifier at somernorth.co.uk/verify (no install; it runs in your browser and makes no network requests), choose the bundle, and paste the fingerprint above into the key box.
2. Or run the command-line verifier from the evidence kit on kessler-evidence.json.

APPENDIX

No engineering evidence report was supplied with this bundle.