PRIVACYSCAN

User Manual

Website privacy, data-protection and security auditing for NDPA, GAID 2025 and GDPR.

Complete reference — v1.0 · Updated 2026-07-26

Contents

1. Welcome & overview

PrivacyScan audits any public website for privacy, data-protection and security compliance in under a minute. It crawls sampled pages non-intrusively, evaluates cookie consent, privacy policy quality, HTTPS/TLS posture, security headers, third-party trackers, subject-rights surfaces and personally identifiable information (PII) exposure, then returns a scored dashboard with citations, evidence excerpts and remediation guidance.

REGULATIONS COVERED

- Nigeria Data Protection Act 2023 (NDPA)
- Nigeria General Application and Implementation Directive 2025 (GAID)
- EU General Data Protection Regulation (GDPR)
- Advisory best-practice: WCAG 2.1 AA accessibility, HSTS preload eligibility, SPF / DMARC / CAA hardening.

SCAN TYPES

Web scan	Live crawl of a URL. Respects robots.txt. Never bypasses authentication.
Static scan	Upload source code (single file or bundle text) for pattern-based checks — API keys, tokens, SDK misconfigurations.
Scheduled scan	Recurring web scan on a fixed cadence. Pro+ plans.
Regression suite	Admin-only run across a fixed corpus to track detector accuracy over time.

INFO
Findings that require authorised source code, mobile app packages or runtime access are labelled Not Verifiable (NV) rather than pass or fail. Upload artefacts on Static Analysis to complete them.

2. Getting started

CREATE AN ACCOUNT

1. Open the landing page and click Start free — sign in.
2. Choose Continue with Google, or enter an email + password (minimum 8 characters; the profile page enforces 12+ later).
3. You are placed in a personal workspace automatically. You can create additional workspaces from the Workspace screen.

ENABLE TWO-FACTOR AUTHENTICATION (RECOMMENDED)

Open Profile ⁉ Two-factor authentication and choose Enable authenticator. Sca
Authenticator, 1Password or Authy, then enter the 6-digit code to activate. On future sign-ins, you will be prompted for the code after your password.

CHOOSE YOUR LANGUAGE

PrivacyScan ships with English, French and Spanish interface translations. Sv
Language. Scan findings themselves are always emitted in English.

WORKSPACE BASICS

A workspace groups scans, findings, teammates and billing. Every user has one default personal workspace. Owners can invite additional members and admins by email — see Section 12.

3. Running a scan

START A WEB SCAN

1. Navigate to /scan from the bottom nav.
2. Paste a full URL (e.g. https://example.com). If you omit the scheme, https:// is added automatically.
3. Press Scan. The engine resolves DNS, refuses private / loopback / link-local ranges, then begins the audit.

LIVE PROGRESS

The scan streams via Server-Sent Events (SSE). You will see live check names as they run, a rolling live Score Index, and a trend sparkline comparing against your last scan of the same hostname. If the stream stalls for more than 30 seconds, PrivacyScan automatically reconnects using full-jitter expon 12 attempts).

DEVICE PROFILE

PrivacyScan detects the device you initiated the scan from (mobile, tablet or desktop) and mirrors that viewport + User-Agent server-side so the target site serves you the same layout a real user would receive. The device profile is stored with the scan and shown as a badge on the results screen.

CANCEL OR LEAVE

Closing the tab does not abort the scan — results continue to accumulate server-side and appear under History when complete. Reopen /scan and it will resume streaming from the current position.

WARN

PrivacyScan will never bypass authentication, submit forms, or ignore robots.txt. Sites that hard-block bots may return partial results — those checks are marked NV.

4. Reading the results

SCORE INDEX & GRADE

The Score Index is a 0–100 number with an accompanying letter grade A–F. It is computed from every check outcome using severity- and category-weighted deductions, with a per-category penalty cap (max 12 points off any single category), a global deduction ceiling (max 85 total), a warn credit of 70%, and a gentle 10% curve to reward broad partial coverage. Realistic well-maintained sites typically land between 65 and 90.

A (90–100)	Excellent. Strong compliance posture across categories.
B (80–89)	Good. Minor gaps or advisory items only.
C (70–79)	Acceptable. Address highest-severity fails first.
D (60–69)	At risk. Multiple category failures.
F (< 60)	Failing. Immediate remediation required.

STATUS CODES

Pass	Check verified compliant on evidence collected.
Warn	Partial credit — implementation present but weak, or advisory best-practice missed.
Fail	Non-compliant. Deduction applied per severity.
NV	Not verifiable via public scan — requires source or runtime access.

EVIDENCE

Every finding shows the exact affected URL and an evidence snippet — the raw HTML, header value, cookie name, DOM excerpt or policy text that triggered the result. Click a finding row to expand full evidence, remediation copy and regulator citations.

5. Findings deep dive

CATEGORIES

Consent	CMP presence, reject-all parity, no pre-checked boxes, no non-essential cookies before consent, Google Consent Mode v2 signals.
Cookies	Full inventory with names, domains, expiry, HttpOnly / Secure / SameSite flags, purposes disclosed in the cookie policy.
DPO / Privacy policy	Contact details, lawful basis, retention, subject rights, transfers, complaint mechanism.
PII exposure	Emails, phone numbers, API keys, tokens, JWTs, private-key material found in HTML / JS bundles.
Security headers	HSTS, CSP, X-Frame-Options, Referrer-Policy, Permissions-Policy, X-Content-Type-Options.
TLS	Protocol version, cipher suite quality, certificate validity, mixed content.
Trackers	24-vendor signature set — analytics, ads, fingerprinting, session replay.
Rights	DSAR / subject-request surface: opt-out form, contact email, self-service portal.
Accessibility	WCAG 2.1 AA sample: contrast, alt text, form labels, landmark structure.
Infrastructure	SPF, DMARC, CAA, HSTS preload eligibility, security.txt.

SEVERITY MODEL

- High — regulator-cited requirement. Missing a DPO contact, no cookie banner, transmitted PII in cleartext.
- Medium — significant risk or partial compliance. Weak reject-all UX, missing subject-rights surface.
- Low — advisory best-practice. HSTS preload not eligible, missing security.txt.

IGNORE / ACCEPT-RISK

From any finding row, choose Ignore!’ provide a justification (min 12 characte are excluded from Score Index recomputation and are visible to teammates for audit. Un-ignore from the same menu.

ASSIGN REMEDIATION

Click Assign on a finding and pick a workspace member. They receive a notification and see the item on their Assigned to me screen. Assignments are re-evaluated on the next scan — if the underlying check flips to pass, the assignment auto-closes.

6. DPO & privacy policy checks

PrivacyScan fetches every candidate privacy / cookie policy page linked from the target and looks for the specific disclosures required by NDPA §35, GAID Article 27 and GDPR Articles 13–14.

WHAT IS INSPECTED

- Data Protection Officer name, email or contact form URL.
- Lawful basis statement for each processing purpose.
- Retention periods (specific durations, not just 'as long as necessary').
- Data subject rights list with a working request mechanism.
- International transfer disclosures (recipients, countries, safeguards).
- Regulator complaint mechanism (NDPC / relevant DPA contact).

TIP

Placeholder emails such as dpo@example.com or hello@yourdomain.com are automatically filtered — they do not count as a valid DPO contact.

7. Cookies & consent

CMP DETECTION

PrivacyScan detects OneTrust, Cookiebot, TrustArc, Osano, CookieYes, Termly, Iubenda, Complianz, Didomi, Quantcast Choice, Klaro, Usercentrics, tarteaucitron and custom banners via DOM heuristics. Detection works across iframes and SPA bundles.

PURPOSE DISCLOSURE

The engine accepts three layout patterns for cookie-purpose disclosures:

- Prose form — 'Purpose / why we use / used to … cookie / tracker' within a single sentence.
- Heading form — Category heading (Necessary / Analytics / Marketing Cookies) followed within ~500 chars by 'Purpose:'.
- Table form — 'Purpose:' or 'Used for:' label alongside a nearby cookie name or category token.

CONSENT MODE V2

If Google Consent Mode v2 signals are present, PrivacyScan checks that ad_storage and analytics_storage default to 'denied' before consent and update correctly after Accept. Missing update signals earn a warn.

8. PII detection & filters

PrivacyScan scans every sampled response body for personal or sensitive tokens using a curated pattern set (emails, phone numbers, Nigerian NIN / BVN, US SSN, AWS / Google / Stripe API keys, JWTs, RSA and EC private keys, .env leaks).

FALSE-POSITIVE FILTERS

Common placeholders (example.com emails, all-nines phone numbers, xxxxx.xxxxx.xxxxx JWTs) are excluded by default. You can add or override filter patterns per workspace under /pii-filters.

- Type — email, phone, apikey, jwt, generic.
- Pattern — literal string or JavaScript regex (validated on save).
- Note — reason for the exclusion, shown in audit logs.

WARN
Overly broad patterns can silently mask real leaks. Prefer literal strings or narrow regexes with anchors.

9. Static source scan

The Static screen (/static) accepts up to 500 KB of source per submission — a single file, a concatenated bundle, or an .env-style dump. It runs a curated pattern set for hard-coded secrets, misconfigured SDK calls, insecure storage APIs and known vulnerable library versions.

SUPPORTED CHECKS

- Secret exposure: AWS keys, Stripe keys, Google service accounts, GitHub tokens, JWT secrets, database URIs.
- Insecure storage: localStorage / SharedPreferences with sensitive keys, plaintext credential caching.
- SDK misuse: Firebase debug flags in production, verbose logging shipped, disabled certificate pinning.
- Dependency signals: known-vulnerable version strings for popular libs.

INFO

Static scans consume the same monthly scan quota as web scans.

10. History, search & tags

Every completed scan appears under /history. Use the search box to filter by hostname substring, tag, score band or scan type.

TAGS

Add free-text tags to any scan from its detail page (e.g. 'client-abc', 'q3-audit'). Tags are private to your workspace.

DIFF VIEW

Open any scan and click Compare to previous. PrivacyScan shows: score delta, new findings, resolved findings, and evidence changes. Diff URLs are shareable within the workspace.

11. Sharing & trust

SHARE LINKS

Generate a public read-only link from any scan. Choose a TTL — 1 hour, 24 hours, 7 days or 30 days — and an optional label. Recipients see a sanitised report (findings, evidence, remediation) without your workspace, teammates or billing data. Revoke at any time; view counts and last-viewed timestamps are tracked.

VERIFIED-SCAN PROOF

Every completed scan is signed with HMAC-SHA256 over a canonical string of its metadata + findings digest. The signature and hash are printed on the PDF report and can be verified independently at /verify by pasting the scan ID or the report file.

STATUS BADGE

Embed a live SVG badge on your own site: /api/public/badge/{token}.svg. Colour and grade update automatically after each scheduled scan. Copy the exact snippet from the Share dialog.

DOMAIN VERIFICATION

Under /domains, add a hostname and choose a DNS TXT record or an HTML <meta> tag. PrivacyScan re-checks on demand and stores a verified-since timestamp. Verified hostnames unlock deeper checks (mail authentication, HSTS preload confirmation).

12. Teams & workspaces

ROLES

Owner	Full control including billing, workspace deletion, member removal.
Admin	Invite / remove members, manage schedules, PII filters, regression suite.
Member	Run scans, view all findings, comment, assign.
Viewer	Read-only. Cannot run scans or modify findings.

INVITES

From /workspace, click Invite. Enter an email, choose a role, optionally include a message. PrivacyScan sends an invite email via Resend containing a single-use link valid for 7 days. Every invite (created, sent, accepted, revoked) is recorded to the audit trail under Workspace!' Invite activity.

INFO

Row-Level Security enforces workspace isolation on every read. A member cannot see scans, findings or comments outside workspaces they belong to, even by manipulating IDs.

13. Collaboration

THREADED COMMENTS

Every finding has its own comment thread. @mention a teammate to notify them. Threads carry across scans of the same hostname when the same check ID recurs.

ASSIGNMENTS

See Section 5 — Assign remediation. Assigned tasks appear on the assignee's /assignments screen with due-date, status and a link back to the finding.

14. Scheduled scans

Pro+ plans can schedule a hostname to be re-scanned automatically. Choose a cadence: daily, weekly or monthly. Runs happen from a pg_cron job and consume monthly scan quota just like manual runs.

REGRESSION ALERTS

When a scheduled scan drops the Score Index by 10+ points, or any check flip receive an in-app notification linking directly to the offending scan.

15. Regression suite (admin)

The /regression screen (admin-only) runs a curated corpus of public sites to track detector accuracy over time. It records per-site score deltas, finding digests and pass↔fail flips per run.

MANAGING THE CORPUS

- Add or disable sites from the corpus table.
- Trigger a run — sites are scanned serially, max 8 per run to protect targets.
- Review history: per-site trend chart, delta scores, flipped checks.

The same threshold (10-point drop or any pass↔fail flip) triggers admin notific

16. Reporting & exports

PDF REPORT

From any scan, click Export ’ PDF. The report contains: cover page with hostn executive summary, full findings table grouped by category, evidence excerpts and remediation. Signed with the scan hash for later verification.

CSV EXPORT

Export ’ CSV emits one row per finding with check ID, status, severity, catego and evidence summary. Suitable for import into GRC tools.

REGULATOR-MAPPED FILTERS

Filter the findings table by NDPA, GAID or GDPR to produce a framework-specific view before exporting.

HIGH-RESOLUTION DASHBOARD PNG

Export ’ PNG produces a 4320×5400 image of the scan detail view suitable for

17. Billing & plans

Free	3 scans / month, 1 workspace, no scheduling, single user.
Pro — ¦3,000/mo	50 scans / month, scheduled scans, PDF/CSV export, single user.
Team — ¦12,000/mo	250 scans / month, up to 10 seats, full collaboration, PII filters.
Business — ¦47,000/mo	Unlimited scans, unlimited seats, API access, regression suite, priority support.

Billing is handled by Paddle. Manage subscription, invoices and cards from /billing. Downgrading takes effect at the end of the current billing period.

QUOTA SEMANTICS

Scan quota resets on the 1st of each calendar month at 00:00 UTC. Failed scans do not consume quota; cancelled scans do.

18. Account & security

PROFILE

- Display name — shown on comments, assignments and audit logs.
- Email address — must be reconfirmed if changed.
- Password change — enforced strength meter (12+ chars, mixed case, digits, symbols; blocks common passwords and email-derived strings).

MFA (TOTP)

See Section 2. Removing your factor requires the current password. Losing access to your authenticator means contacting support with proof of account ownership.

ACCOUNT DELETION

Profile !’ Delete account. Type DELETE to confirm. Removes your user record, workspaces where you are the sole owner, and comments authored by you. Cannot be undone.

DATA RETENTION

Scans and findings are retained indefinitely while the account is active. Share-link view logs are purged after 90 days. Backups are encrypted at rest.

19. API & webhooks (Business)

PUBLIC ENDPOINTS

GET /api/public/badge/{token}.svg	Live SVG badge with current score & grade.
GET /api/public/verify/{scanId}	Machine-readable verification of the report signature.
POST /api/public/csp-report	CSP report ingestion endpoint (attach as report-uri on your site).
POST /api/public/hooks/run-scheduled-scans	Cron trigger for scheduled scans (called by pg_cron).

WARN
All public endpoints implement request validation and (where applicable) signature verification. Never expose your workspace secret to a client.

20. Privacy & compliance of PrivacyScan itself

- SSRF guard — every scan target is DNS-resolved server-side and rejected if it lands on a private, loopback or link-local range. Redirects are re-validated.
- Respects robots.txt — Disallow rules are honoured; a partial scan is preferred over an unauthorised one.
- No authentication bypass — PrivacyScan never submits credentials, cookies you did not provide, or bypasses login walls.
- Private history — scan history is scoped to your workspace via Row-Level Security. Public share links are opt-in per scan with TTL and revocation.
- Signed reports — every completed scan carries an HMAC-SHA256 signature verifiable at /verify.

21. Troubleshooting & FAQ

THE SSE STREAM STALLED

PrivacyScan reconnects automatically after a 30-second silence, up to 12 attempts with exponential backoff. If the scan never completes, reload /history — most scans finish server-side even without an open connection.

A SCAN FAILED WITH 'INVALID URL'

Check that the URL parses. Bare hostnames without a scheme are prefixed with https:// automatically. IP addresses in private ranges are blocked by design.

INVITE EMAIL DID NOT ARRIVE

Ask the recipient to check spam. From /workspace!' Invite activity, verify the
Invites are single-use and expire after 7 days.

I LOST MY AUTHENTICATOR

Contact support@privacyscan with proof of account ownership (billing email + a signed report you generated). Recovery is manual and can take up to 48 hours.

SHARE LINK SAYS 'EXPIRED' OR 'REVOKED'

Regenerate a fresh link with a longer TTL from the source scan's Share dialog.

SCORE DROPPED AFTER A SMALL CHANGE

Open the Compare view against the previous scan. The diff highlights exactly which checks flipped and why — usually a header removed, a cookie added before consent, or a policy page edited.

22. Glossary

Score Index	0–100 weighted compliance score. See Section 4.
Grade	A–F letter grade derived from the Score Index.
NV	Not verifiable — a check that cannot be answered from public scanning alone.
CMP	Consent Management Platform — the cookie banner + preference store.
DSAR	Data Subject Access Request — a user asking to see / delete / port their data.
DPO	Data Protection Officer — statutory contact for privacy matters.
HSTS	HTTP Strict Transport Security header enforcing HTTPS.
HSTS preload	Baked-in browser list of HSTS-enforced hosts. Eligibility has strict criteria.
SPF / DMARC	DNS records that authenticate outbound email and prevent spoofing.
CAA	DNS record restricting which CAs may issue certificates for the domain.
Consent Mode v2	Google's signalling protocol for pre/post-consent tag behaviour.
RLS	Row-Level Security — Postgres policies that enforce per-row access.

23. Support & feedback

Email support@privacyscan for account issues, bug reports and detector feedback. Feature requests are logged in the public changelog at /changelog. When reporting an issue, include your scan ID (or the report signature) so we can reproduce the exact evidence you saw.

TIP
The fastest way to improve PrivacyScan is to send us a URL where a check produced the wrong outcome. We add it to the regression corpus and use it to prevent future regressions.