Information Security Whitepaper

Dark Web Metrics — Security, Privacy, and Operational Controls

Dark Web Metrics — [LEGAL ENTITY]

May 2026 — Version 1.0

Contents

1 1. Purpose

This whitepaper describes the security and privacy controls Dark Web Metrics ("DWM", operated by [LEGAL ENTITY], a [JURISDICTION] limited liability company) implements to protect the confidentiality, integrity, and availability of customer information and the data products derived from public and licensed sources.

It is intended for:

- Prospective and active customers conducting vendor security assessments;
- Customers' internal security, privacy, and procurement teams;
- Auditors operating under NDA on behalf of customers;
- Underwriters of customers' cyber-liability or E&O policies who may review vendor controls.

This document is **not** a SOC 2 report. DWM intends to pursue SOC 2 Type I in the twelve (12) months following its first commercial revenue and SOC 2 Type II thereafter. This whitepaper describes the controls that will be evaluated under that audit.

2 2. About Dark Web Metrics

DWM produces decision-grade exposure indices from public and licensed sources covering torrent networks, public Telegram channels, public forum content, breach feeds (under license), and certificate-transparency logs. Outputs are delivered to customers as a web dashboard, machine-readable API endpoints, scheduled PDF reports, and signed evidence packages.

DWM does **not** offer offensive-security services, penetration testing, take-down execution, threat-actor engagement, or any service that requires authentication bypass on third-party systems. Every collection method described in this document is either (a) observation of public data sources, or (b) consumption of data made available to DWM under a current commercial license from a named provider.

3 3. Governance and roles

Role	Responsibilities
Founder / CEO	Ultimate owner of the security and privacy program.
Security Lead (acting: Founder)	Owns the controls described in this document; reviews quarterly.
Privacy Officer (acting: Founder)	Owns data subject rights handling, DPIA, and lawful-basis review.
Engineering	Implements technical controls; conducts code review.
External Counsel	Reviews contracts, data flows, and incident response; on retainer.
External Auditor (future)	Performs SOC 2 examinations and pen tests.

Until DWM's headcount reaches a size that makes role separation operationally meaningful (estimated 6+ FTEs), the Founder holds the Security Lead and Privacy Officer

roles concurrently. This concentration is disclosed; customer review is expected during procurement.

The Security Program is reviewed at least annually and after any material change in service, infrastructure, or applicable law.

4 4. Asset management

4.1 4.1 Asset inventory

DWM maintains an up-to-date inventory of:

- **Information assets**: customer-owned data (catalog uploads, configuration), DWM-owned datasets (signal records, computed indices, derived metrics), and operational data (logs, telemetry).
- **Software assets**: production code repositories, infrastructure-as-code definitions, third-party libraries and their pinned versions.
- **Hardware assets**: managed cloud accounts (Supabase, Vercel) and self-managed virtual servers (Hetzner) hosting collector workloads. No physical assets are operated by DWM.

Assets are classified by sensitivity:

Class	Description	Examples
Public	May be disclosed freely.	Marketing materials, this whitepaper.
Internal	Operational data; default class.	Internal runbooks, source-list configurations.
Confidential	Customer-derived or customer-identifying.	Catalog uploads, computed indices, customer contacts.
Restricted	High-impact secrets and PII.	Signing keys, API credentials, plaintext PII pending redaction, tenant salts.

4.2 4.2 Data flow

Customer data and externally-collected data are processed in segregated logical pipelines. Customer-supplied catalog data is uploaded over TLS, stored in a tenant-isolated Postgres schema, and exposed only via authenticated APIs subject to row-level security (RLS).

Externally-collected data is ingested by isolated collector processes (Section 7.2), normalized, hashed where it contains PII, and stored in DWM-controlled schemas with no inbound customer access except through derived indices.

5 5. Access control

5.1 5.1 Identity

Customer access to the DWM application is authenticated via Supabase Auth using one or more of: passwordless email magic links, single sign-on (SSO) via the customer's IdP (Fortress tier), and OAuth providers where contractually agreed. Multi-factor authentication (MFA) using TOTP is mandatory for all administrative and Fortress-tier user accounts.

Internal DWM access to production infrastructure is authenticated via the operating provider's IdP (Hetzner, Supabase, Vercel, GitHub) and requires hardware-backed MFA (FIDO2 / WebAuthn) for the founder and any future engineering staff.

5.2 5.2 Authorization

All customer-facing data tables enforce Postgres row-level security keyed on tenant identifier. Internal collector and admin tables are not exposed to any customer-facing JWT.

Service accounts used by collectors and Edge Functions use Supabase service-role credentials with table-level grants scoped to the minimum necessary. Service-role credentials are rotated quarterly or upon staff change.

5.3 5.3 Privileged access

Privileged operations — production database migrations, secret rotations, customer data deletion — require dual control where personnel allow. While DWM operates with a single founder, dual control is approximated by automated peer review (every change is committed to a Git repository, signed, and CI-validated; emergency runbook deviations are logged with after-the-fact review).

5.4 5.4 Joiner / mover / leaver

A documented JML procedure governs onboarding, role-change, and offboarding events. Off-boarding triggers immediate revocation of all production credentials, rotation of any shared secrets touched, and an access-review attestation. While DWM has no employees other than the founder, the procedure is in place for future hires and contractors.

6 6. Cryptography

6.1 6.1 Encryption in transit

All HTTP traffic to customer-facing endpoints is served exclusively over TLS 1.2 or higher, with TLS 1.3 preferred. HSTS is enabled with `max-age=31536000; includeSubDomains; preload`. Cipher suites and certificates are reviewed twice per year against the latest Mozilla recommendations.

Internal traffic between collectors and Supabase, and between Edge Functions and Supabase, occurs over TLS-encrypted Postgres connections (`sslmode=verify-full`).

6.2 6.2 Encryption at rest

Customer data and DWM-managed data stored in Supabase Postgres are encrypted at rest using AES-256 (Supabase-managed encryption). Object storage (PDF reports, evidence packages) is encrypted at rest in Supabase Storage. Collector-host filesystems on Hetzner are LUKS-encrypted; collectors persist no customer-identifying data to disk by default.

6.3 6.3 Hashing of PII

Email addresses observed in any pipeline (stealer logs, breach feeds, public posts) are hashed at ingest using SHA-256 with a per-tenant rotating salt. Salts are rotated annually or on demand following a security incident. Plaintext emails are retained no longer than thirty (30) days and only when required for customer-side notification obligations under applicable law (e.g., GDPR Art. 34).

6.4 6.4 Signing keys

Evidence packages are signed with an Ed25519 keypair held in a hardware-backed key store. The public key is published at `https://darkwebmetrics.com/.well-known/evidence-pubkey.txt` so that any party may verify the signature without contacting DWM. Key rotation occurs annually with an overlap window during which both the old and new public keys are published.

6.5 6.5 Secrets management

Production secrets (API keys, salts, signing keys) are stored in the secret stores native to each provider (Supabase Vault, Vercel Environment Variables, Hetzner Cloud Secrets) and never committed to source control. Pre-commit hooks and CI scan for credential patterns. Discovered credentials are rotated immediately and the discovery is logged.

7 7. Operations security

7.1 7.1 Production architecture

The production stack is:

- **Frontend**: Next.js (App Router) deployed on Vercel.
- **API and database**: Supabase (Postgres, Auth, Storage, Edge Functions).
- **Collectors**: Node 22 processes on Hetzner virtual servers, geographically distributed across at least two regions.
- **PDF generation**: Headless Chromium running on a dedicated Hetzner host.
- **Email delivery**: Resend (or equivalent reputable provider).
- **Observability**: Better Stack (logs), Grafana Cloud (metrics), Sentry (errors).
- **Source control and CI**: GitHub.

A single-region failure of either Vercel or Supabase results in service degradation; both providers operate multi-region active-active for their respective services on DWM's plans.

7.2 7.2 Isolation

Each collector class (DHT, indexer, Telegram, forum, licensed feeds, DNS) runs as an isolated process with its own credentials, its own egress network controls, and no shared state at the operating-system level. A compromise of one collector does not by design grant access to another collector's credentials, to customer data, or to the Supabase service-role secret.

7.3 7.3 Logging and monitoring

All production services emit structured logs to Better Stack. Logs are retained for at least one (1) year in hot storage and longer in cold storage where required by contractual or regulatory obligation. Logs do not contain plaintext credentials, full PII, or signing-key material.

Metrics covering request latency, error rate, queue depth, collector throughput, and pipeline lag are visible on internal dashboards. Alert thresholds are calibrated against historical baselines; alerts are routed to the on-call rotation (currently: founder).

7.4 7.4 Change management

All production changes are introduced through pull requests against the main branch of the source repository. CI runs typing, linting, unit tests, and dependency vulnerability scans. Merges to main trigger automatic deployment to a staging environment; production deployment requires explicit promotion.

Database schema changes are introduced through versioned migrations applied via Supabase's migration tooling. Migrations are run against staging first and then against production with an explicit confirmation step.

7.5 7.5 Backups

Supabase performs daily encrypted backups of the production database with point-in-time recovery (PITR) over the trailing seven (7) days. Backups are retained for thirty (30) days. Backup restoration is tested quarterly.

7.6 7.6 Vulnerability management

Third-party dependencies are scanned by GitHub Dependabot and `npm audit` in CI. Critical and high-severity vulnerabilities are patched within seven (7) and thirty (30) days respectively, where a patch is available. When no patch is available, mitigations are documented and tracked to resolution.

Infrastructure components (operating system, runtime versions on collectors) are updated monthly or upon a critical CVE, whichever is sooner. Automated reboots for kernel security patches are enabled.

8 8. Communications security

8.1 8.1 Network

Customer-facing services are reachable only via TLS on port 443. Collectors initiate outbound connections to Supabase, to source endpoints, and to monitoring services;

inbound connections to collectors are restricted to administrative SSH from a small allowlist of operator IP ranges with key-only authentication and Fail2ban.

8.2 8.2 Email

Email originating from DWM domains is sent through SPF-, DKIM-, and DMARC-aligned configurations with a strict `p=reject` DMARC policy. Inbound email is routed through a hosted provider with spam and phishing filtering enabled.

9 9. System acquisition, development, and maintenance

9.1 9.1 Secure development lifecycle

All production code is held in a private GitHub repository. Commits are signed (where supported by the contributor's environment). Pull requests require at least one reviewer; in the single-founder phase, the founder reviews their own changes against a documented checklist that includes security implications.

Threat modeling (STRIDE-based) is performed for each new feature that introduces a new data flow, a new authentication surface, or a new third-party integration. The threat model is committed to the repository alongside the feature.

9.2 9.2 Third-party libraries

Production runtime dependencies are pinned by version and hash. New dependencies undergo a lightweight review (publisher reputation, maintenance status, license compatibility) before adoption. License inventory is reviewed annually.

9.3 9.3 Penetration testing

DWM will commission an annual external penetration test by a reputable third party once it exceeds the threshold of 25 active paying customers or once it pursues SOC 2 Type I, whichever occurs first. Findings are tracked to resolution; an executive summary is made available to customers under NDA.

10 10. Supplier and sub-processor management

10.1 10.1 Selection

Suppliers handling customer data or DWM-confidential data are selected based on (a) security and privacy posture, including their own attestations (SOC 2, ISO 27001, or equivalent), (b) data-residency and transfer mechanisms, and (c) contractual willingness to enter into a Data Processing Agreement with adequate sub-processor and incident terms.

10.2 10.2 Sub-processor list

The current list of sub-processors is maintained as a separate document and published at `https://darkwebmetrics.com/trust/subprocessors`. Customers are notified in writing of any change to the list at least fifteen (15) days before the change takes effect, with a right to object for material grounds.

10.3 10.3 Ongoing oversight

Each sub-processor's continued use is reviewed at least annually. Material incidents (sub-processor breaches, sub-processor changes of control) trigger an interim review.

11 11. Information security incident management

11.1 11.1 Incident definition

An incident is any event that has, or is reasonably suspected to have, caused (a) unauthorized access to or disclosure of customer data or DWM-confidential data, (b) unauthorized modification of indices or reports, (c) extended service unavailability, or (d) violation of DWM's stated source-collection limits.

11.2 11.2 Response

The incident response plan is documented separately. Summary lifecycle:

1. **Detection.** Via monitoring, customer report, sub-processor notification, or external researcher disclosure.
2. **Triage.** Incident lead (Founder/Security Lead) classifies severity (P0–P3) within one (1) hour of detection.
3. **Containment.** Immediate action to stop the bleeding; documented and timestamped.
4. **Eradication.** Root cause identified and removed.
5. **Recovery.** Service restored; integrity validated.
6. **Notification.** Affected customers notified without undue delay and in any case within seventy-two (72) hours for incidents reasonably likely to affect them. Regulatory notifications made as required by applicable law.
7. **Post-mortem.** Within fifteen (15) business days, a written post-mortem is produced and shared with affected customers.

11.3 11.3 Customer cooperation

DWM cooperates in good faith with affected customers' own incident response, including providing forensic data, indicators, and timelines, subject to the protection of other customers' data and DWM's legal obligations.

11.4 11.4 Disclosure to researchers

DWM operates a coordinated disclosure policy. Researchers may report vulnerabilities to `security@darkwebmetrics.com` or via a published `.well-known/security.txt`. DWM commits to acknowledging reports within seventy-two (72) hours and to good-faith handling. No legal action will be taken against researchers acting in good faith and within the scope described in the disclosure policy.

12 12. Business continuity and disaster recovery

12.1 12.1 Service tier

DWM's contractual availability target (SLA) is 99.5% calculated monthly for the customer-facing dashboard, with separate availability targets for asynchronous deliverables (board decks, alerts) specified in the SLA document.

12.2 12.2 Recovery objectives

- **Recovery Time Objective (RTO).** Twenty-four (24) hours for customer-facing dashboard restoration from a complete provider outage.
- **Recovery Point Objective (RPO).** Twenty-four (24) hours for non-real-time data; one (1) hour for the most recent computed indices, supported by Supabase point-in-time recovery.

12.3 12.3 Provider redundancy

Where service-tier and cost allow, DWM maintains exportable representations of critical data (computed indices, configuration, customer catalogs) that could be restored against an alternate provider in the event of prolonged Supabase or Vercel unavailability. This is a degraded-mode fallback, not a hot standby.

12.4 12.4 Tabletop exercises

Disaster-recovery scenarios are exercised at least annually. Findings inform the BCP and IRP.

13 13. Privacy and compliance

13.1 13.1 Applicable regimes

DWM operates as a service provider serving customers globally, including the United States, the European Union, the United Kingdom, Brazil, Japan, and other jurisdictions. The following regimes are relevant by default:

- **United States** — California Consumer Privacy Act / California Privacy Rights Act (CCPA/CPRA), Virginia CDPA, Colorado CPA, Connecticut CTDPA, and other state laws as they enter effect.
- **European Union and EEA** — General Data Protection Regulation (GDPR), with Standard Contractual Clauses (2021/914) for cross-border transfers.
- **United Kingdom** — UK GDPR with the International Data Transfer Addendum (IDTA) for transfers.
- **Brazil** — Lei Geral de Proteção de Dados (LGPD).
- **Japan** — Act on the Protection of Personal Information (APPI).
- **Australia, Singapore, Canada (PIPEDA)** — handled by contract where customers are present in those jurisdictions.

13.2 13.2 Roles

In the customer relationship, the customer is the **controller** and DWM is the **processor**. For DWM's own corporate processing (employees, suppliers), DWM is the **controller**.

13.3 13.3 Data minimization

DWM applies the principle of data minimization at ingest. Plaintext personal identifiers from external sources are hashed at ingest with rare exceptions (Section 6.3) and the exceptions are time-bound.

13.4 13.4 Data subject rights

Requests under GDPR Articles 15–22, CCPA $§$1798.105–130, LGPD Articles 17–22, and equivalent provisions are handled per the Data Retention & Deletion Policy. Requests directed at the customer's data are forwarded to the customer; requests directed at DWM's own controller-mode processing are handled by DWM.

13.5 13.5 CSAM handling

Any content classified as Child Sexual Abuse Material at ingest is immediately discarded from indices and from raw storage. Only the infohash plus detection timestamp is retained for the purpose of reporting to the National Center for Missing & Exploited Children (NCMEC) or equivalent reporting bodies in the applicable jurisdiction. The detailed procedure is documented in the CSAM Detection & Reporting SOP, available to customers and regulators on request.

13.6 13.6 Personnel security

All personnel with access to production systems sign a Confidentiality and Acceptable Use agreement. Background checks are performed for any future hires whose role requires direct production access, to the extent permitted by the applicable jurisdiction.

14 14. Acceptable use of DWM outputs by customers

DWM outputs are intended for defensive, due-diligence, brand-protection, anti-piracy, and risk-quantification use cases. The contractual prohibitions on customer use are specified in the Acceptable Use Policy. Examples of prohibited customer use include offensive-security campaigns against third parties, contact or harassment of individuals identified in indices, and resale of raw DWM-derived data without written authorization.

15 15. Compliance roadmap

Milestone	Target
Documented IR and BCP procedures with annual tabletop	Within 6 months of first paying customer
External penetration test	Within 12 months or 25 active customers, whichever is first
SOC 2 Type I readiness assessment	Within 12 months of first paying customer
SOC 2 Type I report	Within 18 months of first paying customer
SOC 2 Type II report	12 months after Type I
ISO 27001 (if customer demand justifies)	Decision point after SOC 2 Type II

16 16. Contact

- Security and vulnerability reports: `security@darkwebmetrics.com`

- Privacy inquiries and DSARs: `privacy@darkwebmetrics.com`
- Contract and procurement: `legal@darkwebmetrics.com`
- Methodology questions: `methodology@darkwebmetrics.com`

End of document. Information Security Whitepaper version 1.0 — May 2026.