# Methodology Document

Dark Web Metrics — How the Five Indices Are Computed

Dark Web Metrics — [LEGAL ENTITY]

May 2026 — Version 1.0

## Contents

# 1 1. Purpose and audience

This document specifies, in auditable detail, how Dark Web Metrics computes the five indices it sells: **Piracy Velocity Score (PVS)**, **Title Distribution Index (TDI)**, **Subscriber Credential Exposure (SCE)**, **Premium Account Black Market Index (PABMI)**, and **Brand Impersonation Reach (BIR)**.

The audience is threefold:

- **Customer-side analysts, security, and legal teams** who must defend the numbers in a board deck, a regulatory filing, or an underwriting renewal.
- **Independent auditors** reviewing Dark Web Metrics' service on behalf of a customer (e.g., as part of a SOC 2 readiness assessment of the customer's vendor stack).
- **Regulators or counterparties** examining how exposure metrics are derived in the event of a dispute or investigation.

Anything not specified here is **not** part of the service. If a behavior is observed in the product that is not described or implied by this document, it is either a bug or a documentation gap and should be reported to `methodology@darkwebmetrics.com`.

# 2 2. Scope and non-scope

## 2.1 2.1 In scope

- Definitions, formulas, inputs, normalization, and windowing for the five indices.
- Source taxonomy and weighting rules.
- Title-to-catalog matching procedure.
- Data exclusions, including content automatically discarded.
- Peer-comparison anonymization rules.
- Quality assurance, back-testing, drift detection, and version control of the methodology.
- Limitations and known biases.

## 2.2 2.2 Out of scope

- Pricing of the service (see Order Form).
- Service-level commitments (see SLA).
- Data processing terms (see DPA).
- Acceptable use of the outputs (see AUP).
- Specific source identities and vendor agreements (confidential; summarized in the Sub-processor List where applicable).

# 3 3. Source taxonomy

All inputs that feed the indices come from one of six source classes. Each class has a default trust weight $w_f \in [0, 1]$ used in the index formulas, calibrated quarterly.

| Class | Examples | Default trust $w_f$ | Notes |
|---|---|---|---|
| **Torrent — DHT** | Mainline DHT swarms observed via our own DHT nodes. | 0.90 | Highest signal: observation of an actual live swarm. |
| **Torrent — indexers** | 1337x, Nyaa, Bitsearch and similar public indexers. | 0.70 | Self-reported by the indexer; complements DHT. |
| **Telegram public** | Public channels (no invite-only). | 0.55 | Subject to channel inflation (forwards, bots); we discount accordingly. |
| **Public forums** | Public threads on Nulled, Cracked, archive.org snapshots of forums. | 0.60 | High signal for PABMI; lower for PVS. |
| **Licensed feeds** | DarkOwl, SpyCloud, Constella, Flare (subject to current contracts; see Sub-processor List). | 0.85 | Provider-aggregated; trust calibrated per provider. |
| **DNS / Certificate Transparency** | Certstream and equivalent public CT logs. | 0.95 for BIR only | Used only for BIR. Not a signal for PVS/TDI/SCE/PABMI. |

**Calibration.** Trust weights are reviewed quarterly. Adjustment criteria: (a) precision/recall against a held-out validation set of known-pirated releases, (b) provider reliability incidents, (c) systemic biases discovered during drift review. Calibration changes are version-controlled (see Section 11) and disclosed to active customers within fifteen (15) days of taking effect.

# 4 4. Computation windows

A **window** is the temporal unit over which an index is computed. Three granularities are supported:

| Granularity | Length | Default availability |
|---|---|---|
| `day` | 24h, UTC-aligned. | Sentinel and Fortress |
| `week` | ISO 8601 week (Monday 00:00 UTC – Sunday 23:59 UTC). | All tiers. |
| `month` | Calendar month, UTC. | All tiers. |
| `quarter` | Calendar quarter, UTC. | All tiers (board deck r |

Indices computed at a finer granularity (e.g., `day`) are deterministically aggregated to coarser granularities (`week`, `month`, `quarter`). The aggregation rule differs by index and is specified in Sections 6–10.

**Replay.** All indices are idempotent: re-running the computation on a historical window yields the same result for the same input data and the same methodology version. Methodology version pinning is enforced at the window level.

# 5 5. Title-to-catalog matching

Several indices (PVS, TDI) require matching observed strings (e.g., a torrent name like `Brazzers.24.05.10.Jane.Doe.XXX.1080p.WEBRip`) against the customer's catalog of canonical titles. Matching is performed in three layers:

## 5.1 5.1 Layer 1 — normalization

Observed strings are lowercased, stripped of diacritics, stripped of common quality and codec tags (`1080p`, `2160p`, `xxx`, `webrip`, `hdrip`, `mp4`, `mkv`, …), and tokenized into trigrams. The same normalization is applied to the customer's catalog at upload time to produce the `fingerprints` field stored against each title.

## 5.2 5.2 Layer 2 — trigram Jaccard

The Jaccard similarity between the observed trigram set and each catalog title's trigram set is computed. Matches with similarity below a configured threshold $\theta_{\text{jaccard}}$ (default `0.55`) are discarded. Matches at or above the threshold are passed to Layer 3 (or accepted directly if no Layer 3 is in scope for the tenant).

## 5.3 5.3 Layer 3 — semantic re-ranker (Fortress tier only)

For matches scoring in the ambiguous band $[0.45, 0.65]$, a sentence-embedding re-ranker (currently `gte-small`-class via managed inference) computes cosine similarity between the observed string and the catalog title. The Layer 3 score is blended with the Layer 2 score at $0.6 \cdot \text{semantic} + 0.4 \cdot \text{jaccard}$. The final acceptance threshold is $\theta_{\text{final}} = 0.62$.

## 5.4 5.4 Multi-match handling

When an observation matches more than one catalog title, the observation is attributed to **the single highest-scoring title** unless the difference between top-1 and top-2 is below 0.04, in which case the observation is split proportionally across the tied titles (50/50 in the binary case).

## 5.5 5.5 Performer- and series-level rollups (Fortress tier)

When a customer requests performer-level or series-level indices, observations attributed to titles are rolled up via the `performers` and `series` fields of the catalog. Rollups are additive (a single observation may count once at the title level and once at each parent rollup).

# 6 6. Piracy Velocity Score (PVS)

## 6.1 6.1 Definition

PVS measures how quickly a specific release accumulates pirated exposure relative to its release time. PVS is computed **per title, per window**.

## 6.2 6.2 Formula

$$PVS_{t,W} = \max\Big(0,\ \log_{10}(1 + E_{t,W}) - \alpha \cdot \log_{10}(1 + h_{t,W})\Big)$$

where:

- $E_{t,W}$ is the weighted total exposure for title $t$ in window $W$, defined below;
- $h_{t,W}$ is the elapsed hours between the title's `release_date` and the start of window $W$, floored at $0$;
- $\alpha$ is a time-decay coefficient, default $\alpha = 0.7$, calibrated per tenant tier (Section 6.4).

## 6.3 6.3 Exposure aggregation

$$E_{t,W} = \sum_{s \in S_W(t)} w_{f(s)} \cdot e(s)$$

where $S_W(t)$ is the set of signals observed within $W$ that match title $t$ per Section 5; $w_{f(s)}$ is the trust weight of the source class of signal $s$; and $e(s)$ is the per-signal exposure, defined per source class:

| Source class | $e(s)$ |
|---|---|
| Torrent DHT | `seeders + 0.3 * leechers` |
| Torrent indexers | `0.5 * (seeders + 0.3 * leechers)` (halved to avoid double-counting against DHT) |
| Telegram public | `log10(1 + views) + 2 * log10(1 + forwards)` |
| Public forums | $5 \cdot \mathbb{1}_{\text{thread present}} + \log_{10}(1 + \text{reply\_count})$ |
| Licensed feeds | Provider-normalized exposure, mapped to a 0–100 scale by an internal calibration table. |

## 6.4 6.4 Calibration of $\alpha$

Default $\alpha = 0.7$. Recalibrated quarterly using a held-out set of titles with known release dates and observed pirate uptake. The objective is for the median PVS at $h = 48$h for a "typically-pirated" tier-1 release to fall in the $[5, 7]$ band on the 0–10 reporting scale (Section 6.6).

## 6.5 6.5 Interpretation

| PVS band | Practical meaning |
|---|---|
| 0.0 – 2.0 | No or marginal pirated distribution detected. |
| 2.0 – 4.0 | Background piracy; typical for non-flagship titles >30 days old. |
| 4.0 – 6.0 | Active piracy; standard for any popular release in its first week. |

| PVS band | Practical meaning |
|---|---|
| 6.0 – 8.0 | High piracy intensity; routine for flagship releases or material with strong demand. |
| 8.0 – 10.0 | Extreme piracy intensity and/or pre-release leak; warrants P0 escalation. |

### 6.6 6.6 Reporting scale clipping

PVS is reported clipped to the $[0, 10]$ range. Values above 10 occur but are reported as `10.0` to preserve the interpretability of the scale. The unclipped value is retained in the underlying record for forensic review.

### 6.7 6.7 Aggregation across windows

When aggregating PVS from a finer to a coarser window (e.g., daily PVS into weekly PVS), the aggregation is the **per-title maximum** within the period, not the average. Rationale: a single 24h spike of viral distribution is the operationally relevant signal, not the mean.

## 7 7. Title Distribution Index (TDI)

### 7.1 7.1 Definition

TDI ranks the titles in a customer's catalog by total pirated exposure within a window. Unlike PVS, TDI is not normalized by release date; it answers "what is being pirated most this period?", not "what is being pirated fastest relative to its release?".

### 7.2 7.2 Formula

$$TDI_{t,W} = E_{t,W}$$

with $E_{t,W}$ as defined in Section 6.3.

Titles are then **ranked** in descending order of $TDI_{t,W}$. The reported metric is the rank, the raw exposure, and the normalized exposure (TDI as a percentile of the customer's catalog in $W$).

### 7.3 7.3 Top-N reporting

Standard reports surface top-10 by default, configurable up to top-100 per window. Below top-100, individual titles are not surfaced; aggregate statistics (P50, P90, P99) are reported instead, to preserve interpretability and reduce noise.

### 7.4 7.4 Long-tail decay

For multi-window TDI (e.g., quarterly), titles with exposure below the P50 of the customer's catalog in $W$ are flagged as "long tail" and excluded from the top-N visualization (but retained in the raw export).

# 8 8. Subscriber Credential Exposure (SCE)

## 8.1 8.1 Definition

SCE measures the volume of distinct credentials (email + secret of any kind) tied to the customer's monitored domains observed in stealer logs, combolists, and breach dumps within a window.

## 8.2 8.2 Formula

$$SCE_W = N_W^{\text{new}} + \beta \cdot N_W^{\text{high-sev,fresh}}$$

where:

- $N_W^{\text{new}}$ is the count of **new** unique email hashes observed in $W$ on the customer's domains (deduplicated against the rolling 180-day history of email hashes seen);
- $N_W^{\text{high-sev,fresh}}$ is the subset of new exposures where (a) the severity classification is `stealer-cookie` or `stealer-session` (active session credentials, not just email/password pairs), and (b) the capture freshness is $\leq 7$ days;
- $\beta = 2.5$ is the multiplier for high-severity, fresh exposures, reflecting the substantially higher account takeover risk.

## 8.3 8.3 Severity classification

| Class | Definition | Severity |
|---|---|---|
| `combolist` | Email + password observed in a combolist (recombined from prior breaches). | Low |
| `breach` | Email + password from a discrete identified breach (not the customer's own). | Medium |
| `stealer-cookie` | Stealer log containing cookies for the customer's domain. | High |
| `stealer-session` | Stealer log containing active session tokens / authenticated headers for the customer's domain. | Critical |

Severity is assigned at ingest by a deterministic classifier described in the internal `severity_classification.md` (available to customers under NDA on request).

## 8.4 8.4 Privacy treatment

All emails are hashed at ingest with SHA-256 plus a per-tenant rotating salt. Plaintext emails are retained no longer than thirty (30) days and only when explicitly required for the customer to fulfil a notification obligation under GDPR Art. 34, CCPA $§§$1798.82, or analogous statutes. Plaintext-access events are logged and reviewable by the customer.

### 8.5 8.5 Aggregation

When aggregating SCE across windows, the aggregation is **sum of $N^{\text{new}}$ across constituent windows**, not deduplication across the full superset. Rationale: an exposure newly observed in week 3 of a quarter is operationally distinct from one newly observed in week 1, even if the same hash is seen earlier in the historical horizon.

## 9 9. Premium Account Black Market Index (PABMI)

### 9.1 9.1 Definition

PABMI measures the size and liquidity of the public secondary market for the customer's premium accounts. It is computed per window, across all of the customer's monitored brands.

### 9.2 9.2 Formula

$$PABMI_W = L_W \cdot \bar{p}_W \cdot \phi(V_W)$$

where:

- $L_W$ = count of distinct active listings observed in $W$ that mention a brand owned by the customer;
- $\bar{p}_W$ = mean price of those listings in USD (listings priced in cryptocurrency are converted at the daily closing rate from a vendor-published index);
- $\phi(V_W)$ = a vendor-diversity factor based on $V_W$, the number of distinct vendor handles offering listings in $W$:
  - $\phi = 0.7$ if $V_W \leq 2$,
  - $\phi = 1.0$ if $3 \leq V_W \leq 7$,
  - $\phi = 1.25$ if $V_W \geq 8$.

### 9.3 9.3 Listing deduplication

Identical listing texts from the same vendor handle observed in $W$ are counted once. Listings re-posted by different vendors are counted as distinct listings (this reflects market reality: an account being re-offered by multiple resellers indicates a wider distribution).

### 9.4 9.4 Brand mapping

Brand names are normalized against the customer's declared brand list (provided at onboarding and updateable). Partial matches require human-in-the-loop confirmation for the first $N = 20$ occurrences before being added to the automated brand vocabulary.

## 10 10. Brand Impersonation Reach (BIR)

### 10.1 10.1 Definition

BIR measures the count and estimated reach of domains and social-media handles impersonating the customer's brand.

### 10.2 10.2 Formula

$$BIR_W = \sum_{d \in D_W} \log_{10}(1 + r(d)) + \sum_{h \in H_W} \log_{10}(1 + r(h))$$

where $D_W$ is the set of impersonating domains active in $W$, $H_W$ is the set of impersonating social handles active in $W$, and $r(\cdot)$ is the estimated reach: monthly traffic for domains (sourced from a public traffic estimator), follower count for handles.

### 10.3 10.3 Detection rule

A domain $d$ is considered impersonating if (a) its Levenshtein-normalized similarity to one of the customer's primary domains is in $[0.78, 1.0)$, **or** (b) it appears in a curated `known-impersonator` list maintained by the customer or by Dark Web Metrics, **and** (c) it resolves to an active HTTP service whose homepage content classifier indicates content similar to or referencing the customer's brand (false-positive guard).

### 10.4 10.4 Decay and resurrection

A domain that was previously active and stops resolving for $\geq 14$ consecutive days is removed from $D_W$. If it resumes resolving thereafter, it is re-counted as a new entry. Rationale: takedown-and-resurrection patterns are operationally relevant signals.

## 11 11. Cross-tenant peer comparison

### 11.1 11.1 Purpose

The board deck includes a "peer comparison" view that contextualizes a customer's indices against an anonymized peer set. This section specifies the anonymization rules.

### 11.2 11.2 Cohort construction

A peer cohort is defined by (a) tenant tier (Watchtower, Sentinel, Fortress), (b) industry vertical (adult content production / adult content distribution / adult content creator-aggregator / non-adult adjacent), and (c) revenue band where self-declared.

### 11.3 11.3 Minimum cohort size

A cohort is only published in a customer's report if it contains $N \geq 5$ tenants other than the customer. If the cohort is smaller, the comparison is suppressed and replaced with industry-wide aggregates ($N \geq 25$).

### 11.4 11.4 Anonymization

Peer indices are reported as cohort percentiles (P25, P50, P75, P90) without revealing individual tenant identities, names, or revenue figures. Index values are clipped to the same scales used in customer-facing reports.

### 11.5 11.5 Withdrawal

Any customer may opt out of contributing to peer cohorts in writing at any time. Opt-out takes effect at the next window boundary. No retroactive recomputation is performed.

## 12 12. Data exclusions

The following inputs are categorically excluded from indices, regardless of source:

- **CSAM (Child Sexual Abuse Material).** Any observation flagged as CSAM by the ingest classifier is discarded immediately from indices and from raw storage, with only an infohash plus timestamp retained for legal-reporting purposes (NCMEC). See the Data Retention & Deletion Policy for retention rules.
- **Material under active legal hold.** Inputs subject to an internal legal hold (e.g., due to a pending dispute) are excluded from indices for the duration of the hold and flagged in the audit trail.
- **Material from sources known to be compromised.** If a source is identified as having been hijacked, spoofed, or systematically poisoned, all inputs from that source within the compromise window are quarantined and excluded.
- **Retracted observations.** Customer-requested redactions of specific observations (for documented good-faith reasons, e.g., a false-positive performer match) are honored at the next window boundary and trigger recomputation of affected indices.

## 13 13. Quality assurance and drift detection

### 13.1 13.1 Back-testing

Each new methodology version is back-tested against a held-out validation set spanning at least the prior eight (8) weeks of data. The back-test compares the new version's outputs to the prior version's outputs and to a fixed set of "ground-truth" events curated internally (known pre-release leaks, known mass-credential drops, etc.).

A new version is released only if (a) regressions on the ground-truth set are zero, and (b) systemic drift in customer-visible indices (defined as mean absolute deviation across all active tenants' top-decile titles for PVS, and equivalent metrics for the other indices) is below $5\%$.

### 13.2 13.2 Drift monitoring

Daily, the difference between each index's current value and its trailing 28-day exponentially-weighted moving average is computed. Deviations exceeding three standard deviations trigger an internal alert and a root-cause review before customer-facing data is recomputed.

### 13.3 13.3 Source health monitoring

Each source class has a health signal: throughput, latency, and rejection rate. Sustained health degradation (defined per source) triggers fallback to a lower-trust mode (reduced $w_f$) and a data-gap annotation on customer reports, so that drops in observed exposure are not misinterpreted as improvements in the customer's exposure.

# 14 14. Limitations and known biases

These are disclosed to all customers and are part of the methodology by design.

- **Source survival bias.** Sources that disappear (e.g., due to takedown) cease to contribute observations. Without compensating action, this can produce apparent improvements in indices that reflect source loss, not exposure loss. We mitigate this with explicit data-gap annotations on reports.
- **Language bias.** Coverage of non-English sources (especially Russian, Chinese, Spanish, Portuguese, and Vietnamese forums and Telegram channels) is partial. Indices may underestimate exposure for content with strong non-English audiences. Coverage expansion is part of the product roadmap.
- **Indexer self-reporting.** Seeder/leecher counts on public indexers are self-reported and inflatable. We discount indexer counts by 0.5 against DHT counts (Section 6.3) but residual bias remains.
- **Telegram-forwards inflation.** Public Telegram channels can inflate forward counts via bot networks. Our weighting (Section 6.3) caps the contribution of forward counts logarithmically but does not eliminate the bias.
- **Catalog-quality dependency.** The accuracy of PVS and TDI depends on the completeness and quality of the customer's uploaded catalog. Missing titles cannot be matched. Customers are expected to keep their catalog current and accurate; we provide quarterly reconciliation reports surfacing observed-but-unmatched signals that exceed a salience threshold.
- **PII redaction is one-way.** Once a plaintext email is hashed, it cannot be recovered from the hash. Customers needing plaintext for notification obligations must arrange for this within the thirty (30) day plaintext-retention window (Section 8.4).
- **Reach estimators.** BIR depends on third-party traffic estimators with known noise floors and refresh latencies (typically two to four weeks). BIR should be read as a relative-magnitude indicator across domains rather than an absolute traffic figure.

# 15 15. Version control and change management

## 15.1 15.1 Versioning scheme

This methodology document follows semantic versioning: `MAJOR.MINOR.PATCH`.

- **MAJOR** changes alter the mathematical definition of one or more indices in a way that materially affects historical comparability. Customers receive at least sixty (60) days' notice.
- **MINOR** changes add new indices, new optional features, or recalibrate weights within the bounds described in Section 3. Customers receive at least fifteen (15) days' notice.
- **PATCH** changes correct typographical or expository errors without altering computational behavior. Disclosed in the next quarterly report.

## 15.2 15.2 Audit log

Every methodology version, its effective date, and the diff against the prior version are retained indefinitely and made available to active customers on request via `methodology@darkwebmetrics.com`.

### 15.3 15.3 Recomputation policy

Historical recomputation under a new methodology version is performed only when a customer explicitly requests it in writing, and only for periods within the customer's contracted data retention horizon. The original version's outputs are retained alongside the new version's outputs; reports indicate which version was used.

## 16 16. Contact

- Methodology inquiries: `methodology@darkwebmetrics.com`
- Disputes and corrections: `disputes@darkwebmetrics.com`
- Privacy and deletion requests: `privacy@darkwebmetrics.com`

---

*End of document. Methodology version 1.0 — Effective May 2026.*