exposurebureau.com

exposure bureau

DIGITAL-RISK QUANTIFICATION

PUBLIC

Methodology

How Exposure Bureau computes its digital-risk exposure indices — formulas, weights, 80% confidence intervals, and source provenance, audit-replayable end to end.

DOCUMENT DOC 01 VERSION v2.0 DATE 2026-05-23 ENTITY Exposure Bureau LLC

This document specifies, in auditable detail, how Exposure Bureau computes the digital-risk exposure indices it sells. Every index publishes a formula, weights, an 80% confidence interval, and source provenance. Methodology is versioned and audit-replayable: re-running a historical window against the same inputs and the same methodology version yields the same number.

11	5	80%	2+
indices + composite CXI	lead indices on the website	confidence interval per score	sources to corroborate a signal

1 Purpose and audience

Most threat-intelligence platforms deliver alert streams. Decision-makers need numbers they can put in a board deck, an underwriting model, or a diligence memo. This document specifies how Exposure Bureau derives those numbers so that they can be defended at audit and at procurement.

The audience is threefold:

- Client-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 the Exposure Bureau service on behalf of a client (e.g., as part of a SOC 2 readiness assessment of the client'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 hello@exposurebureau.com.

2 Scope and non-scope

2.1 In scope

- Definitions, formulas, inputs, normalization, and windowing for the index family.
- Source taxonomy and weighting rules.
- Cross-source entity resolution (resolving a client to the graph of domains, brands, employees, and vendors it owns).
- Multi-source corroboration, confidence intervals, and the probabilistic loss model that maps a score to dollars at 80% confidence.
- Peer-cohort benchmarking and anonymization rules.
- Quality assurance, back-testing, drift detection, source-status tracking, and version control of the methodology.

- Limitations and known biases.

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 The index family

Exposure Bureau publishes eleven exposure indices plus a composite CXI. The website surfaces five **lead indices** (CES, IMI, RPS, BLV, BIR); the remainder are per-tier add-ons and roadmap. Clients consume the subset relevant to their tier.

INDEX	CODE	WHAT IT MEASURES	PRIMARY BUYER
Credential Exposure Score	CES	Volume, recency, severity of compromised credentials tied to a domain, multi-source-corroborated	CISO, insurer
Public Code Exposure Score	PCES	Leaked secrets, internal source, registry typosquats across GitHub, GitLab, Postman, Docker Hub, package registries	CISO, M&A
External Attack Surface Index	EASI	Externally observable services, cert posture, cloud misconfig, email-auth posture, subdomain sprawl	Insurer
IAB Mention Index	IMI	Frequency and pricing of company access offered on initial-access-broker channels	Insurer, M&A
Ransomware Proximity Score	RPS	KEV-exposed CVEs + EPSS-weighted vuln load + family activity + victim similarity + infra reputation	CISO, board
Brand Leak Velocity	BLV	Rate of brand-asset leakage (source, internal docs, customer lists) across leak channels	Legal, brand
Brand Impersonation Reach	BIR	Lookalike surfaces (typosquats, app clones, social impersonators, extension lookalikes) × audience × persistence decay	Legal, marketing
Supply-Chain Exposure Index	SCEI	Nth-party exposure: vendor criticality × vendor exposure × data sensitivity	M&A, CISO
Regulatory Disclosure Index	RDI	Regulator-confirmed breach history (HHS / SEC / state AG / FTC) over a 60-month window	Insurer, M&A
Sanctions Adjacency Score	SAS	Statistical proximity to OFAC-nexus ransomware-payment risk, as a basis-point multiplier	Insurer

INDEX	CODE	WHAT IT MEASURES	PRIMARY BUYER
Secret Leak Velocity	SLV	New code-leak findings per week, 13-week trailing (companion to PCES)	CISO

A composite **Cyber Exposure Index** (`CXI`) rolls a client-weighted subset of these into a single board-deck-ready number. Underwriters typically weight `CES + EASI + IMI + RPS + SAS`; M&A teams weight `CES + PCES + SCEI + RDI`; mid-market CISOs weight `CES + RPS + BIR + CXI`.

4 Source taxonomy

All inputs come from licensed data partners and public, indexed sources. No authenticated, invite-only, or vouched-access collection. No interaction with sellers, posters, or operators of any source. No HUMINT, no auth bypass, no offensive collection — defensive framing only. Each source class carries a default trust weight `w_f ∈ [0,1]` used in the index formulas, calibrated quarterly.

CLASS	EXAMPLES	DEFAULT TRUST W_F	NOTES
Licensed partners	Have I Been Pwned, DarkOwl Vision, SpyCloud, Constella Intelligence, Flare, Webz.io	0.85	Provider-aggregated; trust calibrated per provider.
Ransomware leak-site metadata	ransomware.live, ransomwatch (aggregated metadata only — we never touch live .onion endpoints)	0.85	Victim-post signal for RPS and BLV.
IAB listings	Initial-access-broker listings surfaced via DarkOwl / Flare	0.60	Core signal for IMI; we never transact or bid.
Public paste / forum / Telegram	Public pastes; public forum archives (archive.org / archive.today); read-only public Telegram channels	0.55	Subject to inflation and recycling; discounted accordingly.
Internet-wide scanning	Shodan, Censys, LeakIX, GrayhatWarfare, certstream/CT logs	0.90	External attack-surface signal for EASI.
Threat-intel / vuln feeds	CISA KEV, EPSS (FIRST), abuse.ch, OFAC SDN, NVD, MITRE ATT&CK	0.90	Core inputs for RPS and SAS.
Regulatory disclosures	HHS OCR, SEC EDGAR 8-K Item 1.05, state AG portals, FTC consent decrees	0.95	Highest-trust signal; feeds RDI.

> **CALIBRATION**
>
> Trust weights are reviewed quarterly. Adjustment criteria: (a) precision/recall against a held-out validation set of confirmed-victim domains, (b) provider reliability incidents, (c) systemic biases discovered during drift review. Calibration changes are version-controlled (Section 13) and disclosed to active clients within fifteen (15) days of taking effect.

5 Computation windows

A window is the temporal unit over which an index is computed. Supported granularities: **day** (24h, UTC-aligned), **week** (ISO 8601, Monday 00:00 – Sunday 23:59 UTC), **month** (calendar, UTC), and **quarter** (calendar, UTC; the board-deck roll-up cadence).

Indices computed at a finer granularity are deterministically aggregated to coarser granularities; the aggregation rule differs by index and is specified per index below.

> 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.

6 Cross-source entity resolution

A "client" is a graph, not a domain. For each engagement we resolve `parent_org → subsidiaries → owned domains → brand strings → employees & contractors → declared vendors → dependent products` from public registry data (OpenCorporates, SEC EDGAR, EU eIDAS), the client's own SBOM, and LinkedIn-resolved affiliate-repo discovery.

Every signal we observe is pinned to the lowest-cardinality node that owns it. This is how a leaked AWS key in a contractor's personal GitHub still scores against the right tenant.

7 Multi-source corroboration as confidence

A signal confirmed by two or more independent sources is weighted higher than a single-source claim. **Single-source claims are leads, not score inputs.** For credential exposure, a hit confirmed by HIBP + DeHashed + IntelX outweighs the same email appearing only in a small forum dump. This is the explicit anti-hallucination control: vendor variance is treated as noise unless corroborated.

8 Credential Exposure Score (CES)

8.1 Definition

CES measures the volume, recency, and severity of distinct credentials (email + secret of any kind) tied to the client's monitored domains, observed in stealer logs, combolists, and breach dumps within a window and corroborated across sources.

8.2 Formula

`CES_W = N_W^new + β · N_W^(high-sev,fresh)` where `N_W^new` is the count of new unique email hashes observed in window `W` on the client's domains (deduplicated against the rolling 180-day history); `N_W^(high-sev,fresh)` is the subset where (a) the severity classification is stealer-cookie or stealer-session and (b)

capture freshness is ≤ 7 days; and $\beta = 2.5$ is the multiplier for high-severity, fresh exposures reflecting substantially higher account-takeover risk. The raw count is mapped to a banded 0–100 score with an 80% confidence interval (Section 12).

8.3 Severity classification

CLASS	DEFINITION	SEVERITY
combolist	Email + password recombined from prior breaches	Low
breach	Email + password from a discrete identified breach (not the client's own)	Medium
stealer-cookie	Stealer log containing cookies for the client's domain	High
stealer-session	Stealer log containing active session tokens / authenticated headers	Critical

Stealer-log families behind this signal include Redline, Raccoon, Vidar, Lumma, StealC, Meta, Rhadamanthys, Atomic, and Mystic. We do not analyze, distribute, or interact with stealer binaries; we consume post-collection log data via licensed partners (SpyCloud, Flare).

8.4 Privacy treatment

All identifiers are hashed at ingest with SHA-256 plus a per-tenant rotating salt (Section 11). Plaintext dereferencing requires verified domain ownership and a documented lawful basis. Plaintext is retained no longer than thirty (30) days and only when explicitly required for a notification obligation under GDPR Art. 34, CCPA §1798.82, or analogous statutes. Plaintext-access events are logged and reviewable by the client.

9 Ransomware Proximity Score (RPS)

9.1 Definition

RPS measures a domain's proximity to ransomware victimization — not whether it has been hit, but how exposed it is — over a trailing 90-day window.

9.2 Formula

$$RPS(D) = 100 \cdot \text{sigmoid}(w1 \cdot KEV\_exposed(D) + w2 \cdot EPSS\_weighted(D) + w3 \cdot FamilyActivity(D) + w4 \cdot VictimSimilarity(D) + w5 \cdot InfraReputation(D))$$

. `KEV_exposed` counts CISA-KEV CVEs matching D's externally-fingerprinted CPEs; `EPSS_weighted` sums EPSS probabilities for unpatched CVEs (capped at 30 to suppress long-tail dominance); `FamilyActivity` counts ThreatFox / MalwareBazaar IOCs tagged to ransomware families touching D's sector in 30 days; `VictimSimilarity` is a Jaccard score between D's firmographic profile and ransomware.live victims posted in the last 90 days; `InfraReputation` penalizes Spamhaus DROP / URLhaus / AbuseIPDB hits on D's ASN. Initial weights `(w1..w5) = (0.30, 0.25, 0.20, 0.15, 0.10)`, calibrated against a held-out set of confirmed-victim domains. Output 0–100, banded Low / Elevated / High / Critical.

10 Other indices — IMI, BLV, BIR, and add-ons

10.1 IAB Mention Index (IMI)

Counts and prices initial-access-broker listings referencing the client, surfaced via licensed providers. We do not transact, bid, vouch, or interact with listed access.

10.2 Brand Leak Velocity (BLV)

Rate of brand-asset leakage (source code, internal docs, customer lists) across leak channels over a trailing window, fed by ransomware leak-site metadata, paste sites, and public Telegram.

10.3 Brand Impersonation Reach (BIR)

`BIR = Σ_surface (count × audience_weight × persistence_decay)` across four surfaces: domain typosquats (dnstwister + crt.sh), mobile impersonators (app stores + sideload catalogs), social impersonators (X / LinkedIn / Instagram / TikTok / Reddit / Bluesky), and browser-extension lookalikes (Chrome / Edge / Firefox / Safari). `persistence_decay` rewards fast takedowns: a clone live for 7 days scores higher than one killed in 24 hours.

10.4 Add-on and roadmap indices

`PCES` aggregates secret hygiene across public code surfaces; `SLV` is its 13-week velocity companion. `EASI` composites external attack-surface posture. `SCEI` scores Nth-party exposure as `Σ (vendor_criticality × vendor_exposure × data_sensitivity) / |V|`. `RDI` scores regulator-confirmed breach history over a 60-month window. `SAS` expresses OFAC-nexus ransomware-payment risk as a basis-point multiplier. Each ships with its own formula, weights, 80% CI, and source provenance.

11 PII hashing and per-tenant salt rotation

Email addresses, usernames, and other identifiers are hashed at ingest with SHA-256 + a per-tenant salt that rotates quarterly. This prevents cross-tenant correlation, bounds the blast radius of any internal leak, and means a database breach on our side never produces a usable cross-tenant identity graph. We store only the hash plus observation metadata. Plaintext email is only ever returned to a verified domain owner via authenticated lookup with a documented lawful basis. Default retention for collected signals is thirteen (13) months, configurable per category.

12 Confidence intervals and loss modeling

12.1 Confidence intervals on every index

Output is `point_estimate ± 80% CI`. Underwriters will not ingest a score without bounds. Confidence is a function of source-coverage breadth, corroboration depth, and recency — a domain with three corroborating sources from the last 30 days has a tight CI; a domain with one stale source has a wide one. The CI is published alongside the score and feeds the underwriter's own pricing-uncertainty model.

12.2 Probabilistic loss modeling — score to dollars

Every index has a calibration model that maps score → expected-loss distribution, calibrated against insurer-validated incident data. The output an underwriter consumes is **$ exposed at 80% confidence**, not just an abstract number. This is the difference between a vendor dashboard and an actuarial input.

12.3 Peer benchmarking

Every score is rendered against a peer cohort (industry × revenue band × geography). The CISO does not get "CES = 47" — they get "CES = 47, peer median = 68, 65th percentile in your sector, trending down 12% QoQ." A cohort is only published if it contains `N ≥ 5` tenants other than the client; otherwise the comparison is suppressed and replaced with industry-wide aggregates (`N ≥ 25`). Peer indices are reported as cohort percentiles (P25, P50, P75, P90) without revealing individual identities. Any client may opt out of contributing to cohorts in writing; opt-out takes effect at the next window boundary with no retroactive recomputation.

13 Quality assurance, drift detection, and source status

13.1 Cross-feed deduplication at ingest

IOCs are normalized to canonical keys before storage (URL → `scheme + eTLD+1 + path-hash`; IP → `ASN + /24`; hash → SHA-256; malware family → MITRE Software ID + Malpedia slug). A `sightings` table is keyed on canonical IOC with one row per `(source, first_seen, last_seen, confidence, family_tag)`. The same payload from URLhaus and MalwareBazaar collapses into one campaign cluster. Deduplication runs at ingest; raw provenance is retained for audit, but scoring consumes deduped sightings only.

13.2 Temporal scoring with category-specific half-lives

Each signal has its own half-life: stealer-log credentials 30 days, breach data 13 months, ransomware victim posts 90 days, IAB listings 60 days, brand impersonation 7 days (until takedown). Scores are time-weighted aggregates with explicit decay constants published in this paper. Nothing scores forever.

13.3 Back-testing and drift monitoring

Each new methodology version is back-tested against a held-out validation set spanning at least the prior eight (8) weeks. A new version is released only if (a) regressions on a curated ground-truth set are zero and (b) systemic drift in client-visible indices is below 5%. Daily, the deviation of each index from its trailing 28-day EWMA is computed; deviations exceeding three standard deviations trigger an internal alert and root-cause review before client-facing data is recomputed.

13.4 Source-status tracking

When a source disappears (forum takedown, API deprecation, channel ban), we record a **data gap** in `ingest.source_status` rather than letting the absence depress a client's score. Otherwise a takedown would be misread as the client's exposure "improving." Source status is part of the methodology trail every report ships with.

13.5 Sector and geo weighting matrices

Different sources matter differently for different industries: healthcare CES weights HHS OCR heavily; finance weights SEC 8-K and NYDFS enforcement; defense weights state-AG portals and SEC Item 1.05 over EU regulators; DTC brands weight Telegram and leak-site mentions. The weighting matrix is published per tier and audited quarterly against incident data.

14 Evidence chain and data exclusions

14.1 Evidence chain with cryptographic anchoring

Every finding is logged with timestamp, source URL, content hash, and chain-of-custody fields suitable for civil-litigation discovery. Quarterly cryptographic anchor to a public timestamping authority (OpenTimestamps) gives any subsequent dispute an independently verifiable proof of when we observed a signal.

14.2 Data exclusions

- **CSAM.** Any observation flagged as Child Sexual Abuse Material at ingest is discarded immediately from indices and raw storage, with only an infohash plus timestamp retained for legal reporting (NCMEC). See the CSAM Detection & Reporting SOP.
- **Material under legal hold.** Inputs subject to an internal legal hold are excluded for the duration of the hold and flagged in the audit trail.
- **Compromised sources.** If a source is identified as hijacked, spoofed, or systematically poisoned, all inputs from that source within the compromise window are quarantined and excluded.
- **Retracted observations.** Client-requested redactions of specific observations (documented good-faith reasons, e.g., a false-positive match) are honored at the next window boundary and trigger recomputation of affected indices.

15 Limitations and known biases

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

- **Source survival bias.** Sources that disappear cease to contribute observations; without compensating action this can produce apparent improvements that reflect source loss, not exposure loss. Mitigated with explicit data-gap annotations (Section 13.4).
- **Language bias.** Coverage of non-English sources (Russian, Chinese, Spanish, Portuguese, Vietnamese forums and Telegram) is partial. Coverage expansion is on the roadmap.
- **Self-reported telemetry.** Listing counts and forward counts on public channels are inflatable; weighting caps their contribution logarithmically but residual bias remains.
- **Entity-resolution dependency.** Index accuracy depends on a complete client graph. Missing subsidiaries, brands, or vendors cannot be matched; we provide quarterly reconciliation surfacing observed-but-unattributed signals above a salience threshold.
- **PII redaction is one-way.** Once an identifier is hashed it cannot be recovered. Clients needing plaintext for notification obligations must arrange it within the 30-day plaintext-retention window (Section 8.4).

- **Reach estimators.** BIR depends on third-party traffic estimators with known noise floors and refresh latencies; read it as a relative-magnitude indicator across surfaces rather than an absolute figure.

16 Version control and change management

16.1 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. Clients receive at least sixty (60) days' notice.
- **MINOR** changes add new indices, new optional features, or recalibrate weights within the bounds described in Section 4. Clients receive at least fifteen (15) days' notice.
- **PATCH** changes correct typographical or expository errors without altering computational behavior. Disclosed in the next quarterly report.

16.2 Audit log and recomputation

Every methodology version, its effective date, and the diff against the prior version are retained indefinitely and made available to active clients on request. Historical recomputation under a new version is performed only on explicit written request, only for periods within the contracted retention horizon, with the original version's outputs retained alongside the new version's; reports indicate which version was used.

17 Contact

- Methodology inquiries, disputes and corrections, privacy and deletion requests: hello@exposurebureau.com

End of document. Methodology version 2.0 — effective 2026-05-23. Exposure Bureau LLC, exposurebureau.com.