Rankiteo Cyber Risk Scoring Algorithm

Jeremy Canale* Maochao Xu

Last Updated: October 2025

1 Overview & Motivation

We produce a single, interpretable score $s \in [100, 1000]$ for each entity, where higher values indicate lower estimated cyber risk. The framework integrates three principal components: (i) *time-decayed incident exposure* that aggregates all reported cyber events with category-specific exponential decay and quantitative severity; (ii) an *industry adjustment* A_{ind} reflecting historical NAICS-level resilience, applied only to clean or near-clean entities; and (iii) a *market-cap–based baseline and dampening* that scale both the starting level and effective penalty according to organizational size.

Traditional cyber-risk ratings often *overreact* to incident disclosures (penalizing transparent, high-visibility firms) or *underreact* to baseline sector risk (treating high-exposure industries like low-exposure ones). Our design balances evidence, context, and interpretability through the following principles:

- **Evidence-driven exposure.** Every confirmed incident contributes to an aggregate penalty P_{inc}, weighted by recency and scaled by quantitative severity (financial loss, records exposed, and ransomware recurrence). Category-specific base weights reflect practical salience: ransomware (100), data breach (60), cyber attack (20), and vulnerability (5). Those category is created by the AI algorithm. Each category also decays at a different rate, roughly three years for ransomware and data breaches, two years for cyber attacks, and 18 months for vulnerabilities, so low-impact, short-lived disclosures fade more rapidly, while severe or repeated events retain longer influence.

- **Sector-aware normalization.** Industry-level resilience enters through a bounded additive term A_{ind}, computed from standardized NAICS v2 incident-rate z-scores. This adjustment applies only to clean or nearly clean firms; once recent exposure is present, the realized record dominates the sector prior.

- **Scale-aware fairness.** The model embeds two continuous size effects: a logistic baseline between 750 and 850, and a multiplicative dampening $A_{\text{size}}(x)$ that attenuates effective penalties for very large firms. In the October 2025 tuning, the lower bound of the dampening curve was reduced to 0.15 and its midpoint shifted to x_0=10.9 to better capture absorption capacity at trillion-dollar scales. Thus, market capitalization influences *both* where a clean entity starts and how strongly incidents pull it downward, providing a principled alternative to discrete "floor" mechanisms.

- **Transparent composition.** The score decomposes as

$$s \approx \underbrace{\text{baseline}(x)}_{\text{market-cap anchor (750–850)}} - \underbrace{A_{\text{size}}(x)\, P_{\text{inc}}}_{\text{time–severity penalties}} + \underbrace{A_{\text{ind}}}_{\text{sectoral normalization}},$$

  followed by clipping to $[100, 1000]$ to preserve interpretability.

- **Stability without opacity.** Deterministic jitter prevents artificial clustering near round thresholds, improving rank smoothness and visual differentiation without altering underlying ordering or comparability.

*Email: contact@jeremycanale.com. https://www.rankiteo.com/

Market-cap baseline (logistic, 750–850). Let $x = \log_{10}(\text{market cap in USD})$ and define

$$h(x) = \sigma\big(\gamma\,(x - x_0)\big), \qquad \sigma(u) = \frac{1}{1 + e^{-u}},$$

with parameters

$$\gamma = 1.82, \qquad x_0 = 10.38.$$

The baseline depends only on market cap:

$$\text{baseline}(x) = 750 + 100\,h(x) \in [750, 850].$$

This maps approximately as

$$\text{baseline}(10^9) = 754, \quad \text{baseline}(10^{10}) = 777, \quad \text{baseline}(10^{11}) = 818, \quad \text{baseline}(10^{12}) = 848.$$

Values are naturally bounded by the logistic, clipping to $[750, 850]$.

Final score. With market-cap dampening $A_{\text{size}}(x)$ applied to the incident load,

$$P_{\text{inc}}^{(\text{eff})} = A_{\text{size}}(x)\,P_{\text{inc}}, \qquad \text{score}_{\text{final}} = \text{clip}\Big(\text{baseline}(x) - P_{\text{inc}}^{(\text{eff})} + A_{\text{ind}},\, 100,\, 1000\Big).$$

Intended use. The score supports cross-industry and longitudinal assessment of organizational cyber exposure. It is designed for portfolio triage, benchmarking, and temporal monitoring, not as a direct breach-frequency predictor. The framework prioritizes: (1) responsiveness to recent, consequential events, (2) cross-sector fairness via NAICS normalization, and (3) scale-aware stability for large firms without masking genuine risk. Together, these properties make the score empirically grounded, context-aware, and operationally robust.

2 Market-Capitalization Estimation for Missing Values

Many privately held or mid-sized entities lack reliable market-capitalization data (mc) in public sources. Because both the logistic baseline and the market-cap dampening factors depend on $\log_{10}(mc)$, we employ a statistically consistent regression estimator to infer $\widehat{mc}$ from observable proxies such as the company's follower count ($foll$) and number of employees (emp).

Hybrid fallback structure. For each company, the algorithm selects among three tiers:

1. If a valid market-cap value is available, use it directly.
2. If both follower count and employee count are available, estimate $\widehat{mc}$ using a polynomial regression on $\log(foll)$ and $\log(emp)$.
3. If only follower count is available, use a follower-only model.

These models were trained on a calibration sample of publicly listed firms with known capitalization values, yielding an adjusted $R^2 \approx 0.49$ and residual standard error ≈ 1.6.

Polynomial estimation model. Using raw (non-orthogonal) polynomial terms, the fitted relationships are:

$$\log(\widehat{mc}) = 23.081 - 1.150\,\log(foll) + 0.0646\,\big[\log(foll)\big]^2 + 0.441\,\log(emp),$$

and, when employee counts are unavailable,

$$\log(\widehat{mc}) = 22.298 - 0.731\,\log(foll) + 0.0624\,\big[\log(foll)\big]^2.$$

Implementation details. Logarithms are computed as $\log(\max(x, 1))$ to prevent undefined values for zero followers or employees. The estimated capitalization is then obtained by $\widehat{mc} = \exp\{\log(\widehat{mc})\}$, and capped at \$1 trillion:

$$\widehat{mc}_{\text{final}} = \min\{\widehat{mc}, 10^{12}\}.$$

No lower bound is imposed, allowing small firms to retain appropriately high penalty sensitivity. For numerical stability, only the argument of $\log_{10}(mc)$ within the logistic baseline and dampening functions is bounded below by 1.

Interpretation. This regression-based estimator preserves monotonicity and scale fairness: larger social presence and workforce size produce larger $\widehat{mc}$, leading to higher baselines and milder penalty dampening. The \$1 trillion cap ensures that all entities at or above that scale are treated equivalently, consistent with the empirical saturation of market-cap effects among globally dominant firms. When true market-cap data are missing, the estimator provides a coherent, empirically calibrated substitute that maintains internal consistency throughout the scoring framework.

3 Incident Scoring (P_{inc})

The component P_{inc} measures the cumulative exposure of a company to adverse cyber events observed in the historical record. The model applies a continuous exponential time decay, so that older incidents gradually lose influence while recent events retain greater weight. Each incident contributes a penalty determined by its type category and empirical severity indicators (e.g., records exposed, financial loss), allowing high-impact breaches to dominate over numerous low-consequence disclosures.

The resulting penalty is additive across all recorded incidents, with diminishing weights applied to successive events to reflect decreasing marginal informational value. Individual incident penalties are bounded by a category-specific cap, and the total P_{inc} is globally limited to prevent extreme cases from overwhelming the overall score. This design ensures that the penalty term remains interpretable, severity-aware, and smoothly responsive to both frequency and intensity of cyber incidents over time.

3.1 Incident Categories & Sector-Sensitive Base Points

Each cyber event is assigned to one of four canonical categories, selected to balance interpretability, frequency coverage, and empirical distinctiveness. The *base points* represent the nominal impact of a fully credible, high-severity event before adjustments for recency, quantitative severity, and market-cap dampening. To reflect that identical technical incidents can have vastly different practical consequences across industries, the canonical base points are modulated by sector-specific multipliers $M_c(g)$ derived from the NAICS v2 code g.

Functional form. For incident category $c \in \{\text{ransomware, data breach, cyber attack, vulnerability}\}$ and NAICS v2 sector g, the sector-adjusted base points are

$$B_i' = B_c \times M_c(g), \qquad B_c \in \{100,\ 60,\ 20,\ 5\}.$$

Multipliers $M_c(g)$ encode the *practical impact* of an incident in its industry context. They are not frequency-based but instead reflect domain knowledge of (1) safety-of-life risk, (2) service continuity, (3) regulatory or legal exposure, and (4) data sensitivity. For each NAICS group, these four dimensions are rated 0–3 and combined as

$$I(g) = \tfrac{1}{3}\big(0.35\,\text{SL} + 0.30\,\text{SC} + 0.25\,\text{RG} + 0.10\,\text{DS}\big), \qquad M_c(g) = 1 + \theta_c I(g),$$

where θ_{rw}=0.35, θ_{db}=0.25, θ_{ca}=0.15, θ_{vl}=0.10, capped to the interval $[0.9, 1.4]$. Sectors with higher $I(g)$ values—such as hospitals, utilities, or national security—thus amplify the base penalty more strongly than low-criticality sectors.

Table 1: Sector-sensitive impact multipliers $M_c(g)$ by NAICS v2 sector. Values reflect practical consequences (safety-of-life, service continuity, regulatory exposure, and data sensitivity).

NAICS	**Sector (short name)**	SL	SC	RG	DS	$I(g)$	M_{rw}	M_{db}	M_{ca}	M_{vl}
211	Oil & Gas Extraction	1	3	2	1	0.63	1.22	1.16	1.09	1.06
212	Mining (except Oil)	1	2	1	0	0.40	1.14	1.10	1.06	1.04
221	Utilities (Power/Water)	2	3	2	1	0.73	1.26	1.18	1.11	1.07
230	Construction	1	2	1	0	0.37	1.13	1.09	1.06	1.04
311	Food Manufacturing	0	2	1	1	0.33	1.12	1.08	1.05	1.03
312	Beverage/Tobacco Manufacturing	0	2	1	1	0.33	1.12	1.08	1.05	1.03
323	Printing/Related Support	0	1	0	0	0.13	1.05	1.03	1.02	1.01
325	Chemical Manufacturing	1	3	2	1	0.67	1.23	1.17	1.10	1.07
326	Plastics/Rubber Products	0	2	1	0	0.27	1.09	1.07	1.04	1.03
333	Machinery Manufacturing	1	3	1	1	0.53	1.19	1.13	1.08	1.05
334	Computer & Electronic Products	0	2	1	1	0.33	1.12	1.08	1.05	1.03
335	Electrical Equipment & Appliances	1	3	1	1	0.53	1.19	1.13	1.08	1.05
336	Transportation Equipment	2	3	1	1	0.57	1.20	1.14	1.09	1.06
339	Misc. Manufacturing / Medical Devices	2	2	2	2	0.67	1.23	1.17	1.10	1.07
420	Wholesale Trade	0	1	1	1	0.23	1.08	1.06	1.03	1.02
441–454	Retail Trade (avg.)	0	1	1	2	0.30	1.10	1.07	1.04	1.03
481–488	Transportation & Warehousing	0	3	1	1	0.50	1.18	1.13	1.08	1.05
511	Publishing / Software	0	2	2	2	0.53	1.19	1.13	1.08	1.05
515	Broadcasting	0	2	2	1	0.47	1.17	1.12	1.07	1.05
517	Telecommunications	1	3	2	1	0.67	1.23	1.17	1.10	1.07
518	Data Processing / Hosting / Cloud	0	3	2	2	0.60	1.21	1.15	1.09	1.06
520	Monetary Authorities / Central Bank	1	3	3	3	0.87	1.30	1.22	1.13	1.09
522	Banking / Credit Intermediation	1	2	3	3	0.73	1.26	1.18	1.11	1.07
523	Securities / Investment	0	2	3	3	0.67	1.23	1.17	1.10	1.07
524	Insurance / Actuarial Services	1	2	3	3	0.70	1.25	1.18	1.11	1.07
525	Funds / Trusts / Pensions	0	2	3	2	0.60	1.21	1.15	1.09	1.06
541	Professional, Scientific, Tech. Services	0	2	2	2	0.53	1.19	1.13	1.08	1.05
561	Admin. & Support / Facilities Services	0	2	1	1	0.33	1.12	1.08	1.05	1.03
611	Education Services	0	1	1	2	0.27	1.10	1.07	1.04	1.03
621	Ambulatory Health Care	2	2	3	2	0.73	1.26	1.18	1.11	1.07
622	Hospitals	3	3	3	2	0.87	1.30	1.22	1.13	1.09
623	Nursing / Residential Care	2	2	2	1	0.63	1.22	1.16	1.09	1.06
624	Social Assistance / NGOs	1	2	2	1	0.53	1.19	1.13	1.08	1.05
711–713	Arts, Recreation, Entertainment	0	1	1	1	0.23	1.08	1.06	1.03	1.02
721–722	Accommodation & Food Services	0	1	1	1	0.23	1.08	1.06	1.03	1.02
810	Religious / Civic Organizations	0	1	1	0	0.17	1.06	1.04	1.02	1.02
920	General Public Administration	1	3	3	2	0.80	1.28	1.20	1.12	1.08
921	Executive / Legislative / Government	1	3	3	2	0.80	1.28	1.20	1.12	1.08
922	Justice, Public Order, Safety	2	3	3	2	0.87	1.30	1.22	1.13	1.09
923	Human Resource / Welfare Programs	1	3	3	2	0.80	1.28	1.20	1.12	1.08
928	National Security / Defense	2	3	3	1	0.83	1.29	1.21	1.12	1.08

Interpretation.

- The canonical base points B_c establish cross-category comparability: ransomware (100), data breach (60), cyber attack (20), and vulnerability (5).
- Multipliers $M_c(g)$ embed domain-informed judgment about safety-of-life, continuity, and systemic criticality. Thus, identical incidents (e.g., a ransomware outbreak) have greater impact for hospitals, utilities, or defense entities than for retail or manufacturing firms.
- The structure is intentionally stable: updates occur only when new regulations or technologies materially shift the real-world consequence profile.
- These sector-adjusted base points feed directly into the quantitative severity multipliers (m_L, m_R), followed by temporal decay and aggregation into the entity-level exposure P_{inc}.

3.2 Category-Specific Time Decay

Let age_days_i denote the number of days since incident i occurred. Each event is attenuated by a continuous exponential factor

$$w_{\text{time},i} = \exp(-\lambda_{c_i}\,\text{age\_days}_i)\,, \qquad \lambda_{c_i} = \frac{\ln 2}{h_{c_i}},$$

where h_{c_i} is the half-life (in days) associated with the incident's category c_i. Thus every h_{c_i} days, the effective weight of an event in category c_i is reduced by one half.

Assigned half-lives by incident type.

$$h_{\text{ransomware}} = h_{\text{breach}} = 1095 \text{ d (3 years)}, \quad h_{\text{attack}} = 730 \text{ d (2 years)}, \quad h_{\text{vulnerability}} = 540 \text{ d (18 months)}.$$

This refinement ensures that long-lasting, high-impact incidents such as ransomware and large breaches persist longer in memory, whereas transient vulnerabilities and generic cyber attacks decay more rapidly.

Illustration for representative categories ($P_i^{\max} = 100$).

$$P_i(t) = 100 \cdot \exp\!\left(-\tfrac{\ln 2}{h_{c_i}}\, t\right)$$

Time since event	$w_{\text{time}}^{(\text{rw/db})}$	$w_{\text{time}}^{(\text{attack})}$	$w_{\text{time}}^{(\text{vuln})}$	Interpretation
0 years (today)	1.000	1.000	1.000	full impact
1 year (365 d)	0.77	0.70	0.62	moderate decay begins
2 years (730 d)	0.60	0.50	0.38	vulnerability largely faded
3 years (1095 d)	0.50	0.35	0.24	half-life for ransomware/breach
5 years (1825 d)	0.32	0.19	0.09	long-past residual trace

Compared to the uniform 3-year decay previously used, this differentiated structure allows vulnerability disclosures to fade after roughly a year and a half, while high-severity ransomware and breach events remain influential for a full three years, better matching observed reputational and regulatory persistence.

3.3 Recurrence Escalation for Ransomware

Ransomware recurrence is a robust indicator of weak cyber hygiene, persistent adversarial foothold, or failure to remediate root causes. To reflect this heightened systemic risk, repeated ransomware events are subject to a bounded, time-weighted *recurrence multiplier* that can double the effective ransomware penalty when clustering is severe.

Recent ransomware mass. Let $\mathcal{R}$ index all ransomware incidents for a company, and define

$$w_{\text{time},i} = \exp\left(-\tfrac{\ln 2}{1095}\,\text{age\_days}_i\right), \qquad S_{\text{rw}} = \sum_{i\in\mathcal{R}} w_{\text{time},i}.$$

Here, a 3-year half-life ($h = 1095$ days) ensures that an event's weight halves every three years. The soft count S_{rw} thus behaves as an "effective number" of recent ransomware: a fresh event contributes 1.0, a 3-year-old event contributes 0.5, etc.

Revised recurrence multiplier (October 2025 tuning). We moderate escalation so that multiple recent ransomware incidents still increase the ransomware component, but the total multiplier is capped at a 50% boost rather than a full doubling. The specification is

$$M_{\text{rw}} = 1 + \rho\,(S_{\text{rw}} - 1)_+, \qquad (x)_+ \equiv \max(0, x), \quad M_{\text{rw}} \le M_{\max},$$

with

$$\rho = 0.25, \qquad M_{\max} = 1.5.$$

Thus:

- a single fresh ransomware (S_{rw}=1) $\Rightarrow$ M_{rw}=1 (no escalation);
- two fresh ransomware ($S_{\text{rw}}\approx$2) $\Rightarrow$ M_{rw}=1.25 (25% boost);
- three fresh ransomware ($S_{\text{rw}}\approx$3) $\Rightarrow$ M_{rw}=1.5 (capped at 50%).

Worked example. Suppose a firm experienced three ransomware events: one today (t=0), one a year ago (t=365 days), and one four years ago (t=1460 days). Then

$$w_A = 1.000, \quad w_B = 2^{-365/1095} \approx 0.794, \quad w_C = 2^{-1460/1095} \approx 0.397,$$

so $S_{\text{rw}} = 1.000 + 0.794 + 0.397 = 2.191$. With $\rho = 0.25$ and $M_{\max} = 1.5$,

$$M_{\text{rw}} = 1 + 0.25\,(2.191 - 1) = 1.297, \quad \text{rounded to } M_{\text{rw}} \approx 1.30.$$

If the decayed ransomware subtotal is $P_{\text{rw,sub}} = 219$, then the adjusted ransomware load is

$$P_{\text{rw}} = M_{\text{rw}} \times P_{\text{rw,sub}} \approx 1.30 \times 219 = 284.7.$$

Illustrative scenarios.

Scenario	S_{rw}	$M_{\text{rw}} = 1 + 0.25(S_{\text{rw}} - 1)_+$	Escalation?
One fresh incident	1.00	1.00	No
Two fresh incidents	2.00	1.25	Mild
Three fresh incidents	3.00	1.50	Moderate (capped)
Fresh + 3-year-old	1.50	1.13	Minimal
Four or more clustered events	≥ 3	hits cap 1.5	Moderate

Interpretation. This specification:

- leaves single ransomware cases unpenalized beyond their base weight;
- penalizes recurrence roughly in proportion to effective count, up to a 1.5× limit;
- better aligns with observed practice where repeated ransomware within short horizons suggest partial containment weaknesses but not necessarily systemic failure.

3.4 Quantitative Severity: Market-Cap–Adjusted Financial Loss and Records Exposed

Numeric disclosures of monetary or data impact provide an objective basis to modulate incident severity. When available, these values scale the base category penalty upward, reflecting that more costly or far-reaching events indicate greater systemic failure. If quantitative details are missing, the base severity remains unchanged (*i.e.*, no downward adjustment).

Each incident's categorical score is multiplied by two continuous, log-scaled severity factors—one for *financial loss* (USD) and one for *records exposed*. Both factors are normalized by firm size to ensure that identical dollar or record counts have proportionally larger influence on smaller organizations. The final multiplier is their product, bounded above by a global cap of 3.0×, which prevents extreme outliers from dominating aggregate exposure while preserving proportionality across orders of magnitude.

Financial-loss multiplier (relative to market capitalization). Let L be the disclosed loss in USD and mc the firm's market capitalization. Define the relative loss ratio $r_L = L/\mathrm{mc}$, and the scaled argument

$$x = \log_{10}\!\left(1 + 10^9\, r_L\right) = \log_{10}\!\left(1 + 10^9\,\frac{L}{\mathrm{mc}}\right).$$

The multiplier is then

$$m_L(x) = \begin{cases} 1 + \dfrac{x}{12}, & x \le 6 \quad (\text{loss} \lesssim 0.001 \text{ of market cap}), \\ \min\!\left\{\, 1.5 + 0.5(x-6),\, 3.0 \,\right\}, & x > 6. \end{cases}$$

This formulation anchors $m_L = 1.5$ when the loss equals roughly 10^{-3} of market capitalization (*e.g.*, a \$1 M loss at a \$1 B firm or a \$100 M loss at a \$100 B firm), and grows toward the hard cap of 3.0 for proportionally larger losses.

Table 2: Market-cap–adjusted financial-loss multiplier m_L (examples).

Market Cap	**Loss (USD)**	$r_L = L/\mathbf{mc}$	$\log_{10}(1 + 10^9 r_L)$	m_L
\$100 M	\$1 M	0.010	7.00	2.00
\$1 B	\$1 M	0.001	6.00	1.50
\$10 B	\$1 M	0.0001	5.00	1.42
\$100 B	\$10 M	0.0001	5.00	1.42
\$1 T	\$100 M	0.0001	5.00	1.42

Records-exposed multiplier (relative to market capitalization). Let R be the number of records exposed. Normalize record counts by market capitalization through a policy constant η representing the typical record capacity per \$1 B of firm value (default $\eta = 5$, corresponding to 5 million records per \$1 B):

$$r_R = \frac{R}{\eta \times 10^6 \times (\mathrm{mc}/10^9)}.$$

Compute

$$y = \log_{10}\!\left(1 + 10^6 r_R\right), \qquad m_R(y) = \begin{cases} 1 + \dfrac{y}{12}, & y \le 6, \\ 1.5 + 0.5\,(y-6), & 6 < y \le 9, \\ 3.0, & y > 9\ . \end{cases}$$

This structure preserves the familiar logarithmic progression while automatically moderating large firms with proportionally greater data holdings.

Table 3: Market-cap–adjusted records multiplier m_R (examples, $\eta = 5$).

Market Cap	**Records**	$r_R = R/\text{capacity}$	$\log_{10}(1 + 10^6 r_R)$	m_R
\$50 M	1,000,000	4.0	6.60	1.80
\$1 B	1,000,000	0.20	5.30	1.44
\$100 B	10,000,000	0.02	4.30	1.36
\$1 T	100,000,000	0.02	4.30	1.36

Computation. For each incident with adjusted base category weight B_i, the total severity multiplier is

$$m_{\text{sev}} = m_L \times m_R, \qquad P_i = B_i \times m_{\text{sev}}.$$

Both m_L and m_R are smooth, monotone functions of the relative loss and exposure, each capped at 3.0 to maintain comparability. Incidents without quantitative disclosure default to $m_L = m_R = 1.0$, so qualitative classification alone determines baseline impact.

Interpretation.

- Severity now reflects the *proportional* impact of an incident relative to the firm's financial and operational scale.
- Small and mid-cap entities experience stronger amplification for identical losses or record counts, while trillion-dollar firms see modest increments unless the exposure is truly exceptional.
- The global 3× cap limits tail inflation but preserves meaningful differentiation across four orders of magnitude.

3.5 Aggregation Across Incidents

All incident contributions P_i are first adjusted for quantitative severity (financial loss and records exposed) and then time–decayed according to their category-specific half-life:

$$w_{\text{time},i} = \exp\left(-\frac{\ln 2}{h_{c_i}} \,\text{age\_days}_i\right), \qquad h_{c_i} \in \{1095, 730, 540\} \text{ for ransomware/breach, attack, vulnerability.}$$

The resulting unbounded exposure is

$$P_{\text{raw}} = \sum_i P_i \, w_{\text{time},i}.$$

Ransomware incidents receive an additional recurrence adjustment as described previously:

$$M_{\text{rw}} = 1 + \rho \left(S_{\text{rw}} - 1\right)_+, \qquad \rho = 0.25,\ M_{\max} = 1.5,$$

applied to the decayed ransomware subtotal. The adjusted ransomware component $P_{\text{rw}} = M_{\text{rw}} \sum_{i \in \mathcal{R}} P_i w_{\text{time},i}$ replaces the simple sum of ransomware terms.

The final aggregate penalty is then

$$P_{\text{inc}} \;=\; \min\Big(C_{\text{ent}},\; P_{\text{rw}} + \sum_{i \notin \mathcal{R}} P_i \, w_{\text{time},i}\Big), \qquad C_{\text{ent}} = 600.$$

The entity-level cap C_{ent} maintains interpretability on the 100–1000 scale and prevents large, disclosure-rich firms from accumulating unbounded penalties.

Interpretation.

- Only ransomware incidents trigger multiplicative escalation through M_{rw}, reflecting recurrence risk.
- All penalties accumulate additively after time decay, until the global cap C_{ent} is reached.
- The combination of time decay, recurrence adjustment, and sector-specific scaling yields smooth, bounded exposure trajectories that remain comparable across industries and firm sizes.

4 Industry Adjustment (A_{ind})

The industry adjustment now serves as a lightweight, reputation-based normalization applied *only* to firms that exhibit a clean or near-clean historical record. Rather than continuously blending historical breach rates, the updated design grants a modest, sector-informed offset to reflect baseline industry resilience, but withdraws it once any material or recent incident occurs.

Applicability. Let N_{5y} be the count of incidents within the last five years and N_{old} the count of older incidents. The industry adjustment is applied only when

$$N_{5y} = 0 \quad \text{or} \quad \big(N_{5y} = 0,\ N_{\text{old}} = 1,\ \text{and the single incident is non-ransomware}\big).$$

In all other cases—for any recent event or repeated history— A_{ind} is set to zero so that realized performance dominates industry priors.

Base scaling. Each three-digit NAICS v2 group g carries an empirically estimated standardized incident-rate z-score $z_{\text{rate}}(g)$, from which a baseline offset is computed as

$$A_{\text{ind,base}} \;=\; \text{clip}\big(-25,\, +25,\, -10 \times z_{\text{rate}}(g)\big),$$

where high historical breach prevalence (e.g., information, utilities) yields negative values, and historically resilient sectors (e.g., education, manufacturing) yield positive ones.

Application rule. For qualifying "clean" entities, the active adjustment is a softened version of $A_{\text{ind,base}}$:

$$A_{\text{ind}} = \text{sign}(A_{\text{ind,base}})\,|A_{\text{ind,base}}|^{0.9}, \qquad \text{if the firm meets the clean-history criteria;}$$

otherwise,

$$A_{\text{ind}} = 0.$$

Rationale. This policy reflects that sector-level resilience matters most for firms with no demonstrated exposure history. Once a company has any recent or repeated incident—even a minor disclosure within five years—firm-specific outcomes dominate, and the industry adjustment is neutralized. This preserves interpretability and fairness by rewarding sectoral cleanliness only where it is credible.

Illustration. Suppose a firm in a relatively safe sector has $A_{\text{ind,base}} = +20$:

$$A_{\text{ind}} = \text{sign}(+20)\,|20|^{0.9} \approx +15.9.$$

If the same firm records a single minor breach six years ago (non-ransomware), the same offset applies. However, any incident within the last five years or multiple events of any age reset $A_{\text{ind}} = 0$.

Interpretation.

- A_{ind} now functions as a *sectoral cleanliness bonus* rather than a continuous correction.
- Clean or long-stable firms inherit modest credit for operating in safer industries, improving comparability across sectors.
- Firms with any recent or repeated incident receive no industry adjustment, so realized risk experience drives the score.

5 Market-Cap Dampening of Incident Load

To preserve comparability across firms of vastly different scale while avoiding the discontinuities of hard "floors," we apply a continuous *market-cap dampening* to the total incident load P_{inc}. Large-cap organizations tend to experience higher disclosure frequency due to transparency and brand visibility, yet are typically better resourced to contain and remediate events. The dampening therefore reduces the effective penalty in proportion to firm size without masking genuinely severe or systemic risk.

Functional form. Let $x = \log_{10}(\text{market cap in USD})$. Define a smooth, size-dependent carry factor

$$A_{\text{size}}(x) = a_{\min} + (1 - a_{\min}) \frac{1}{1 + \exp\{\gamma\,(x - x_0)\}}, \qquad a_{\min} = 0.15,\ \gamma = 1.52,\ x_0 = 10.9.$$

The adjusted incident load is

$$P_{\text{inc}}^{(\text{eff})} = A_{\text{size}}(x)\, P_{\text{inc}}.$$

The factor A_{size} lies in $[0.15, 1]$, approaching 1 for small or mid-cap firms and declining smoothly toward 0.25 for trillion-dollar firms.

Illustration.

Market Cap	$\log_{10}(\text{mc})$	A_{size} (carry)
\$1B	9.0	0.96
\$10B	10.0	0.83
\$100B	11.0	0.46
\$1T	12.0	0.28
\$3T	12.5	0.22

Application. After computing the base incident penalty P_{inc} (including severity, time decay, and ransomware recurrence), we apply the dampening:

$$\text{score}_{\text{final}} = \text{baseline}(x) - P_{\text{inc}}^{(\text{eff})} + A_{\text{ind}}, \qquad P_{\text{inc}}^{(\text{eff})} = A_{\text{size}}(x) P_{\text{inc}},$$

followed by clipping to $[100, 1000]$. The cleanliness factor $C = 1 - P_{\text{inc}}/C_{\text{ent}}$ used in the industry adjustment continues to rely on the *undampened* P_{inc}, ensuring that large firms cannot retain sector credit once material incident burden accumulates.

Interpretation.

- The logistic form ensures a smooth, continuous moderation of incident penalties with scale—no abrupt thresholds or tiers.
- Small and mid-sized entities (\$¡10 B) are essentially unaffected, while \$100 B–\$1 T firms retain roughly 25–50 % of their raw incident load.
- The adjustment recognizes structural resilience and disclosure intensity without obscuring genuinely high-severity or repeated events, maintaining fairness across sectors and firm sizes.

6 Worked Example: Logistic Baseline and Market-Cap Dampening under Ransomware

Company: MegaTech (NAICS 511, market capitalization $\approx$ \$1 trillion; $\log_{10}$ mc = 12.0). Two comparable ransomware cases illustrate how the logistic baseline and market-cap dampening interact to moderate—but not erase—the impact of severe incidents.

Scenario A: Severe Ransomware

A confirmed double-extortion ransomware event occurred 9 months ago (270 days), with a financial loss of \$5 million and 50,000 records exposed.

- **Baseline (logistic):** $\text{baseline}(x{=}12.0) = 750 + 100\,\sigma[1.82(x - 10.38)] = 848$.
- **Base category:** $B_i = 100$.

- **Time decay:** $w_{\text{time}} = 2^{-270/1095} = 0.84$.
- **Market-cap dampening:** $A_{\text{size}}(x{=}12.0) = 0.28$.
- **Severity multipliers (market-cap adjusted):**
$$m_L = 1.85, \quad m_R = 1.39, \quad m_L m_R = 2.57.$$
- **Credibility (confirmed, double-extortion):** $\times 1.5$
$\Rightarrow M_{\text{total}} = 0.84 \times 2.57 \times 1.5 = 3.24$.
- **Incident contribution:** $P_i = 100 \times 3.24 = 324$.
- **Industry adjustment:** recent ransomware $\Rightarrow A_{\text{ind}} = 0$.
- **Effective penalty (after dampening):** $P_{\text{inc}}^{(\text{eff})} = 0.28 \times 324 = 90.7$.
- **Final score:**
$$s = 848 - 90.7 + 0 = \mathbf{757.3}.$$

Interpretation. Despite its trillion-dollar scale, MegaTech's severe ransomware event lowers the score to roughly 757, reflecting a major but not catastrophic penalty after size-based moderation.

—

Scenario B: Minor Ransomware

Same company and timing, but the ransomware caused only \$0.5 million in loss and no records were exposed.

- **Baseline:** 848.
- $B_i = 100, \quad w_{\text{time}} = 0.84, \quad A_{\text{size}} = 0.28$.
- **Severity multipliers:** $m_L = 1.20,\ m_R = 1.00 \Rightarrow m_L m_R = 1.20$.
- **Credibility (confirmed, single event):** $\times 1.5$
$\Rightarrow M_{\text{total}} = 0.84 \times 1.20 \times 1.5 = 1.51$.
- **Incident contribution:** $P_i = 100 \times 1.51 = 151$.
- **Industry adjustment:** recent ransomware $\Rightarrow A_{\text{ind}} = 0$.
- **Effective penalty:** $P_{\text{inc}}^{(\text{eff})} = 0.28 \times 151 = 42.3$.
- **Final score:**
$$s = 848 - 42.3 + 0 = \mathbf{805.7}.$$

Interpretation. A contained, low-loss ransomware incident reduces the score modestly, keeping MegaTech near 806—still within the upper large-cap range. The framework recognizes proportional harm without over-penalizing scale.

—

Comparison Summary.

Scenario	Loss (USD)	Baseline	Market-Cap Factor	Final Score
A. Severe ransomware	\$5 M	848	0.28	757
B. Minor ransomware	\$0.5 M	848	0.28	806

Key Insight. The logistic baseline anchors large, clean entities near 850, while market-cap dampening reduces—but never cancels—the effect of realized incidents. Severe ransomware still produces a substantial score drop, whereas minor, isolated events yield proportionally smaller penalties. This provides a continuous, interpretable alternative to the discrete "floor" logic of earlier designs.