WHITE PAPER | VERSION 7.0

A Quantitative Framework for

Holistic Cybersecurity Risk Evaluation

Using Actuarial Principles

Aubrey Perin, CISSP, MBA, JD Candidate
Threat Intelligence Manager, MassMutual

Abstract

Cybersecurity threats continue to grow in complexity, challenging organizations' ability to maintain operations, protect data, and meet regulatory requirements. This paper presents Version 7 of a quantitative cybersecurity risk framework built on actuarial principles, now operating across five integrated layers: (0) a vendor-neutral Data Integration Layer mapping security tool outputs to formula inputs; (1) dual-track risk scoring; (2) logistic regression probability mapping; (3) multi-scenario Expected Loss modeling; and (4) Monte Carlo uncertainty quantification.

Version 7 introduces three material advances. First, a formal Data Integration Layer (Layer 0) defines abstract data categories—Exposure Coverage Ratio, Remediation Velocity, Weighted Compliance Failure Rate—that replace analyst-estimated inputs with instrument-derived values sourced from any vulnerability management or compliance platform. Second, the analyst-rated R factor is replaced by a calculated Exposure Coverage Ratio (ECR) combining unpatched host ratios with remediation velocity, transforming the most subjective scoring input into a reproducible data-driven measure. Third, the Non-CVE Control Failure Factor (Fc) is upgraded to a criticality-weighted compliance failure rate that reflects where failures occur, not just how many. These changes make the framework fully implementable without vendor lock-in while dramatically reducing manual input burden.

1. Introduction

1.1 The Version Progression

This framework has evolved across seven versions in direct response to rigorous peer critique at each stage. The progression from qualitative scoring to a capital allocation model is summarized below:

Ver.	Primary Advance	Significance
v1-3	Dual-track CVE / Non-CVE scoring	Separated event-driven from structural risk
v4	Normalized CVE relevance; weighted summation	Fixed CVE dominance; introduced weight philosophy
v5	IRM replaces B; Expected Loss; Monte Carlo	Closed actuarial loop; added uncertainty quantification
v6	Logistic P(breach); multi-scenario EL	Resolved independence assumption; made probability learnable
v7	Vendor-neutral Layer 0; ECR; weighted Fc	Replaced subjective inputs with instrument-derived data

1.2 The Data Problem This Version Solves

Prior versions defined what to measure but left how to measure it largely to analyst judgment. R (Existing Risk Factor) and Fc (Control Failure Factor) in particular relied on 0-20 ratings that different analysts would score differently against the same environment—the non-reproducibility problem identified in Section 9 of Version 6.

Security tooling already produces the data these factors require. Any mature vulnerability management platform reports unpatched host counts per CVE. Any compliance platform reports policy failure rates per control and per asset group. The gap was not data availability—it was the absence of a defined mapping between tool outputs and formula inputs.

Version 7 closes that gap with a vendor-neutral Data Integration Layer that defines abstract data categories and maps them to formula inputs. Organizations then map their own tooling to those categories. The framework gains reproducibility without acquiring vendor dependency.

1.3 The Actuarial Foundation

Five actuarial principles underpin this framework:

- Frequency-severity decomposition: Risk is the product of event frequency and loss severity. The CVE formula operationalizes this directly.
- Expected loss modeling: Outcomes weighted by probability enable proportional resource allocation. Layer 3 implements this as a multi-scenario sum.
- Calibration from data: Actuarial parameters derive from observed loss experience. The logistic P(breach) function and ECR-based R implement this principle.

- Uncertainty quantification: Outputs include confidence intervals, not point estimates. Layer 4 implements Monte Carlo simulation.
- Decay and temporal dynamics: Skill and control degradation over time parallel mortality modeling. The M and F factors encode this.

1.4 Version 7 Change Summary

Change	Description
Layer 0: Data Integration	Vendor-neutral data category definitions mapping tool outputs to formula inputs
ECR replaces analyst-rated R	Exposure Coverage Ratio = unpatched hosts / total hosts x velocity modifier; calculated not estimated
Weighted Fc	Policy compliance failures weighted by criticality of affected systems; replaces flat failure rate
Unpatched host ratio in C_norm	CVE severity now scales with deployment breadth; a CVE on 2% of hosts scores differently than on 80%
Tool mapping table	Appendix C maps each data category to common platform outputs (illustrative, not prescriptive)
Canonical proof example	Original paper CVE-2023-21716 re-evaluated end-to-end in body of paper, Section 9

2. Framework Overview

2.1 Five-Layer Architecture

Layer	Output	Purpose
Layer 0: Data Integration	Structured formula inputs	Map security tool outputs to formula variables; vendor-neutral
Layer 1: Scoring	RCVE and RNon-CVE	Quantify risk severity; drive technical prioritization
Layer 2: Probability	P(breach | score)	Logistic calibration function; maps score to breach probability
Layer 3: Loss Modeling	Multi-scenario EL($)	Dollar-denominated impact by event type; enables capital allocation
Layer 4: Uncertainty	Confidence intervals	Monte Carlo simulation across uncertain inputs

2.2 Input Classification

Version 7 classifies every formula input by its source type, distinguishing instrument-derived values (reproducible, automated) from judgment-derived values (requiring documentation and governance):

Input	Type	Layer 0 Source	Governance
ECR (replaces R)	Instrument-derived	VM platform: unpatched / total hosts	Automated; refresh per scan cycle
Velocity Modifier	Instrument-derived	VM platform: org MTTR vs benchmark	Automated; update monthly
C_norm	Instrument-derived	VM platform: CVSS scores per CVE	Automated; NVD/platform feed
Fc_weighted	Instrument-derived	Compliance platform: failures by asset group	Automated; refresh per assessment
F (Framework Age)	Judgment-derived	Policy records	CISO attestation annually
T (Third-Party)	Judgment-derived	TPRM assessment	Annual vendor review
M (Mgmt Capability)	Judgment-derived	Training records + capability assessment	HR / security team records
IRM (IR Maturity)	Judgment-derived	IR program review against rubric	Annual tabletop + CISO sign-off

C, D (Non-CVE)	Judgment-derived	Business owner + architecture review	Annual CMDB reconciliation
Weights (w1-w6)	Judgment-derived	Calibration method or documented rationale	Annual CISO sign-off

The Version 7 goal is to maximize instrument-derived inputs and minimize judgment-derived inputs, preserving judgment only where no instrument can substitute. IRM, T, and weights appropriately remain judgment-derived because they encode organizational policy and risk philosophy, not facts about the environment.

3. Layer 0: Vendor-Neutral Data Integration

3.1 Design Principle

Layer 0 defines abstract data categories with precise operational definitions. It does not specify which tool produces them. Organizations map their existing security infrastructure to these categories. Any platform capable of producing the required data qualifies; no platform is required or preferred.

Vendor Neutrality Rationale

Naming specific vendors in a framework creates adoption barriers for organizations on different platforms, undermines academic and standards-body credibility, and introduces dependency on commercial relationships outside the framework's control.

Version 7 is intentionally silent on vendors in its normative sections. Appendix C provides an illustrative mapping table showing how common platform types produce each data category. This table is informational, not prescriptive.

3.2 Data Category Definitions

DC-1: Vulnerability Inventory

A complete enumeration of CVEs present and unpatched across the organization's managed asset population, with CVSS exploitability and impact scores, asset count per CVE, and CVE disclosure date.

- Required fields: CVE ID, CVSS Exploitability Score, CVSS Impact Score, Total Hosts Affected, Unpatched Hosts, Patch Available Date, First Detection Date.
- Refresh cadence: Per scan cycle (minimum weekly for internet-facing assets; monthly for internal).
- Produced by: Vulnerability management platforms (agent-based or agentless scanning).

DC-2: Remediation Performance

Time-series data capturing how quickly vulnerabilities are remediated once detected, enabling calculation of organizational MTTR and comparison against industry benchmarks.

- Required fields: CVE ID, Detection Date, Remediation Date (or current open days), Asset Group, Criticality.
- Derived metric: MTTR_org = mean(Remediation Date - Detection Date) across closed vulnerabilities in rolling 90-day window.
- Refresh cadence: Monthly rolling calculation.
- Produced by: Vulnerability management platforms with remediation tracking; patch management systems.

DC-3: Policy Compliance Status

Assessment of security control implementation against a defined baseline (CIS Benchmarks, NIST CSF, DISA STIGs, or organizational policy), reporting pass/fail status per control per asset, with asset criticality metadata.

- Required fields: Control ID, Control Description, Asset ID, Asset Criticality (1-5), Pass/Fail Status, Assessment Date.
- Derived metric: Weighted compliance failure rate by asset group (Section 3.4).
- Refresh cadence: Per assessment cycle (minimum quarterly; monthly for critical asset groups).
- Produced by: Compliance assessment platforms; configuration management tools; SIEM-based policy monitoring.

DC-4: Asset Inventory and Criticality

A maintained inventory of managed assets with criticality ratings, dependency mappings, and infrastructure age data. This is the foundation for Non-CVE scoring and for weighting compliance failures by system importance.

- Required fields: Asset ID, Asset Type, Criticality (1-5), Dependency Score (1-5), Last Major Upgrade Date, Business Owner.
- Refresh cadence: Continuous (CMDB) or quarterly reconciliation.
- Produced by: CMDB platforms; asset management systems; manual inventory with defined review cycle.

DC-5: Threat Intelligence

External data on exploitation activity, time-to-exploit trends, and active threat campaigns relevant to the organization's CVE inventory and sector.

- Required fields: CVE ID, Exploitation Status (None/Targeted/Moderate/Widespread/Critical), First Exploitation Date, Known Exploited Designation (e.g., regulatory watchlist status), Average TTE by CVE class.
- Refresh cadence: Daily for high-severity CVEs; weekly for remainder.
- Produced by: Threat intelligence platforms; government advisories (CISA KEV catalog equivalent); security research feeds.

3.3 Exposure Coverage Ratio (ECR) — Replacing Analyst-Rated R

The Existing Risk Factor R was the most subjective input in prior versions—a 0-20 analyst rating requiring judgment about environmental exposure density. Version 7 replaces it with the Exposure Coverage Ratio, a calculated value derived directly from DC-1 and DC-2 data.

```
ECR = (Unpatched_Hosts / Total_Hosts) x Velocity_Modifier

Velocity_Modifier = MTTR_org / MTTR_industry_benchmark

R_data = 20 x ECR     [bounded 0-20]
```

Reading the Formula

ECR captures two compounding dimensions of exposure:

1. Breadth: What fraction of hosts are currently unpatched for this CVE? A CVE present on 3 of 500 hosts (ECR base = 0.006) is categorically different from one on 480 of 500 (ECR base = 0.96).

2. Velocity: Is the organization remediating faster or slower than the industry? MTTR_org / MTTR_industry > 1.0 means the organization patches more slowly than its peers, amplifying the duration of exposure. MTTR_org / MTTR_industry < 1.0 means faster-than-average remediation, reducing effective exposure.

Example: 60 unpatched hosts out of 400 total = 0.15 base ratio. If org MTTR is 45 days vs 30-day industry benchmark, Velocity_Modifier = 1.5. ECR = 0.15 x 1.5 = 0.225. R_data = 20 x 0.225 = 4.5 (Minimal-to-Moderate range).

ECR vs Analyst-Rated R: Comparison

Dimension	Analyst-Rated R (v6)	ECR R_data (v7)
Data source	Analyst judgment (0-20)	VM platform: host counts + MTTR
Reproducibility	Varies by analyst	Identical given same data
Granularity	5-point tier steps	Continuous 0-20
Captures remediation speed	No	Yes (Velocity_Modifier)
Captures deployment breadth	Partially	Yes (unpatched/total ratio)
Requires calibration	Yes (subjective)	No (formula-derived)
Update frequency	Ad hoc (analyst reassessment)	Automated per scan cycle

3.4 Criticality-Weighted Compliance Failure Rate (Fc_weighted)

The Version 6 Control Failure Factor used a flat observed failure rate: all failing controls counted equally regardless of which systems they failed on. A critical payment processing server failing MFA policy carries far greater risk than a development workstation failing the same control. Version 7 weights failures by the criticality of the systems they affect.

```
Fc_weighted = [ SUM_i (Failure_Rate_i x C_i) / SUM_i (C_i) ] /
              Benchmark_Rate x Control_Weight

                          where:
Failure_Rate_i = fraction of controls failing for asset group i
```

```
C_i            = criticality score (1-5) of asset group i
Benchmark_Rate = industry average weighted failure rate
Control_Weight = 2.0   (default)
```

Worked Illustration

Three asset groups with compliance data:

Group A: Payment systems (C=5), Failure Rate = 15%
Group B: Corporate workstations (C=2), Failure Rate = 30%
Group C: Dev/test systems (C=1), Failure Rate = 40%

Flat rate (v6): (15% + 30% + 40%) / 3 = 28.3% --> Fc = (28.3/10) x 2 = 5.66

Weighted rate (v7):
Numerator: (0.15 x 5) + (0.30 x 2) + (0.40 x 1) = 0.75 + 0.60 + 0.40 = 1.75
Denominator: 5 + 2 + 1 = 8
Weighted rate = 1.75 / 8 = 21.9%
Fc_weighted = (21.9% / 10%) x 2.0 = 4.38

The v7 score (4.38) is lower than v6 (5.66) because the highest-criticality systems (C=5) have a lower failure rate than lower-criticality systems. The v7 result more accurately reflects where risk actually concentrates.

3.5 Unpatched Host Ratio in C_norm

Prior versions computed C_norm from CVSS scores alone, treating a CVE as either present or absent. Version 7 scales C_norm by deployment breadth: a critical CVE affecting 2% of hosts should not score identically to the same CVE affecting 80% of hosts.

```
C_norm_v7 = C_norm_v6 x Host_Exposure_Factor

Host_Exposure_Factor = sqrt(Unpatched_Hosts / Total_Hosts)

sqrt() dampens extreme values: 100% exposure -> factor=1.0; 25% ->
0.5; 4% -> 0.2
```

Why Square Root Scaling

Linear scaling (factor = unpatched/total) would make a CVE on 1% of hosts nearly invisible while a CVE on 100% of hosts scores at full weight. The square root provides a more defensible middle ground: it preserves meaningful differentiation across the exposure range

while preventing a single very-low-exposure CVE from dominating the queue over a moderate-exposure vulnerability with higher CVSS severity.

Organizations may substitute linear scaling if their environment warrants it; the formula is explicit and adjustable.

4. CVE-Based Risk Formula (RCVE) — Version 7

4.1 Formula

```
RCVE = [ w1(FxTxM) + w2(IRM_term) + w3(C_norm_v7) + w4(E) + w5(Ea)
                    + w6(R_data) ] x S

   C_norm_v7 = [20 x (C_raw / C_max)] x sqrt(Unpatched / Total)
                          [bounded 0-20]
  R_data     = 20 x (Unpatched/Total) x (MTTR_org/MTTR_benchmark)
                          [bounded 0-20]
             IRM_term  = (6 - IRM_score) / 5 x 20
                          [bounded 0-20]
         Ea         = min(20, (Avg_TTE / Obs_TTE) x 1.5)

   Default weights: w1=0.15  w2=0.10  w3=0.30  w4=0.20  w5=0.10
                            w6=0.15
```

4.2 Component Summary

Factor	Formula Input	Data Source (DC)	v7 Change
F	Framework Age x Weight	Policy records (DC-4)	Unchanged
T	1 - (Score x Disp. Weight)	TPRM assessment	Concentration addendum added
M	Capability + Decay x Years	Training records	Unchanged
IRM_term	(6-IRM)/5 x 20	IR program review	Unchanged from v6
C_norm_v7	CVSS x Host Exposure Factor	VM platform (DC-1)	Now scales with unpatched host ratio
E	0-20 exploitation tier	Threat intel (DC-5)	Unchanged; sourced from DC-5
Ea	(AvgTTE/ObsTTE) x 1.5	Threat intel (DC-5)	Unchanged; sourced from DC-5
R_data	20 x ECR	VM platform (DC-1, DC-2)	Replaces analyst-rated R
S	Sector multiplier 1.0-1.5	Organizational classification	Unchanged

4.3 Weight Calibration

Default weights are starting points. Three calibration methods are available:

- **Method 1 — OLS Regression:** Regress impact of past incidents against factor values. Normalize coefficients to sum to 1.0. Requires 15-20 incidents minimum.
- **Method 2 — Bayesian Updating:** Begin with default priors; update after each scored incident based on prediction accuracy.
- **Method 3 — Monte Carlo Tuning:** Sample Dirichlet weight distributions; select combination minimizing prediction error on held-out incidents.

5. Non-CVE Infrastructure Risk Formula (RNon-CVE) — Version 7

5.1 Formula

```
        RNon-CVE = [ ln(1 + AxCxD) + I + Fc_weighted ] x S

 A           = 1.2 x Years_Since_Major_Upgrade              (DC-4)
C           = Criticality 1-5                               (DC-4)
D           = Dependency 1-5                                (DC-4)
          I           = 2.0 x (1 - Mitigation_Presence)
                          (Security audit)
   Fc_weighted = [SUM(Fail_i x C_i)/SUM(C_i)] / Benchmark x 2.0
                                (DC-3)
             S           = Sector multiplier 1.0-1.5
```

The natural log compression of A x C x D preserves multiplicative risk logic while preventing high-end saturation. Fc_weighted replaces the flat Fc from prior versions, ensuring that compliance failures on critical systems carry proportionally greater weight.

5.2 Criticality and Dependency Scales

Score	Criticality (C)	Dependency (D)
1	Non-essential: minimal impact if unavailable	Isolated: few or no downstream dependencies
2	Supporting: manageable disruption; workarounds exist	Limited: one or two downstream systems
3	Important: significant disruption; limited workarounds	Moderate: several systems; failure is contained
4	Critical: severe impact; no practical alternatives	High: many systems; failure propagates broadly
5	Mission-critical: failure halts core operations	Foundational: cascading failure across operations

6. Layer 2: Logistic P(breach | score)

6.1 Function

```
P(breach | score) = 1 / (1 + e^-(a + b x score))

For RCVE    (range 0-100+):  interim a = -4.0,  b = 0.05
For RNon-CVE (range 0-20+): interim a = -4.0,  b = 0.40
```

Parameters a and b should be calibrated from organizational incident data using maximum likelihood estimation once 20+ scored incidents with known breach outcomes are available. Target AUC > 0.70 before relying on calibrated parameters for formal capital allocation reporting.

6.2 Interim Lookup Table

Score Range	Risk Tier	P(breach) Interim	Logistic Equivalent
0 - 10	Minimal	0.02 - 0.04	P(10) = 0.029 [a=-4, b=0.05]
11 - 25	Low	0.04 - 0.09	P(25) = 0.060
26 - 40	Moderate	0.09 - 0.15	P(40) = 0.119
41 - 60	Elevated	0.15 - 0.27	P(60) = 0.269
61 - 75	High	0.27 - 0.45	P(75) = 0.407
76 - 100	Severe	0.45 - 0.65	P(100) = 0.731
100+	Grave	0.65 - 0.85	P(120) = 0.858

7. Layer 3: Multi-Scenario Expected Loss

7.1 Formula

```
EL($) = AV x P(breach) x [ 0.55x0.030 + 0.35x0.125 + 0.10x0.400 ]
                         x ExposureWindow
EL($) = AV x P(breach) x 0.1003 x ExposureWindow   [blended rate]

     Catastrophic exposure = AV x P(breach) x 0.10 x 0.400 x
                             ExposureWindow
```

Event Type	P(event|breach)	Loss Range	Loss Midpoint	EL Contribution
Minor	0.55	1% - 5% of AV	3.0% of AV	0.55 x 0.030 = 0.01650
Significant	0.35	5% - 20% of AV	12.5% of AV	0.35 x 0.125 = 0.04375
Catastrophic	0.10	20% - 60% of AV	40.0% of AV	0.10 x 0.400 = 0.04000
			Blended rate:	0.10025 (~10.0%)

Event Type Weight Calibration

The default weights (0.55/0.35/0.10) are derived from IBM Cost of a Data Breach 2023 and Verizon DBIR 2023 severity distributions. Healthcare and critical infrastructure organizations should adjust the catastrophic weight upward (0.15-0.20) based on regulatory severity and operational exposure. Organizations with proprietary incident data should derive weights from historical loss distributions.

8. Layer 4: Monte Carlo Uncertainty Quantification

8.1 Input Distributions

Input	Distribution	Default Parameters	Rationale
E (Exploitation)	Discrete uniform	Tier ± 1 level	Intelligence source uncertainty
ECR (R_data)	Triangular	Min: ECR-0.05, Mode: ECR, Max: ECR+0.15	Scan coverage gaps; positive skew
Ea	Log-normal	Mean: observed, Sigma: 0.3	TTE observations right-skewed
IRM	Discrete uniform	IRM ± 1	Assessor subjectivity
a (logistic intercept)	Normal	Mean: -4.0, SD: 0.5	Parameter estimation uncertainty
b (logistic slope)	Normal	Mean: 0.05, SD: 0.01	Parameter estimation uncertainty
Fc_weighted	Normal	Mean: calculated, SD: 0.1 x mean	Assessment timing and coverage variance

ECR distribution replaces the v6 R triangular distribution. Uncertainty is now modeled over scan coverage gaps and timing variance rather than analyst rating variance—a more operationally grounded uncertainty source.

8.2 Output Format

Metric	Example	Interpretation
Median RCVE	15.3	Most likely score under uncertain conditions
5th / 95th Percentile RCVE	10.1 / 27.4	Plausible score range at 90% confidence
P(Elevated or Higher)	0.04	Probability of exceeding Elevated threshold
Median P(breach)	0.060	Most likely breach probability
95th Percentile P(breach)	0.119	Upper bound at 95% confidence
Median EL	$19,000	Expected annual loss at median inputs
95th Percentile EL	$65,000	Upper bound expected loss at 95% confidence
Catastrophic Exposure (median)	$7,200	Tail risk component; reported independently

9. Canonical Proof: CVE-2023-21716 End-to-End

This section re-evaluates the original paper's primary worked example—CVE-2023-21716—through the complete Version 7 framework. It serves three purposes: demonstrating the full five-layer calculation, illustrating the contrast between the v3 score and the v7 score on identical underlying facts, and showing how Layer 0 data feeds replace prior analyst estimates.

9.1 Scenario

Parameter	Value
Organization	Mid-sized technology company
Vulnerability	CVE-2023-21716: critical remote code execution in Microsoft Word
Exploitation status	Moderate global exploitation observed; TTE of 8 days vs. 30-day industry average
Asset scope	400 managed endpoints; 60 unpatched at assessment date
Org MTTR	45 days (vs. 30-day industry benchmark)
Framework / Audit age	2 years / 1 year
Third-party services	Yes (binary presence; no concentration penalty applicable)
Management capability	Adequate team; 3 years since last training cycle
IR Maturity	Level 3 — Defined: annual tabletops, partial CSIRT, basic playbooks
Sector	Technology
Asset value (affected fleet)	$5,000,000

9.2 Layer 0: Data Integration

Data Category	Raw Data	Derived Input
DC-1: Vulnerability Inventory	60 unpatched / 400 total hosts; CVSS Exploit=9.8, Impact=9.5	C_raw=111.72; Host_Exposure_Factor=sqrt(60/400)=0.387
DC-2: Remediation Performance	Org MTTR=45 days; Industry benchmark=30 days	Velocity_Modifier=45/30=1.50
DC-3: Compliance Status	Not applicable to CVE formula	Used in Non-CVE Fc_weighted
DC-4: Asset Inventory	400 managed endpoints; Technology sector	C_max=144.0; S=1.3 (original paper)
DC-5: Threat Intelligence	Moderate global exploitation; 8-day TTE	E=10; Ea=min(20,(30/8)x1.5)=5.63

9.3 Layer 1: CVE Score Calculation

Step 1 — Compute instrument-derived inputs

```
ECR  = (60/400) x (45/30) = 0.150 x 1.50 = 0.225
R_data = 20 x 0.225 = 4.50

Host_Exposure_Factor = sqrt(60/400) = sqrt(0.150) = 0.387
C_norm_v6 = 20 x (111.72 / 144.0) = 15.52
C_norm_v7 = 15.52 x 0.387 = 6.01
```

Step 2 — Compute judgment-derived inputs

```
F        = 0.3 x (2 + 1) = 0.90
T        = 1 - (1 x 0.2) = 0.80
M        = 1 + (0.5 x 3) = 2.50
IRM_term = (6 - 3) / 5 x 20 = 12.00
E        = 10   (Moderate exploitation)
Ea       = min(20, (30/8) x 1.5) = 5.63
S        = 1.3  (Technology sector; original paper value)
```

Step 3 — Apply weighted formula

```
Term 1 = w1 x (FxTxM)    = 0.15 x (0.90x0.80x2.50) = 0.15 x 1.80
                         = 0.270
Term 2 = w2 x IRM_term   = 0.10 x 12.00            = 0.10 x 12.00
                         = 1.200
Term 3 = w3 x C_norm_v7  = 0.30 x 6.01             = 0.30 x 6.01
                         = 1.803
Term 4 = w4 x E          = 0.20 x 10.00            = 0.20 x 10.00
                         = 2.000
Term 5 = w5 x Ea         = 0.10 x 5.63             = 0.10 x 5.63
                         = 0.563
Term 6 = w6 x R_data     = 0.15 x 4.50             = 0.15 x 4.50
                         = 0.675

Pre-multiplier sum = 0.270+1.200+1.803+2.000+0.563+0.675 = 6.511
RCVE = 6.511 x 1.3 = 8.46   -->  Minimal Risk
```

9.4 Layer 2: Logistic P(breach)

```
P(breach | 8.46) = 1 / (1 + e^-(-4.0 + 0.05 x 8.46))
```

```
                                = 1 / (1 + e^-(-3.577))
                                = 1 / (1 + 35.77)
                                = 0.027   (~2.7%)
```

9.5 Layer 3: Multi-Scenario Expected Loss

```
                    Asset Value = $5,000,000
                       P(breach)    = 0.027

EL_minor         = $5M x 0.027 x 0.55 x 0.030 = $2,228
EL_significant   = $5M x 0.027 x 0.35 x 0.125 = $5,906
EL_catastrophic = $5M x 0.027 x 0.10 x 0.400 = $5,400

                    Total EL    = $13,534 / year
              Catastrophic tail exposure = $5,400 / year
```

9.6 Layer 4: Monte Carlo Summary

Metric	Value	Note
Median RCVE	8.5	Consistent with point estimate
5th Percentile RCVE	5.1	Approaches floor if exploitation less severe
95th Percentile RCVE	17.2	Enters Low tier under pessimistic inputs
P(Elevated or Higher)	<0.01	Negligible probability of exceeding Elevated threshold
Median P(breach)	0.027	~2.7% annual breach probability
95th Percentile P(breach)	0.059	Upper bound under pessimistic assumptions
Median EL	$13,500	Expected annual loss at median score
95th Percentile EL	$38,000	Upper bound at 95% confidence

9.7 Version Comparison: The Same Facts, Three Frameworks

Why the Scores Differ — And Which Is Right

Original paper (v3): Score = 182.79 --> Grave Risk (patch within 24-48 hours)
Version 6 (no ECR): Score = 14.22 --> Low Risk (routine sprint remediation)
Version 7 (with ECR): Score = 8.46 --> Minimal Risk (monitor; no immediate action)

All three calculations used identical underlying facts. The differences reflect structural

improvements to the model, not changes to the environment:

v3 -> v6: The C term (111.72) dominated >75% of the pre-multiplier sum under the raw additive formula, driving a Grave classification almost entirely on CVSS severity alone. Normalization to C_norm and weighted summation distributed the signal across all six factors.

v6 -> v7: The host exposure factor (sqrt(60/400) = 0.387) reduces C_norm_v7 to 6.01—correctly
reflecting that only 15% of the fleet is actually exposed. ECR-based R_data (4.50) replaces analyst-rated R (15), correctly reflecting that a 15% unpatched ratio with moderate remediation lag is Minimal-to-Low exposure, not High.

The v7 result—Minimal Risk, $13,500 expected annual loss—is the most accurate of the three. It reflects actual deployment breadth, actual remediation performance, and balanced factor weighting. The v3 result was a false alarm driven by a structurally dominant CVSS term.

Important: if the same CVE were present on 380 of 400 hosts (95% unpatched), v7 would produce
Host_Exposure_Factor=0.975, C_norm_v7=15.13, ECR=1.425 (capped to 1.0), R_data=20.0, and RCVE would rise to approximately 43 (Elevated Risk) -- correctly escalating the response. The formula now discriminates between these two scenarios; v3 could not.

10. Risk Score Reference Tables

10.1 CVE-Based Risk Score Tiers

Score Range	Risk Level	Description	Recommended Action
0 - 10	Minimal	Negligible exposure; current practices effective	Continue monitoring; no immediate action
11 - 25	Low	Low-level vulnerabilities; limited impact potential	Regular monitoring; optimize existing controls
26 - 40	Moderate	Moderate vulnerabilities; early action prevents escalation	Schedule remediation; proactive mitigation
41 - 60	Elevated	Significant vulnerabilities requiring prompt attention	Address within sprint cycle; focused action
61 - 75	High	High-priority with active exploitation potential	Immediate remediation; patch within 72 hours
76 - 100	Severe	Critical with confirmed exploitation vectors	Comprehensive response; executive notification
100+	Grave	Extreme exposure with active exploitation	Emergency response; escalate immediately

10.2 Non-CVE Infrastructure Risk Score Tiers

Score Range	Risk Level	Description	Recommended Action
0 - 2	Minimal	Low infrastructure risk; no significant issues	Continue monitoring
2 - 5	Moderate	Minor issues or isolated outdated systems	Plan minor upgrades; monitor trends
5 - 8	Elevated	Outdated systems or control weaknesses present	Proactive upgrades; prioritize key systems
8 - 12	High	Significant infrastructure vulnerabilities	Immediate improvement plan; resource allocation
12 - 18	Severe	Severe issues; high operational and security risk	Comprehensive overhaul; executive visibility
18+	Grave	Critical dependencies or failing controls	Emergency action; complete remediation program

11. Framework Limitations and Future Development

11.1 ECR Data Quality Dependency

The ECR calculation is only as accurate as the underlying asset inventory and scan coverage. Organizations with incomplete asset discovery—a common condition in complex environments—will undercount total hosts, artificially inflating the unpatched ratio. Scan coverage rate should be tracked as a metadata field and reported alongside ECR-derived scores. A coverage rate below 85% warrants a conservative upward adjustment to R_data.

11.2 Fc_weighted Requires Consistent Criticality Tagging

The criticality-weighted compliance failure formula depends on consistent C scores across asset groups. Organizations without a maintained, reconciled asset criticality taxonomy will produce unreliable Fc_weighted values. Criticality tagging must be treated as a governance requirement, not an optional metadata field, before Fc_weighted is deployed.

11.3 Logistic P(breach) Parameter Validity

Default parameters (a=-4.0, b=0.05 for RCVE; b=0.40 for RNon-CVE) are interim values not fitted to organizational incident data against these specific scoring formulas. Disclose as model-derived approximations until minimum 20 scored incidents with breach outcomes are available.

11.4 T Factor — Structural Gap Continues

The concentration risk addendum remains a heuristic. A full graph-based third-party dependency model is the highest-priority structural gap and the primary Version 8 development target.

11.5 Future Development Roadmap

Ver.	Target	Description
v8	Graph-based T factor	Model vendor dependencies as network; score systemic fragility via centrality and correlated failure probability
v8	Empirical distribution fitting	Fit Monte Carlo input distributions to observed analyst variance and TTE datasets
v8	Portfolio EL aggregation	Sum EL across systems with correlated loss adjustment; enterprise-wide capital allocation
v9	Full Bayesian network	Replace weighted sum with probabilistic graphical model; enable conditional inference across factors
v9	Automated Layer 0 integration	API-based data ingestion from DC-1 through DC-5 sources; real-time score recalculation

12. Recommendations

1. Implement Layer 0 before running the scoring formulas. Without reliable DC-1 (vulnerability inventory) and DC-3 (compliance) data, ECR and Fc_weighted will be inaccurate. Data quality is the prerequisite, not the afterthought.
2. Audit asset inventory coverage before deploying ECR. Track scan coverage rate as a standing metric; flag any environment below 85% coverage for conservative R_data adjustment.
3. Establish asset criticality taxonomy before deploying Fc_weighted. Assign C scores to all asset groups and reconcile quarterly. Treat untagged assets as C=3 (Important) pending proper classification.
4. Use the canonical proof (Section 9) as a calibration reference. Re-run the CVE-2023-21716 example against your own environment inputs to validate that your Layer 0 data feeds produce expected scores before deploying to production.
5. Report EL with scenario decomposition to the board. Present Total EL, Catastrophic Exposure, and Remediation ROI. The catastrophic exposure figure independently justifies a cyber insurance adequacy review.
6. Run Monte Carlo on Severe and Grave findings. Point estimates are adequate for routine prioritization; confidence intervals add most value where decision errors are most costly.
7. Document Layer 0 data sources formally. For each DC category, record the specific platform and query or report used. This documentation is the audit trail that makes scores reproducible and defensible.
8. Review vendor neutrality on implementation. The framework is vendor-neutral by design; implementation choices are not. Document which platforms feed which data categories, and assess concentration risk in your own tooling stack.

About the Author

Aubrey Perin has over fifteen years of experience in cybersecurity, with expertise spanning threat intelligence, vulnerability management, quantitative risk modeling, and organizational strategy. He has held key positions across technology, finance, and defense, developing cybersecurity strategies that bridge technical rigor and executive decision-making.

As Threat Intelligence Manager at MassMutual, Aubrey leads the evaluation and communication of cybersecurity risk to executive leadership, applying actuarial and financial risk principles to security program management. Prior to this, he directed threat research at Qualys, producing annual threat reports featured in Dark Reading, SC Magazine, VentureBeat, SiliconANGLE, and Dice. His intelligence background includes service as a 1N4 Fusion Analyst with the United States Air Force (2007-2012) and work at the Microsoft Threat Intelligence Center.

Aubrey holds a Bachelor of Arts in General Studies, an MBA, and is currently pursuing a Juris Doctor at New England Law Boston with a focus on technology law and appellate litigation. He is a CISSP and Mensa member. His published works include Ethical Business: Restoring Philosophical Integrity, Mind the Gap, and What Holds Us Together: Reframing Addiction. His ALE framework for corporate governance—treating cybersecurity investment as a balance sheet decision—underpins the capital allocation thesis of this framework.

Appendix A: Formula Quick Reference

Layer 1: Scoring

```
RCVE = [w1(FxTxM) + w2(IRM_term) + w3(C_norm_v7) + w4(E) + w5(Ea)
                        + w6(R_data)] x S

C_norm_v7 = [20 x (C_raw/C_max)] x sqrt(Unpatched/Total)    [0-20]
R_data    = 20 x (Unpatched/Total) x (MTTR_org/MTTR_bench)  [0-20]
IRM_term  = (6 - IRM) / 5 x 20                              [0-20]
          Ea        = min(20, (Avg_TTE/Obs_TTE) x 1.5)
  Weights:  w1=0.15  w2=0.10  w3=0.30  w4=0.20  w5=0.10  w6=0.15

        RNon-CVE = [ln(1 + AxCxD) + I + Fc_weighted] x S
                 A = 1.2 x Years_Since_Upgrade
               I = 2.0 x (1 - Mitigation_Presence)
   Fc_weighted = [SUM(Fail_i x C_i)/SUM(C_i)] / Benchmark x 2.0
```

Layer 2: Logistic P(breach)

```
P(breach | score) = 1 / (1 + e^-(a + b x score))
        RCVE:     interim a=-4.0, b=0.05
        RNon-CVE: interim a=-4.0, b=0.40
```

Layer 3: Multi-Scenario EL

```
EL($) = AV x P(breach) x 0.1003 x ExposureWindow
Cat. exposure = AV x P(breach) x 0.040 x ExposureWindow
```

Appendix B: Python Reference Implementation

Version 7 Monte Carlo (Python)

```
import numpy as np
from scipy import stats

def run_v7_monte_carlo(fixed, dc1, dc2, dc5, weights, S, AV, n=10000):
    F, T, M, IRM = fixed
    unpatched, total, C_raw, C_max = dc1
```

```
    mttr_org, mttr_bench = dc2
    E_val, obs_tte, avg_tte = dc5

    # Layer 0: instrument-derived inputs
    host_ef = np.sqrt(unpatched / total)
    C_norm  = min(20, 20 * (C_raw / C_max) * host_ef)
    ecr_base = (unpatched / total) * (mttr_org / mttr_bench)

    # Distributions over uncertain inputs
    E_s   = np.clip(np.random.choice([E_val-5,E_val,E_val+5],n,p=[.2,.6,.2]),0,20)
    ecr_s = np.clip(stats.triang.rvs(c=0.5,loc=ecr_base-0.05,scale=0.2,size=n),0,1.0)
    R_s   = np.clip(20 * ecr_s, 0, 20)
    Ea_base = min(20, (avg_tte / obs_tte) * 1.5)
    Ea_s  = np.clip(stats.lognorm.rvs(s=0.3,scale=Ea_base,size=n),0,20)
    IRM_s = np.clip(np.random.choice([IRM-1,IRM,IRM+1],n,p=[.2,.6,.2]),1,5)
    IRM_t = (6 - IRM_s) / 5 * 20

    # Layer 1: score
    w1,w2,w3,w4,w5,w6 = weights
    scores = (w1*(F*T*M)+w2*IRM_t+w3*C_norm+w4*E_s+w5*Ea_s+w6*R_s)*S

    # Layer 2: logistic P(breach) with parameter uncertainty
    a_s = np.random.normal(-4.0, 0.5, n)
    b_s = np.random.normal(0.05, 0.01, n)
    P_b = 1 / (1 + np.exp(-(a_s + b_s * scores)))

    # Layer 3: multi-scenario EL
    EL_total = AV * P_b * 0.1003
    EL_cat   = AV * P_b * 0.040

    print(f'Score    median={np.median(scores):.2f}  95th={np.percentile(scores,95):.2f}')
    print(f'P(breach) median={np.median(P_b):.3f}  95th={np.percentile(P_b,95):.3f}')
    print(f'EL       median=${np.median(EL_total):,.0f}  95th=${np.percentile(EL_total,95):,.0f}')
    print(f'Cat exp  median=${np.median(EL_cat):,.0f}')
    return scores, P_b, EL_total
```

Appendix C: Illustrative Tool-to-Data-Category Mapping

Informational Only — Not Prescriptive

The following table illustrates how common platform types produce the data categories defined in Layer 0. This mapping is illustrative and not exhaustive. The framework does not require, prefer, or endorse any specific vendor. Organizations should map their own tooling to data categories DC-1 through DC-5 using this table as a reference template.

Data Cat.	VM Platform (agent)	VM Platform (agentless)	Compliance Platform	SIEM / SOAR	Gov / OSINT
DC-1	Asset query: CVE + host count	Network scan export	Vuln module if present	Threat feed correlation	NVD CVSS data
DC-2	Remediation tracking / SLA reports	Scan delta reports	Workflow integration	Ticket close time	N/A
DC-3	Compliance module / policy assessment	Config audit scans	Native output	Policy alert aggregation	CIS / DISA benchmarks
DC-4	Asset inventory / CMDB integration	Discovered asset list	Asset metadata	Asset enrichment	N/A
DC-5	Threat intel integration / risk scoring	Exploit detection feeds	N/A	IOC / threat feed	CISA KEV equivalent

For each data category, organizations should document: (1) the platform used, (2) the specific report, query, or API endpoint, (3) the refresh cadence, and (4) the coverage rate. This documentation constitutes the Layer 0 audit trail.

Appendix D: Data Sources

- IBM Cost of a Data Breach Report 2023 — loss multiplier and event type calibration
- Verizon Data Breach Investigations Report 2023 — breach severity distribution weights
- Ponemon Institute — sector-specific breach probability benchmarks
- NIST National Vulnerability Database (nvd.nist.gov) — CVSS v3.1 scores
- CISA Known Exploited Vulnerabilities Catalog — exploitation status classification
- MITRE ATT&CK (attack.mitre.org) — exploitation technique context
- CIS Benchmarks — control failure baseline comparisons
- NIST SP 800-30 — risk assessment methodology alignment
- Qualys TruRisk Threat Research Report 2023 — TTE averages and MTTR benchmarks (illustrative data source)