Prevention Credits: **Quantifying Avoided Loss in Renewable Energy Operations — Whitepaper v1.0 (October 2025)**

Mark Ivankovich — Founder, SkinnyCowboy.ai

White Paper v1.0 | DOI: https://doi.org/10.65331/wp-CEKD7 | October 2025 |

Abstract

Every day, intelligent maintenance systems quietly prevent costly failures, truck rolls, and downtime events that never appear in accounting ledgers. These *non-events* represent measurable operational and financial value; today that value is invisible.

This paper introduces the **Prevention Credit (PC)**: a standardized, risk-adjusted unit that converts verified avoided loss into an auditable metric of economic benefit. Drawing on reliability engineering, probability theory, and AI governance principles, we derive a general equation linking baseline hazard rates, intervention effects, expected loss magnitude, and verification confidence:

Eq. 1: General Prevention Credit Equation (Summary Form)

$$PC = a \cdot q \cdot (1-h) \cdot \left[\mu_y (1-\theta) \left(\int \lambda_0(t)\, dt - K \right) \right]$$

Where each term captures attribution, confidence, conservatism, mean loss, hazard reduction, and cost.

We outline how operational data, from SCADA, CMMS, or telemetry logs, feeds the model, and how governance rules ensure transparency and replicability.

By making prevention quantifiable, Prevention Credits enable operators, insurers, and investors to value reliability improvements the same way markets already price carbon reductions.

This work reframes prevention as a line item of verifiable avoided cost, enabling CFOs and ESG officers to quantify reliability improvements within financial statements. The framework generalizes across any data-rich operational environment, enabling consistent valuation of reliability across sectors.

White Paper v1.0 | DOI: https://doi.org/10.65331/wp-CEKD7 | October 2025 |

1 Introduction & Context

Industrial operators have long measured production and emissions, yet *prevention*; the act of averting loss before it happens, remains uncounted. In sectors like solar O&M, predictive analytics routinely lower fault frequency, shorten repair times, and reduce resource waste. However, existing accounting frameworks capture only what *did* happen, never what *did not*. The result is a blind spot where proactive reliability delivers real savings but receives no recognized credit.

As one O&M technician remarked, “We only get credit when things break; never for the days everything runs perfectly.”

The **Prevention Credit Framework** seeks to close that gap. Inspired by the evolution of the carbon-credit market, it formalizes avoided operational loss as a verifiable economic quantity. Each credit represents a statistically confirmed reduction in expected loss relative to a defined baseline, adjusted for confidence and attribution.

This theoretical foundation transforms prevention from a narrative claim into a measurable, tradable unit of value; anchored in data, governed by transparent rules, and extensible across energy, manufacturing, and infrastructure domains.

The sections that follow derive the governing equation, describe measurement and governance protocols, and propose a pilot implementation pathway within SkinnyCowboy.ai’s **OpenRoper AI Strategy Ecosystem**. Together they establish the groundwork for a new standard: a market that rewards prevention as tangibly as production. This paper builds the technical and governance foundation for an economy that finally prices prevention.

White Paper v1.0 | DOI: https://doi.org/10.65331/wp-CEKD7 | October 2025 |

2 Theoretical Foundation

2.1 From Intuition to Equation

Every operational environment contains a measurable *hazard rate*; the statistical likelihood that a fault, failure, or costly event will occur within a given time window.

Let the baseline rate of such events be $\boldsymbol{\lambda_0(t)}$. When an intelligent preventive system is introduced; whether an AI model, automated workflow, or predictive maintenance rule; the new rate becomes $\boldsymbol{\lambda_1(t)}$. The fractional reduction between the two rates defines the *effectiveness* of prevention:

Eq. 2: Hazard Reduction Model

$$\lambda_1(t) = \theta \lambda_0(t), 0 < \theta \leq 1$$

Defines the fractional effectiveness of prevention: the post-intervention hazard rate is θ times the baseline rate. A value of $\theta = 0.7$, for example, means the intervention lowers expected events by 30 percent.

2.2 Expected Prevented Events

The expected number of prevented incidents equals the time integral of the difference between baseline and post-intervention hazard rates.:

Eq. 3: Expected Prevented Events

$$E[\Delta N] = \int (\lambda_0(t) - \lambda_1(t)) dt$$

This expression converts prevention from a qualitative claim ("we had fewer faults") into a quantitative estimate of *how many incidents were avoided*.

White Paper v1.0 | DOI: https://doi.org/10.65331/wp-CEKD7 | October 2025 | © SkinnyCowboy.ai |

2.3 Translating Incidents into Value

Each avoided incident carries an expected cost or loss magnitude, $\boldsymbol{\mu_\gamma}$, for example, the average truck-roll cost, production loss, or penalty exposure per event.

Eq. 4: Expected Avoided Loss

$$E[\mathit{Avoided\ Loss}] = \mu_y E[\Delta N]$$

The expected avoided loss is the product of the mean loss per event and the expected number of prevented events. This formulation maintains statistical rigor by treating event reduction as a random variable.

Eq. 5: Net Preventive Benefit

$$B_{net} = E[\mathit{Avoided\ Loss}] - K$$

Subtracting the incremental cost K of running the preventive system (edge AI inference, human review, false-positive handling) yields the *net preventive benefit*:

2.4 Adjusting for Real-World Uncertainty

Real operations involve multiple overlapping programs, noisy data, and imperfect verification. To ensure each issued credit reflects **only verified, attributable, conservative value**, three coefficients are applied:

Table 1: Governance Coefficients and Descriptions

Symbol	Meaning	Typical Range
a	Attribution factor – share of	0 – 1

Symbol	Meaning	Typical Range
	benefit assignable to the specific intervention	
q	Verification confidence – probability that the measured reduction is genuine	0 – 1
h	Haircut – conservative deduction for model or data uncertainty	0 – 0.5

Applying these adjustments produces the **Prevention Credit equation**:

Eq. 6: Risk-Adjusted Prevention Credit Value

$$PC = a \cdot q \cdot (1 - h) \cdot B_{net}$$

This equation clearly shows how the *governance coefficients* (a, q, h) scale the net preventive benefit (B_{net}) to yield the final verified Prevention Credit (PC).

Each term is observable or estimable from existing CMMS and SCADA data, making the equation operationally verifiable.

2.5 Interpretation

Together they define a *risk-adjusted expected value of prevention*; a quantity that can be recorded, audited, and ultimately traded as a Prevention Credit.

White Paper v1.0 | DOI: https://doi.org/10.65331/wp-CEKD7 | October 2025 |

Table 2: Interpretation of Core Parameters in the Prevention Credit Equation

Symbol	**Meaning/Interpretation**
λ_0	captures baseline operational risk.
θ	expresses intervention performance.
μ_Y	monetizes each prevented event.
K	enforces economic realism.
$a\ q\ (1-h)$	ensures integrity and comparability across deployments.

White Paper v1.0 | DOI: https://doi.org/10.65331/wp-CEKD7 | October 2025 |

3 Operationalization & Measurement

3.1 From Equation to Evidence

The Prevention Credit framework becomes actionable when each variable in the equation is tied to data already captured inside an operator's existing systems. Modern field operations generate millions of structured records; from SCADA, CMMS, ERP, and ticketing tools that describe every event the model references. By linking these datasets, prevention can be *observed, verified, and priced* without adding new instrumentation.

3.2 Primary Data Inputs

Table 3: Primary Data Inputs and Sources.

Data Source	**Key Fields**	**Parameter Mapping**
SCADA / Telemetry	Timestamped alarms, inverter status, power output, irradiance	Baseline and observed hazard rates ($\boldsymbol{\lambda_0}$, $\boldsymbol{\lambda_1}$)
CMMS / Ticket Logs	Fault type, severity, downtime minutes, crew hours, part cost	Event frequency and loss magnitude ($\boldsymbol{\mu_Y}$)
Financial / ERP	Labor rates, travel cost, replacement part cost	Cost basis ($\boldsymbol{\mu_Y}$, $\boldsymbol{K}$)
Preventive System Logs	Alerts, model confidence, false positives	Verification and attribution ($\boldsymbol{q}$, $\boldsymbol{a}$, $\boldsymbol{h}$)

All data types shown are already recorded in standard O&M workflows; the framework re-uses rather than expands the data burden.

Each input is measurable with existing infrastructure; the framework simply re-interprets these fields through a risk-reduction lens.

White Paper v1.0 | DOI: https://doi.org/10.65331/wp-CEKD7 | October 2025 |

3.3 Baseline and Intervention Modeling

To compute λ_0 and θ:

Baseline Hazard (λ_0):
Estimate the baseline hazard rate using pre-deployment data or matched control assets. Techniques may include Poisson regression, a Cox proportional-hazard model, or a simple moving-average of historical incident rates.

Treatment Effect (θ):
Quantify the fractional change between pre- and post-intervention rates:

$$\theta = \lambda_1 / \lambda_0 = 0.09/0.12 = 0.75$$

Example: The post-intervention rate is 75 % of the baseline rate, representing a 25 % reduction in event frequency. Where λ_1 is the observed hazard rate after deployment and λ_0 is the baseline rate before deployment.

Confidence (q):
Derived from statistical significance testing or Bayesian posterior probability that $\theta < 1$, indicating a true reduction in event frequency.

Attribution (a):
Determined through causal analysis or Shapley-style feature attribution when multiple programs overlap.

Haircut (h):
A policy parameter, typically 10–30 %, applied to ensure conservative valuation of the credited benefit.

3.4 Worked Example

Table 4: Worked Example Variables.

Variable	Description	Value
$\int \lambda_0 \, dt$	Baseline incidents / month	12
θ	Rate reduction factor	0.70
μ_y	Mean loss per incident	$5000

White Paper v1.0 | DOI: https://doi.org/10.65331/wp-CEKD7 | October 2025 |

Variable	Description	Value
K	Operating cost	$2000
a	Attribution	0.6
q	Verification confidence	0.8
h	Haircut	0.1

Expected Prevented Events

$$E[\Delta N]=(1-0.70)\times 12=3.6$$

Avoided Loss

$$Avoided\ Loss=3.6\times 5000=18000$$

Net Preventive Benefit

Net Benefit = 18000 − 2000 = 16000

Risk-Adjusted Prevention Credit Value

$$PC=0.6\times 0.8\times(1-0.1)\times 16000=6912$$

Result:
$6912 of verified, risk-adjusted prevention value issued as one month of Prevention Credits.

3.5 Data Integrity & Verification

Each PC issuance requires:

White Paper v1.0 | DOI: https://doi.org/10.65331/wp-CEKD7 | October 2025 |

1. Immutable time-stamped logs for baseline and intervention data.
2. Independent replication or automated cross-check of rate calculations.
3. Storage of all assumptions (λ_0 window, θ estimation method, haircut value).
4. Optional third-party or community validation prior to tokenization or monetization.

These evidence requirements feed directly into the five-stage governance protocol in Section 4.

This procedure ensures that Prevention Credits are **auditable assets**, not inferred claims; anchoring prevention value in the same evidentiary rigor used for energy generation or emissions accounting.

4 Governance & Verification Protocol

4.1 Purpose of Governance

For Prevention Credits to be trusted, each issued unit must be **traceable, reproducible, and policy-compliant**. Governance provides the rule set that guarantees integrity across different fleets and operators. It defines *how data are verified*, *who attests to them,* and *what qualifies a valid Prevention Credit*. Without this layer, the metric would remain theoretical; with it, PCs become auditable assets.

4.2 Verification Stages

Prevention Credit issuance follows a transparent five-stage sequence, mirroring reliability and ESG audit practices:

1. **Baseline Registration (Compass):**
   Establish pre-deployment performance targets, asset scope, and reference datasets. Record λ_0 and corresponding cost structures.

2. **Blueprint Model Validation (Blueprint):**
   Document algorithms, inference logic, and data-quality checks that will estimate θ and μ_γ. Specify confidence and attribution methods before deployment.

3. **Operational Monitoring (TrailMap):**
   Continuously collect CMMS + SCADA data, track live incidents, and record system alerts and human interventions.

4. **Credit Calculation (FieldAssistant):**
   Apply the verified equation to observed data, compute adjusted net benefit, and generate a verifiable PC ledger entry.

5. **Governance Review (StrategyBuilder):**
   Independent or automated review of parameters (a, q, h, K), reproducibility test, and sign-off before issuance or market listing.

White Paper v1.0 | DOI: https://doi.org/10.65331/wp-CEKD7 | October 2025 |

Figure 1: Governance at Every Stage within the OpenRoper AI Strategy Ecosystem.

*Each Prevention Credit passes through a five-stage governance loop: **Compass** defines value and alignment, **Blueprint** quantifies ROI, **TrailMap** designs the proof of performance, **FieldAssistant** captures operational intelligence, and **StrategyBuilder** governs verification and scaling. Governance remains embedded across all five stages, linking quantitative models to transparent, auditable execution*

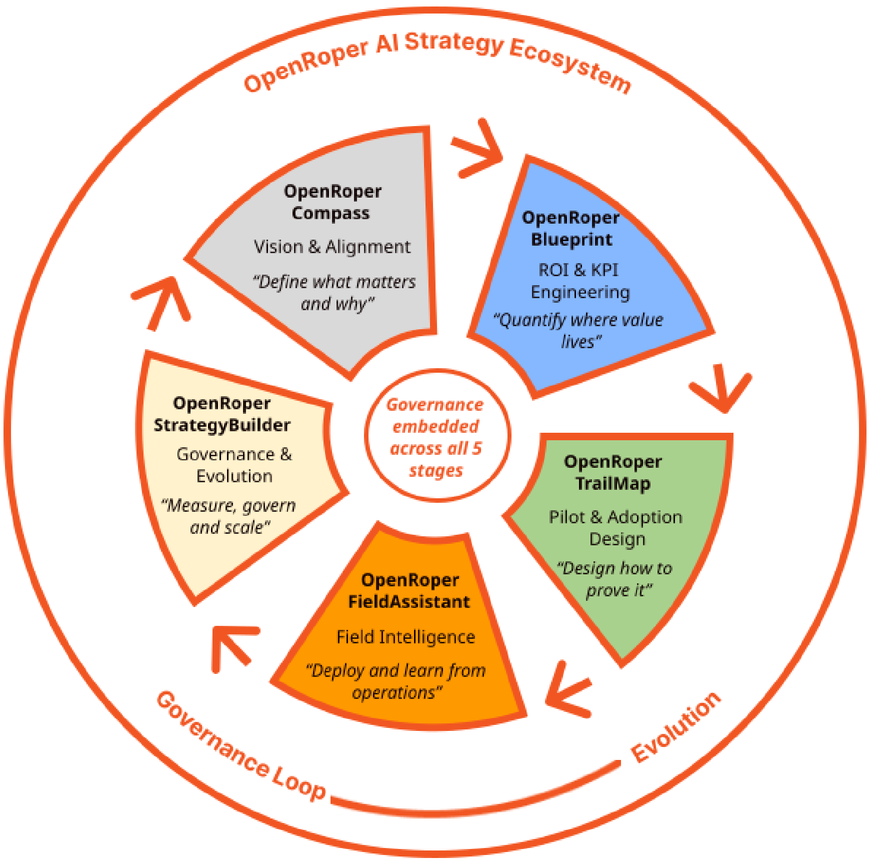

Each layer writes metadata to a distributed, immutable log, enabling full lifecycle traceability from source data to issued credit.

4.3 Minimum Verification Requirements

The table below summarizes mandatory controls ensuring each PC meets audit and disclosure expectations consistent with ISO 42001.

Table 5: Minimum Verification Requirements.

Category	Requirement	Purpose
Data Integrity	Raw logs retained ≥ 24 months, checksum verified	Reproducibility
Statistical Confidence	90 % posterior probability that $\theta < 1$	Scientific credibility
Attribution Disclosure	Statement of overlapping programs	Transparency
Governance Haircut (h)	≥ 10 % default reduction	Conservatism
Audit Trail	JSON or PDF certificate linked to PC ID	External verification

4.4 Ethical & Policy Alignment

Governance aligns with emerging AI-risk and sustainability standards, including:

- **ISO 42001 - AI Management Systems**
- **ISO 31000 - Risk Management**
- **LF Energy Governance Principles**
- **EU AI Act Transparency Provisions**

White Paper v1.0 | DOI: https://doi.org/10.65331/wp-CEKD7 | October 2025 |

By embedding these controls directly in each layer of the OpenRoper ecosystem, prevention value becomes not just measurable, but **trustworthy and compliant**.

4.5 Outcome

A governed Prevention Credit carries the same audit weight as a kilowatt-hour or a verified carbon credit. It transforms "model accuracy" into *market confidence*; a signal that prevention has been proven, priced, and permanently recorded.

Section 5 operationalizes this protocol through a live pilot to demonstrate repeatable issuance and audit results.

White Paper v1.0 | DOI: https://doi.org/10.65331/wp-CEKD7 | October 2025 |

5 Pilot Demonstration Framework

5.1 Purpose of a Pilot

The first Prevention Credit pilot is intended to **demonstrate feasibility and credibility**. It validates three things:

1. That the governing equation produces stable, reproducible results on live operational data.
2. That the governance workflow (Compass → Blueprint → TrailMap → FieldAssistant → StrategyBuilder) functions end-to-end.
3. That the verified credits correspond to recognizable business value; downtime avoided, truck rolls reduced, or dollars saved.

The pilot's output is not a token or tradeable asset yet; it is a **proof certificate** showing that Prevention Credits can be measured and audited with real data. Successful validation of all metrics will constitute the first empirical issuance of a Prevention Credit under governed AI operations.

5.2 Candidate Environment

Select a medium-sized solar O&M operator or asset owner with:

- At least **100–400 MW** of monitored PV capacity.
- SCADA and CMMS data history of ≥ **12 months**.
- Active use of predictive maintenance or AI alerting systems.
- Willingness to open anonymized logs for analysis.

This scale yields sufficient statistical power while keeping the governance overhead manageable.

White Paper v1.0 | DOI: https://doi.org/10.65331/wp-CEKD7 | October 2025 | © SkinnyCowboy.ai |

5.3 Methodology

1. **Baseline Definition (3 months prior data)**
   – Estimate λ_0 from recorded incidents per inverter or subsystem.
   – Quantify average cost per incident (μ_Y) from maintenance and production records.

2. **Intervention Window (3 months post-deployment)**
   – Deploy OpenRoper preventive models and enable alert logging.
   – Record observed incident rate λ_1 and compute $\theta = \lambda_1 / \lambda_0$.

3. **Validation & Attribution**
   – Run Difference-in-Differences or Bayesian rate-ratio tests to derive q.
   – Assign attribution based on overlapping initiatives.
   – Apply default haircut $h = 0.1$ unless statistical uncertainty > ±10 %.

4. **Credit Computation**
   – Calculate PC using the standard equation.
   – Generate audit packet (JSON + PDF) including data sources, model code, and summary statistics.

5. **Governance Review & Publication**
   – Independent review through the StrategyBuilder board.
   – Publish an anonymized summary and invite peer replication.

5.4 Success Criteria

Table 6: Pilot Success Metrics.

Metric	Target	Validation Method
Rate Reduction $(1 - \theta)$	≥ 20 % reduction in incident frequency	Statistical test ($p < 0.05$ or $q \geq 0.9$)
Verified Avoided Loss	Positive net benefit $(\mu_Y(1-\theta)\int\lambda_0 dt - K > 0)$	Accounting cross-check

White Paper v1.0 | DOI: https://doi.org/10.65331/wp-CEKD7 | October 2025 |

Metric	Target	Validation Method
Audit Pass Rate	100 % of sample records reproducible	Independent replay
Stakeholder Acceptance	CFO and O&M manager sign-off	Qualitative

Meeting these thresholds demonstrates that Prevention Credits can be **measured, governed, and valued** with the same confidence as existing performance metrics. A positive net-benefit validation provides CFO-grade financial evidence of prevention ROI.

5.5 Pilot Deliverables Summary

- Baseline vs. intervention dataset (cleaned and documented).
- Credit-calculation report with equations and confidence intervals.
- Governance certificate issued by StrategyBuilder.
- ROI summary showing avoided cost and potential PC valuation.
- Public "Pilot Summary Paper" suitable for submission to LF Energy, ASES, or IEEE PES.

5.6 Next Steps

Upon successful pilot completion, the methodology can scale to multiple fleets, forming the foundation of a **Prevention Credit Registry**. Each verified pilot adds empirical weight to the theory introduced in Sections 2–4, gradually transforming it from concept to accepted accounting instrument.

These pilots will seed the empirical corpus required for an open Prevention Credit Registry and standard-setting collaboration with LF Energy.

White Paper v1.0 | DOI: https://doi.org/10.65331/wp-CEKD7 | October 2025 |

6 Market Implications & Future Work

6.1 A New Asset Class for Operational Reliability

Just as carbon credits turned *emission reductions* into economic value, Prevention Credits transform *verified risk reductions* into a measurable asset. They do not replace existing performance or insurance metrics; they complement them by pricing what traditional accounting omits; **the value of what never happened**.

For operators, each verified PC represents avoided cost and improved uptime. For insurers, it functions as an actuarial signal of lower risk exposure. For investors, it provides an ESG-aligned measure of proactive stewardship. This convergence of engineering reliability and financial accountability establishes a **new asset class**: tradable, auditable prevention value.

Table 7: Comparison of Credit Types.

Attribute	Carbon Credit	Efficiency KPI	Prevention Credit
Basis	Emissions reduced	Output ratio	Loss avoided
Verification	External registry	Internal audit	Governed operational data
Financial Role	Offset liability	Report metric	Tradable reliability value

6.2 Integration Pathways

1. **Carbon and Engergy Markets**
   Prevention Credits can pair with carbon programs by quantifying operational efficiencies that reduce energy waste or standby losses, producing dual-benefit credits (CO_2 + prevention).
2. **Insurance and Risk Finance**
   PCs offer an evidence base for performance-linked insurance premiums and resilience bonds. Verified avoided losses could be discounted directly into underwriting models.
3. **AI Governance and Compliance**
   Because each PC is derived from traceable AI outputs and governed audits, it aligns naturally with ISO 42001 and emerging AI Act requirements for transparency and accountability.

White Paper v1.0 | DOI: https://doi.org/10.65331/wp-CEKD7 | October 2025 |

4. **Enterprise Reporting**
   PCs can be recorded as *"Verified Avoided Loss"* line items in ESG or sustainability reports, connecting technical reliability improvements to financial statements.
5. **Accounting and ESG Disclosure**
   Prevention Credits can appear as "Verified Avoided Loss (VAL)" within ESG disclosures under emerging SEC climate-risk guidance.

6.3 Standardization and Open Collaboration

To achieve adoption, the methodology will evolve through open collaboration. Next steps include:

1. Publishing this white paper under a Creative Commons license for peer review.
2. Launching an **Open Prevention Credit Registry** prototype to test data-exchange and verification APIs.
3. Engaging standards bodies; LF Energy, ISO/TC 207, and IEEE PES, to explore harmonized data schemas.
4. Forming a cross-industry **Advisory Consortium** of operators, insurers, and regulators to define governance thresholds and auditing norms.

These actions will transform the framework from company innovation to community standard.

By Q2 2026, the team will submit pilot data schema to LF Energy, followed by a registry API draft in Q4 2026.

6.4 Future Research Directions

Future work will expand beyond solar O&M to any domain where predictive maintenance or AI-driven foresight prevents loss; wind, manufacturing, grid assets, logistics, and healthcare. Research priorities include:

- Refining attribution models to separate overlapping interventions.
- Developing confidence calibration metrics for low-frequency events.
- Building an open-source verification engine with reproducible audit trails.

- Exploring tokenized market mechanisms under regulatory guidance.

Each pilot and publication will reduce uncertainty, strengthening the empirical foundation of the Prevention Credit as a standard measurement of avoided risk.

6.5 Conclusion

Prevention Credits make the unseen measurable. They convert reliability, foresight, and governance into economic language, bridging the gap between operational excellence and financial recognition. When prevention becomes provable, it becomes valuable; when it becomes valuable, the market rewards it.

This paper establishes the theoretical and operational framework for that transition; the coming phase of pilots, registries, and peer collaboration will determine how rapidly prevention joins production and carbon as the third verified currency of sustainable industry.

White Paper v1.0 | DOI: https://doi.org/10.65331/wp-CEKD7 | October 2025 |

7 Transparency & Authorship

This paper was conceived, authored, and directed by **Mark Ivankovich**, Founder of **SkinnyCowboy.ai** (Boerne, TX, USA).

All conceptual foundations, including the definition of Prevention Credits, the economic equation for verified avoided loss, and the proposed governance structure, originated with the author.

Drafting assistance was provided by AI-based language tools for clarity and structure. The author retains full responsibility for the content, equations, and conclusions presented herein.

To promote reproducibility and community trust, the author encourages open peer review and independent replication of all calculations and pilot results.

This work will be released under a **Creative Commons Attribution 4.0 International License (CC BY 4.0)**, permitting sharing and adaptation with attribution to the original author.

Acknowledgment

The author thanks the early collaborators, field engineers, and advisors who provided feedback on the operational realities that inspired the Prevention Credit framework, and the open-source and standards communities whose prior work in reliability, risk, and AI governance laid the groundwork for this new domain.

White Paper v1.0 | DOI: https://doi.org/10.65331/wp-CEKD7 | October 2025 |

8 Appendices

Appendix A – Data Schema & Entity-Relationship

This appendix defines how operational data maps to the Prevention Credit (PC) equations and how records are linked for auditability. The schema is designed to reuse existing O&M data (SCADA, CMMS, ERP) and adds only the minimal metadata needed for verification and issuance.

A.1 Scope & Versioning

- Schema name: `PreventionCredit.v1`
- Primary objectives: (1) quantify λ_0, λ_1, θ, μ_γ, K; (2) preserve provenance for a, q, h; (3) produce an immutable, replayable PC record.
- Record keys: use stable, opaque IDs (UUIDv4 or ULID). Timestamps are UTC ISO-8601.

A.2 Core Entities

Table 8: Entity Schema and Equation Mapping.

Entity	Key Fields (type)	Purpose / Equation Link
Asset	asset_id(ID), name(str), type(enum: inverter/tracker/…), capacity_mw(num), latitude(num), longitude(num), operator_id(ID)	Groups incidents & telemetry to estimate λ_0, λ_1.
Incident_Log	incident_id(ID), asset_id(ID), ts_start(ts), ts_end(ts), fault_code(str), severity(enum), downtime_min(num)	Source of event counts and timing → $\lambda(t)$, $\int\lambda_0 dt$.
Telemetry_Record	telemetry_id(ID), asset_id(ID), ts(ts), power_kw(num), irradiance_wm2(num), temp_c(num), status(str)	Validates operating state; aligns incidents with conditions.
Preventive_Action	action_id(ID), asset_id(ID), model_id(ID), ts_alert(ts), confidence(num 0–1), action_type(enum), acknowledged_by(user),	Evidence for attribution a and verification q

Entity	Key Fields (type)	Purpose / Equation Link
	ts_resolution(ts), outcome(enum)	(posterior that $\theta<1$).
Cost_Record	cost_id(ID), asset_id(ID), incident_id(ID? nullable), labor_hours(num), labor_rate_usd_hr(num), parts_usd(num), travel_usd(num), lost_prod_usd(num), period(month)	Aggregates to μ_γ (mean loss/event) and operating cost K.
PC_Run	run_id(ID), period_start(ts), period_end(ts), method_ref(str), lambda0_method(str), theta_method(str), haircut_h(num), ver_conf_q(num), attrib_a(num)	Parameter snapshot ensuring repeatability.
Prevention_Credit	pc_id(ID), run_id(ID), operator_id(ID), portfolio_id(ID?), a(num), q(num), h(num), mu_y(usd), theta(num), lambda0_events(num), k_usd(usd), b_net_usd(usd), pc_value_usd(usd), issued_at(ts), hash(str)	Final auditable record: PC = a·q·(1−h)·B_net.
Audit_Attachment	attachment_id(ID), pc_id(ID), kind(enum: json/pdf/code/log), uri(str), sha256(str)	Links artifacts (code, notebooks, logs) to the PC.
Operator	operator_id(ID), name(str), region(str)	Owner context for access and reporting.

Note: PII is not required; use role accounts (e.g., “O&M-Tech-42”) for acknowledged_by.

A.3 Relationships (text ER diagram)

```
Operator 1———< Asset 1———< Incident_Log
          |          |
          |          └———< Cost_Record >———1 Asset
          |
          ├———< Telemetry_Record
```

White Paper v1.0 | DOI: https://doi.org/10.65331/wp-CEKD7 | October 2025 | © SkinnyCowboy.ai |

└———< Preventive_Action ——> Model (external registry)

PC_Run 1———< Prevention_Credit >———1 Operator
Prevention_Credit 1———< Audit_Attachment

Incident_Log (optional) ——< Cost_Record (when costs are incident-scoped)

Cardinalities

- One Asset has many Incidents, Telemetry, Actions, Costs.
- One PC_Run can issue many Prevention_Credits (e.g., per site, per month).
- Each Prevention_Credit may have many Audit_Attachments.

A.4 Field-Level Mapping to Equations

Eq. 2 ($\lambda_1=\theta\lambda_0$):
- Incident_Log counts per period (pre vs. post) → estimate λ_0, λ_1.
- PC_Run.theta_method stores model (Cox, DiD, Bayesian rate ratio).

Eq. 3 ($E[\Delta N]=\int(\lambda_0-\lambda_1)dt$):
- lambda0_events = $\int\lambda_0 dt$ (baseline count in window).
- theta saved to Prevention_Credit.theta.

Eq. 4 ($E[\text{Avoided Loss}]=\mu_Y \cdot E[\Delta N]$):
- mu_y derived from Cost_Record (labor+parts+travel+lost_prod)/events.

Eq. 5 (B_net=E[Avoided Loss]−K):
- k_usd from Cost_Record tagged operating_cost=true for the window.
- b_net_usd computed and stored in Prevention_Credit.

Eq. 6 (PC=a·q·(1−h)·B_net):
- a,q,h stored in both PC_Run (defaults) and Prevention_Credit (effective).
- pc_value_usd computed and hashed.

White Paper v1.0 | DOI: https://doi.org/10.65331/wp-CEKD7 | October 2025 |

A.5 Constraints, Indexes, and Quality Gates

Primary Keys

- *_id fields are unique, immutable.

Foreign Keys

- Incident_Log.asset_id → Asset.asset_id
- Telemetry_Record.asset_id → Asset.asset_id
- Cost_Record.asset_id → Asset.asset_id (and optional incident_id → Incident_Log.incident_id)
- Prevention_Credit.run_id → PC_Run.run_id
- Prevention_Credit.operator_id → Operator.operator_id
- Audit_Attachment.pc_id → Prevention_Credit.pc_id

Indexes (recommended)

- Incident_Log(asset_id, ts_start)
- Telemetry_Record(asset_id, ts)
- Cost_Record(asset_id, period)
- Prevention_Credit(issued_at, operator_id)

Quality Gates

White Paper v1.0 | DOI: https://doi.org/10.65331/wp-CEKD7 | October 2025 |

- Completeness: no nulls for keys, timestamps, or monetary fields.
- Temporal coherence: ts_start ≤ ts_end; all records within PC_Run.period_start/period_end.
- Double-count protection: unique (asset_id, ts_start, fault_code) for incidents; dedup costs.
- Attribution evidence: if a > 0, Preventive_Action or design documentation must exist in Audit_Attachment.

A.6 Provenance, Hashing, and Retention

- Provenance: Each entity includes `source_system` (enum: SCADA/CMMS/ERP/Manual), `ingested_at`(ts).
- Hashing: `Prevention_Credit.hash = sha256(concat(sorted(core fields)))`. Store the same hash in the certificate PDF/JSON.
- Retention: Raw logs retained ≥24 months; derived tables ≥36 months; certificates permanent.
- Reproducibility: `PC_Run.method_ref` points to a tagged code commit or notebook; attach as `Audit_Attachment(kind=code)`.

A.7 Sample Prevention Credit (JSON certificate)

```
{
  "pc_id": "pc_01JABC4V9K2Z3",
  "operator_id": "op_6F2D",
  "period_start": "2025-08-01T00:00:00Z",
  "period_end": "2025-08-31T23:59:59Z",
  "asset_scope": ["asset_A12","asset_A34","asset_B07"],
  "theta": 0.70,
  "lambda0_events": 12,
  "mu_y_usd": 5000,
  "k_usd": 2000,
  "a": 0.6,
  "q": 0.8,
  "h": 0.1,
```

White Paper v1.0 | DOI: https://doi.org/10.65331/wp-CEKD7 | October 2025 |

```
  "b_net_usd": 16000,
  "pc_value_usd": 6912,
  "method_ref": "pc-methods@v1.3.2",
  "attachments": [
    {"kind":"code","uri":"s3://.../notebook.ipynb","sha256":"..."},
    {"kind":"logs","uri":"s3://.../incident.csv","sha256":"..."}
  ],
  "issued_at": "2025-09-05T12:13:14Z",
  "hash": "a9b1...f2c4"
}
```


A.8 Textual Entity-Relationship Overview

The **Operator** entity owns one or more **Assets**.

Each **Asset** generates multiple **Incident_Log**, **Telemetry_Record**, **Preventive_Action**, and **Cost_Record** entries.

A **PC_Run** captures the parameterization and methods for a defined calculation window.

From each **PC_Run**, one or more **Prevention_Credit** records are issued, each linked to one-to-many **Audit_Attachment** entries that store the supporting artifacts required for independent verification and replay.

A.9 Security & Access Notes

- No customer PII required; use role accounts.
- Row-level security by `operator_id`.
- Cryptographically sign certificates if exchanged externally.

A.10 Minimal Migration Checklist (for pilots)

1. Map SCADA + CMMS exports to Incident_Log and Telemetry_Record.
2. Import labor/parts/production to Cost_Record; compute μ_γ and K.
3. Register code + parameters in PC_Run.
4. Produce Prevention_Credit and attach artifacts.
5. Verify hash and archive certificate.

White Paper v1.0 | DOI: https://doi.org/10.65331/wp-CEKD7 | October 2025 |

Appendix B – Glossary of Terms

- ΔN — Number of prevented incidents
- PC — Prevention Credit
- PCU — Prevention Credit Unit (standardized accounting unit of PC value)
- VAL — Verified Avoided Loss

White Paper v1.0 | DOI: https://doi.org/10.65331/wp-CEKD7 | October 2025 |

Appendix C – References and Standards

- ISO 42001 [1] 2023 — *Artificial Intelligence — Management System Standard*
- ISO 31000 [2] 2018 — *Risk Management Guidelines*
- SASB Standards (2025 Update — ISSB Integration Draft) [3]
- SEC Climate-Related Disclosure Rules (2024 Final — Pending Implementation) [4]
- LF Energy Foundation [5] — Governance Framework and Charter (undated) [5]
- Formal mathematical derivations are available in the supplemental technical documentation maintained by SkinnyCowboy.ai (Technical Supplement, forthcoming).

White Paper v1.0 | DOI: https://doi.org/10.65331/wp-CEKD7 | October 2025 |