Cortex cathodic-protection project

Technical critique and development review

18 October 2026 | Review draft 1 | Prepared for the project team

DRAFT - awaiting CP4 review. Not approved for engineering acceptance or active scoring policy.

The project is useful as a traceable review queue. It is not yet a validated measure of protection adequacy or overall system health. The calculations can be reproduced. The larger issue is whether the measurements, formulas and scoring choices support the conclusions people may draw from the grades.

I would keep the source-linked workbook, the separate source conclusions, the branch summaries and the knowledge repository. Before relying on the grades for decisions, I would resolve the measurement dates, make the calculation dependencies visible, and put an explicit engineering decision ahead of the numeric score.

The findings that matter most

- **The decay calculation joins different survey dates.** Off is under 7 June 2023; Depol is under 28 July 2024. Their difference is calculable, but the available records do not establish a valid polarization-decay sequence.
- **Several factors reuse the same evidence.** Required current and resistance are spreadsheet formulas. Following those formulas, Off is an upstream input to 85% of the rubric's nominal weight. Seven factors do not represent seven independent observations.
- **A better grade can result from less evidence.** A controlled example rises from 66/B to 88/A when an unfavorable decay input is removed. That is a denominator effect, not demonstrated improvement in protection.
- **The letter grade and source conclusion can disagree substantially.** Thirty-three source-Fail records receive B grades. Thirty-four source-Fail records display "No rule flag." These are separate findings, not necessarily the same records.
- **System means describe a partial set of records.** Fuel Gas Supply's 64.1/B uses 23 of 36 records; CCW's 53.8/C uses 47 of 94. Neither mean demonstrates that all connected assets are protected.

How to use this review

Pages 2-5 explain the engineering and scoring issues. Pages 6-7 propose the missing-measurement workflow. Pages 8-9 address the CP4 example and combined systems. Pages 10-13 cover Jev, the wiki, corrective work and review decisions. Pages 14-16 contain sources and reproducibility details.

Scope: original local workbook; current scorer and loader; saved Jev results; system and survey summaries; supplied report screenshot; repository and sharing design. Calculations were checked with read-only audits and isolated synthetic cases. No source cells, production grades, active rules or model judgments were changed, and no fresh Jev evaluation was made for this critique.

This is an AI-assisted technical review for a qualified cathodic-protection specialist to challenge and complete. It is not CP4 certification, a signed engineering report, a field procedure or authorization to alter a protection system.

Evidence notation: **[L]** identifies local files, cells or reproducible tests; **[S]** identifies external primary sources. Proposed rules are recommendations for review, not current runtime behavior.

2. What the engineering decision actually requires

Cathodic protection controls electrochemical corrosion of a metal exposed to an electrolyte by supplying protective current through galvanic anodes or an impressed-current system. Coatings reduce the demand for that current. Geometry, coating condition, resistivity and electrical connections affect where current reaches the structure. A source producing current does not, by itself, show that every relevant surface receives adequate protection. Excessively negative potentials can also cause material or coating problems. Shielding and interference require attention beyond a single meter reading. These are basic physical reasons to retain field context. [S1, sections 3-4 and 6]

A criterion needs a defined measurement basis

For covered ferrous gas-pipeline applications, Appendix D includes alternative criteria with specified measurement conditions. Its -0.85 V provision identifies a saturated copper/copper-sulfate reference and requires consideration of other voltage drops. Its polarization-decay method starts from the potential after the immediate interruption shift. An "Off" label or the subtraction of two readings does not establish those conditions. The regulation is not a universal rule for every industrial plant system. [S2, I.A, II-IV; S3]

Current DoD O&M guidance discusses reference/meter error, electrolyte voltage drop, anode gradients, contact resistance and mixed-metal effects. It also identifies limitations on polarization criteria, including bimetallic conditions. Its March 2026 edition is informational guidance, not automatically binding on this project. [S4, foreword; 2-1.1; 2-1.2.4]

Information this project still needs

Recommended asset-level record: material and relevant metallurgy; buried/submerged surface being assessed; service and environment; coating and temperature; galvanic or impressed-current arrangement; electrical continuity and intended isolation; bonds and grounding; current-source boundaries; applicable criterion, edition and company procedure; person responsible for its selection.

Recommended observation-level record: asset and test-point identity; time and location; instrument and reference-electrode basis/check; units and sign; method; interruption state and timing; interfering-source conditions; operator; raw reading; source attachment; validity decision and reason.

These are proposed fields for this project. Their absence in the current sanitized Jev input is a limitation of what that input can support, not proof that the field team lacks the information elsewhere.

Keep four conclusions separate

Conclusion	What it should answer
Criterion assessment	Does valid, applicable evidence satisfy the selected criterion at this location and time?
Equipment / isolation finding	What component or electrical-boundary issue was observed, and how was it confirmed?
Evidence status	What is missing, contradictory, stale or technically unverified?
Action priority	What should be reviewed or investigated first, considering consequence and uncertainty?

A single A-F score currently mixes parts of these questions. It cannot establish corrosion rate, wall thickness, remaining life or fitness for service from the supplied inputs. No such engineering model or validation evidence was found in the reviewed implementation. [L2-L4]

3. What is actually in the workbook

The source hash matches the saved station results, system summary and survey comparison. There are **155 imported source records**, grouped into 22 branch labels and six reporting parents. These are record counts; they are not established counts of unique physical stations. [L1, L3-L5]

Dates are part of the evidence

Source cells, Current Status sheet	Date / meaning	Why it matters
N3:S3; N4 / O4	7 June 2023; On / Off	The existing condition score uses this Off column.
T3:U3; T4 / U4	28 July 2024; Depol / required current	Pairing T with O needs its test history explained.
V3:W3; V4	18 November 2024; On	First later-survey cohort.
X3:AC3; Y4	4 April 2026; On	Second later-survey cohort; not a replacement Off reading.

The labels establish workbook chronology. They do not prove when every individual reading was taken, how interruption occurred, or that the same reference conditions applied. Ask whether these are true acquisition dates, reporting dates or carried-forward data, and retain the answer with its evidence.

Numerical availability under the present code

Factor	Nominal weight	Available / 155	Unavailable
Instant-off	35	81	74
Polarization decay	25	77	78
Isolation	15	133	22
Current adequacy	10	57	98
On consistency	5	81	74
IR contribution	5	81	74
Resistance peer comparison	5	64	91

These counts mean that the software can calculate the factor. They do not certify measurement validity. Only 30 records have all seven factors numerically available. The saved results contain 81 grades and 74 Unscored records: 1 A, 41 B, 31 C and 8 D. [L3, L5-L6]

Source conclusions and identity

All 155 source status cells are formulas: **88 Fail, 7 Pass, 60 blank**. The formula requires a nonblank Depol before returning an outcome, tests Off against -850 mV, and adds a Depol <= -400 mV condition to its alternative shift test. The basis of that additional condition needs a named owner and reference. Preserve the source result; do not assume it is independently adjudicated ground truth. Eighteen source-Fail records lack finite Off values. [L1, S5 pattern; L5]

Forty-six records have blank station IDs; 45 use the connection placeholder "No ID." "No ID - ANO" occurs at R008 and R021 for different records. Nine substantive isolation conflicts exist between source views. R008, for example, is shorted in Current Status and Good in the inventory. Source precedence does not establish which observation is correct or newer. Use source hash plus row reference for this snapshot and build permanent asset IDs for future imports. [L5]

4. The rubric counts related calculations more than once

The current weights are **35/25/15/10/5/5/5**. Available factors are averaged using only their available weights. An overall score requires Off plus at least 60 weight points. Grade bands are A >= 80, B >= 60, C >= 40, D >= 20 and F below 20, using the rounded overall score. These are project rules, not a demonstrated engineering scale. [L2]

The source formulas expose the dependency

For row 5, and the corresponding formula patterns throughout the imported rows:

```
P = On - Off                          [voltage difference]
R = -(On - Off) / current_output      [derived resistance]
U = ((-850 - Depol)/(Off - Depol)) * current_output
                                      [derived required current]
current_output / U = (Off - Depol)/(-850 - Depol)
```

These simplified expressions omit spreadsheet guards. The cancellation applies where the divisions and positive requirement are valid. Units are mV and mA, making their ratio ohms. It does not establish that the resulting R represents the particular circuit resistance the reviewer needs. Q is labeled **system total** current and repeats across some records. [L1, P5/R5/U5/Q4; L5]

Off directly feeds the first, second, fifth and sixth factors: 70% of nominal weight. In this workbook, R and U also depend on Off, increasing the upstream dependency to **85%**. This is a lineage count, not a statistical correlation coefficient or a claim that every score always changes when Off changes.

Each factor needs a more precise interpretation

Current factor	Main critique	Recommended treatment
Off, 35%	Reference, IR and test validity are not eligibility gates. The linear scale below the threshold is unvalidated.	Evaluate an applicable criterion using validated evidence.
Decay, 25%	Cross-date subtraction is treated as a decay factor; the same Off is reused.	Require the sequence and applicability decision.
Isolation, 15%	Good adds points that can offset other concerns; shorted is not an independent mandatory-review trigger.	Keep the finding and its consequence visible.
Current adequacy, 10%	In this source the ratio algebraically reuses Off/Depol; it is not independent capacity evidence.	Label derivation and validate the physical model.
On consistency, 5%	Ordering two readings is a consistency check, not proof of adequate protection.	Move to data-quality checks.
IR contribution, 5%	A small On-Off fraction earns more points without a demonstrated health relationship.	Show as a diagnostic quantity, subject to validity.
Resistance peer, 5%	Similarity to a same-label median is rewarded; the target contributes to its own median.	Use as an anomaly question, not an acceptance criterion.

The current Off, decay and current factors cap at 80, while several secondary factors can reach 100. This permits secondary evidence to lift a letter grade. There is no engineering justification in the reviewed materials for these weights, caps, linear curves or compensation rules. Calibration must test the intended decision, not merely reproduce this arithmetic. [L2, L6]

5. What the reproduced tests show

The unchanged engine reproduces all 155 saved deterministic results. That confirms implementation consistency for this snapshot. It does not validate the meaning of the score. The following examples deliberately distinguish actual records from isolated experiments. [L6]

Observed in the saved results

Finding	Evidence	Interpretation
33 source-Fail records have B grades	R005 and 32 others	A B is not equivalent to the source's Pass.
3 B records miss both draft numeric thresholds	R021 64.5; R032 63.7; R107 72.1	Other factor points compensate. These are code-threshold comparisons, not complete engineering verdicts.
34 source-Fail records say No rule flag	26 C and 8 D records	That label only describes existing rule triggers. It is not human clearance.
All six source-shorted records are Unscored	R008, R159, R160, R162, R172, R173	A shorted B was not observed in this dataset.

Controlled synthetic examples

These examples call the existing scorer with fabricated inputs and controlled peers. They are software demonstrations, not reported field conditions.

Change in the synthetic input	Before	After	What it exposes
Good isolation becomes shorted	86/A	71/B	A zero isolation factor can be compensated.
Shorted isolation is replaced by missing	71/B	83.5/A	Excluding adverse evidence raises the mean.
Zero decay becomes missing	66/B	88/A	Remaining coverage is 75%, so the low-coverage flag is not triggered.
Zero current output becomes missing	78/B	86.7/A	Missingness can remove a penalty.
On/Off move together from -850 to -3000 mV, with shift held at 100 mV	86/A	86/A	There is no excessive-polarization or plausibility gate.

The last case does not prescribe an upper protection limit. It shows a missing decision path. Whether an extreme value is credible or harmful depends on the material, method and environment.

Sensitivity of the real scores

Excluding the explicit 25-point decay factor and renormalizing coverage over the remaining 75 points leaves 81 records eligible at a 60% threshold. **42 letter grades change**; score changes range from -8.2 to +21.6. Keeping the original minimum of 60 weight points instead leaves 77 eligible, with 38 of those changing grade. These are alternative sensitivity conventions, not approved replacement methods. [L5]

Both variants retain the original current-adequacy factor, which still depends on Off and Depol. They therefore do **not** cleanly remove the disputed date-pairing assumption. The proper correction starts with measurement validity and dependency handling, then recalculates under a reviewed rubric while retaining the historical output.

6. An order of operations for missing measurements

The PEMDAS comparison is useful for explaining sequence: check what applies, then follow the applicable route. The important qualification is that an unavailable measurement is not automatically irrelevant. Skipping a term changes what can be concluded.

Proposed workflow: context -> evidence -> validity -> applicable route -> conclusion -> human decision. This is a design proposal. It does not activate any of the repository's draft policies.

Step 1. Establish the question and the asset

Identify the location, electrical boundary, survey event and intended decision. Select the applicable criterion through the responsible specialist. A plant branch name alone cannot choose the measurement method or establish legal applicability.

Step 2. Give the missing input a meaningful state

Proposed state	Meaning	Reporting effect
Recorded, validity unconfirmed	A value exists; technical checks remain open.	Display it, but do not call the criterion established.
Recorded, validated	Required checks for the selected use are satisfied.	Eligible for that use, with method and revision attached.
Unavailable	Access, interruption or another documented condition prevented the measurement.	Preserve reason, dependencies, owner and next route.
Invalid / conflicting	The observation failed checks or sources disagree.	Preserve both versions; require resolution.
Approved not applicable	The measurement is unnecessary for this specific approved route.	Record scope and approver; do not invent a favorable value.
No information	Neither a value nor a defensible explanation is available.	Keep the uncertainty explicit.

These are proposed meanings, to be mapped carefully to the existing schema. A numeric zero remains a reading until its provenance shows otherwise. A blank is never converted to measured zero.

Step 3. Follow dependencies before choosing a fallback

If Off is invalid, revisit decay, On consistency, the IR fraction and this workbook's derived R and U. If Depol is invalid, revisit decay and U. Keeping downstream results while dropping only the visibly missing column can preserve the same invalid assumption under another label.

Step 4. Select an approved evidence route

Interruption difficulties can change survey options. PHMSA's educational CIS material describes On surveys where CP sources cannot be interrupted, while also describing limitations and different test states. This does not make every On reading equivalent to a valid polarized-potential result. [S5]

The specialist may investigate a different supported measurement arrangement or corroborating assessment. A historical DoD design reference discusses coupon installations when interruption is difficult; that is a designed physical arrangement, not a software-estimated replacement reading. Current design requirements must be checked before specifying it. [S6]

Step 5. Explain, assign and allow intervention

Show the route considered, why it was selected or rejected, remaining limits, affected calculations and the next action. The reviewer must be able to correct identity, challenge validity, select a supported route, return the case for measurements or leave it unresolved. Every decision needs a person, date, reason and evidence link. An AI suggestion cannot close its own hold.

7. Measurement-rule worksheet for the team

Draft decisions to complete. Each row needs a named owner, applicable asset scope, approved method, required evidence, fallback conditions, resulting conclusion, expiry/review date and change history. The table proposes how software should preserve uncertainty; it does not authorize field work.

Input / current role	If it cannot be used	Team decision still needed
On potential	Retain unavailable/invalid state; disable dependent comparisons. Do not substitute another year's value silently.	Which survey conditions and descriptive uses are acceptable?
Off potential	Hold criteria that require it. Invalidate its dependent calculations unless separately established.	Which alternate route is suitable, and what evidence makes it valid?
Depol / decay endpoint	Do not infer decay from an unrelated survey. Hold decay and the source-derived requirement calculation.	Required pairing, interruption history, observation interval and acceptance basis.
Isolation status	Keep unknown or conflict visible. Do not award Good points from absence of a shorted label.	Test method, source precedence, electrical-boundary impact and review urgency.
Current output	Preserve measured zero distinctly from unavailable. Record whether it is a station, branch or system total.	Measurement location, method, aggregation boundary and expected operating context.
Required current	Label calculated versus independently established. Show formula and inherited limitations.	Validity of the current requirement model and design assumptions.
Resistance	Identify the actual quantity: this source uses a voltage/current ratio. Suppress an invalid derived result.	Meaning, units, method and comparable population; whether it belongs in grading.
On-Off / IR difference	Calculate only from compatible valid observations; show units and sign.	Whether it is a useful diagnostic and how it informs follow-up.
2024 On	Preserve a one-sided survey observation when the later value is missing.	Asset match and method comparability for trend use.
2026 On	Do not treat it as Off, decay or proof of protection. No delta without a valid pair.	Same as above, including operating and reference conditions.
2026 AC voltage	Keep a separate interference/safety review question. Never infer current density from volts alone.	Applicable AC-corrosion assessment, personnel-safety procedure and competent reviewer.

The federal interference provision distinguishes current-density triggers for specified transmission pipelines; it is not a volts threshold for this workbook. A PHMSA-sponsored laboratory study also shows that AC and CP effects interact under its tested conditions. Neither source supplies a universal plant acceptance limit. [S7-S8]

Minimum explanation shown to a reviewer

"Off could not be established for this survey. The stated reason is [reason and evidence]. This prevents [specific conclusion] and affects [dependent calculations]. The proposed route is [method], subject to [requirements]. Until reviewed, the result remains [status]. Assigned to [person] for [next action]."

Do not fill these placeholders with a model's invented field history. If the reason or owner is absent, say so. A reviewer may approve an applicable alternate route, but a missing value should never silently increase the apparent certainty of the conclusion.

8. Reading the CP4 example fairly

The supplied screenshot is a useful reporting example: it puts branches, survey dates and potential ranges together. The user's identification of its author as CP4 is respected; credentials and the complete signed report were not independently reviewed. The critique below concerns the visible table and the local workbook arithmetic, not the author's unseen methods or overall conclusions. [L7]

"Avg" needs a definition

The displayed values in several checked rows match **midrange = (minimum + maximum)/2**. A mean instead uses every included reading. A midrange is a legitimate descriptive statistic if labeled; it does not summarize the distribution in the same way as a mean.

Branch / statistic	2024 On	2026 On	Paired change, 2026 minus 2024
Fuel Gas Supply - Main: midrange, mV	-1090.5	-824.0	+357.0
Fuel Gas Supply - Main: mean, mV	-1194.7	-769.7	+425.0
Available observations / valid pairs	3	3	3
CCW - Return: midrange, mV	-1202.0	-905.0	+473.5
CCW - Return: mean, mV	-1283.0	-623.2	+663.0
Available observations / valid pairs	27	23	23

Numbers are rounded for display. Numeric minimum means most negative; numeric maximum means least negative. The screenshot's Low/High arrangement appears magnitude-oriented, so explicit labels avoid ambiguity. A positive paired change here means the later On reading is less negative. Its engineering meaning remains dependent on survey conditions. [L4-L5]

Blanks produce misleading deltas unless guarded

There are 104 numeric 2024 readings, 91 numeric 2026 readings and **90 valid same-row pairs**. Three numeric source deltas do not have two numeric inputs:

Cell	Cached delta	Missing evidence
AC45	0 mV	Both survey values blank
AC109	+625 mV	2026 value blank; 2024 is -625 mV
AC140	-1233 mV	2024 value blank; 2026 is -1233 mV

All 90 valid paired source deltas match subtraction. The new survey comparison already excludes the three invalid pairs. For CCW Unit #3 Supply, the sole valid paired change is **+940 mV**. Combining it with the unpaired -1233 gives the screenshot's -146.5 midrange, which is not a valid two-observation paired statistic. Confirm the original report's source and handling with its author before revising that report. [L1, L4-L5, L7]

CCW Unit #1 Supply and Chemical Feed have no numeric reading in either later-survey column. SWS has one 2024 value and none in 2026. Zero-filled summary rows therefore remain uncertain: absence of a numeric source value does not explain why testing was unavailable. Keep yearly cohorts, paired cohorts and unknown reasons separate. These On-to-On trends are not polarization decay and are not substitutes for the current score's older Off evidence.

9. Branches and combined parent systems

Showing both levels is useful. The current summary correctly computes each parent mean from its underlying graded records, rather than giving a tiny branch the same weight as a large one. It also retains unscored counts, source outcomes and station-level escalation indicators. Those controls should remain. [L4]

Current parent results, reproduced without changes

Reporting parent	Draft mean / grade	Graded / records	Unscored	Mean factor coverage	Source Fail
Fuel Gas Supply	64.1 / B	23 / 36	13	65.6%	28
CCW	53.8 / C	47 / 94	47	52.0%	49
CMG	47.1 / C	3 / 11	8	35.0%	5
BDN	64.2 / B	6 / 6	0	92.5%	3
AQA	63.7 / B	2 / 5	3	45.0%	3
SWS	Unscored	0 / 3	3	15.0%	0

"Graded / records" describes eligibility under the present rubric. "Mean factor coverage" averages available nominal weights across all records. Neither measures percentage of pipe length surveyed, percentage of assets protected, or confidence in measurement validity. Source Fail counts remain the workbook formula's outcomes.

What an aggregate can conceal

CCW contributes 94 of 155 records; Return and Supply together contribute 83. A record-weighted mean therefore reflects survey density. It is not weighted by pipe length, failure consequence, exposure, service or unique assets. Repeated system-total inputs also prevent treating every record as an independent sample. [L1, L5]

Fuel Gas Supply is 64.1 using underlying graded records, versus 65.6 if rounded branch means are averaged equally. CCW is 53.8 versus 52.6 under that alternative. Neither weighting scheme becomes an engineering model merely by being explicit. Parent and branch counts describe the same records and must not be added together.

The parent mappings follow labels, not verified electrical connectivity. AQA is retained separately, while a facility description also mentions it within a broader current-total grouping. A process-service group and a CP electrical circuit can require different hierarchies. Store both, with a reviewed crosswalk. [L5]

Proposed final review for each system

- **Scope:** included assets, branches, CP circuit boundaries, survey date and excluded areas.
- **Findings:** criterion decisions at relevant locations, source conclusions, unresolved isolation or interference issues, and supporting records.
- **Evidence coverage:** eligible records, missing/invalid observations, survey footprint when known, and comparability limits.
- **Actions:** cases requiring attention, reason, named owner, due date set by the responsible team, and review status.
- **Optional index:** the clearly labeled draft mean and distribution, always subordinate to unresolved findings. The maximum station Jev urgency remains a station escalation indicator, not a separately evaluated system judgment.

No approved criterion failure, unresolved safety concern or required measurement should disappear because other records improve the average. Define those review gates with the specialist before turning the rollup into an acceptance decision.

10. What Jev contributes, and what still needs control

The current architecture separates deterministic condition calculations from Jev's advisory urgency. That is a useful boundary. The latest processed export reuses previously saved per-record model responses that passed the software's format and range checks; it is not a new Jev run. Connecting repository context can make instructions more consistent, but it cannot validate the underlying engineering assumptions. [L2-L3, L8]

Controls worth keeping

- The outbound record excludes names, identifiers, GPS and arbitrary source notes. It contains selected numeric measurements, factor scores, source outcome and limited signals.
- The knowledge loader validates a curated structured bundle, including version, supported formulas and hashes. A wiki edit does not automatically become a scoring rule.
- The current cache checks model, rubric/version, knowledge fingerprint and prompt contract. The old source-hash-only fallback is gone; it should not be listed as a current defect.
- Invalid or missing model answers prevent a completed export. Provenance and saved outputs permit comparison across versions.

The model still lacks important context

The sanitized record does not carry a validated measurement sequence, reference-electrode basis, asset/material applicability or full electrical topology. These omissions limit what a judgment can mean. Sending more raw client data is not the immediate answer: first define the minimal reviewed context the decision requires.

Five saved responses have confidence below 0.5. That confidence is a model output, not measured probability that the engineering judgment is correct. NIST's voluntary AI Risk Management Framework distinguishes validity for the intended use from general model capability and stresses evaluation, transparency and appropriate human oversight. Those principles support testing this particular workflow against independently reviewed cases. [S9]

Reproducibility needs an explicit policy

Three repeated sanitized-input groups cover 43 records and contain saved urgency spreads of about 1.7-3.3 points. Reusing the last answer for an identical key would change 34 raw urgency values in a source-order simulation. The latest export correctly preserved its per-record saved judgments; the remaining question is whether future equivalent inputs should share one canonical judgment or retain separate sampled answers. Document the choice and compare stable versions. [L6, L8]

The response validator checks ranges and probability sums, but does not enforce consistency between the numeric score and class probabilities. A synthetic contradictory answer passes. The actual saved deviations are small, at most 0.02 on the 0-4 scale, so this is a contract weakness rather than evidence of large current inconsistency. Decide whether the score is derived from probabilities or has an independently defined meaning. [L6]

Explanations and intervention must be first-class outputs

Every recommendation should identify the source observations, applicable rule/version, omitted or invalid inputs, alternative considered, uncertainty and reason for review priority. The current summary fields and probability output do not by themselves provide that full chain. Let a reviewer accept, amend, reject or defer a recommendation with a recorded reason. Preserve the original model answer and create a separate human decision; do not overwrite history or treat an override as automatic proof that a future model should imitate it.

11. The repository and sharing model

The wiki is a practical way to retain decisions, failed approaches, field exceptions and reusable review rules. Markdown is easy to inspect and edit; structured JSON can express a small, validated subset for software. The strongest design choice is the separation between readable knowledge and activated behavior.

Recommended publication and rule states

State	Human-readable wiki	May alter scoring?
Source evidence	Link or controlled attachment with provenance	Only through a reviewed interpretation
Draft proposal	Clearly labeled, with open questions and author	No
Reviewed guidance	Scope, reviewer, evidence and limitations recorded	Only if separately approved for execution
Approved executable rule	Versioned contract, tests, release and rollback	Yes, within its approved scope
Superseded / rejected	Retained with explanation and successor link	No

The reviewed repository contains 11 draft measurement policies that remain disabled and outside the active Jev export. Keep that boundary while the team fills in the missing engineering decisions. Claude and Kimi proposals are drafts to review, not additional independent field observations or approvals. [L9]

For each promoted rule, record the criterion basis, affected assets, prerequisites, unavailable-data behavior, downstream dependencies, explanation text, reviewer, approval date and replacement history. Keep a single canonical structured definition and generate or check its explanatory page, so the Markdown and JSON cannot silently contradict one another.

A privacy statement does not make a public site private

The currently shared wiki and PDF use publicly accessible web routes. That is how the team can open them without an account. A "private project" label or copying restriction does not enforce access control. Direct PDF links and portable ZIP downloads also remain copies once distributed. This is an implementation observation, not a legal opinion about the license. [L10]

Choose a content boundary: public demonstration material, team-restricted project knowledge, and separately controlled client/source data. Confirm redistribution rights before including third-party reports or identifying site data. If the team needs restricted access, authentication must protect the article, direct files and archives; protecting only the landing page is insufficient. This critique has not changed the site's publication settings.

Make the repository teach the system accurately

Historical runs should name their source hash, scorer version, model, knowledge fingerprint, active policy version and timestamp. Newly discovered limitations should be added as reviewed corrections, while older artifacts remain recognizable as historical results. Avoid calling a curated set of exported rules "everything in the wiki" when the model receives only that subset.

Draft explanations can be useful immediately, but their unresolved status needs to travel into downloads and screenshots. A separate visible approval record is more reliable than relying on a footer that readers may miss. The goal is to make it easy to see what is known, what is proposed and who can authorize the next step.

12. Corrective work, in a useful order

These are recommended work items. They are not changes already made by this review and do not prescribe field response deadlines.

Priority / owner role	Work item	Evidence needed to close it
1 - Responsible CP specialist + data owner	Resolve Off/Depol chronology and reference basis. Validate the R and U derivations and source Pass/Fail rule.	Source procedure, survey records, formula rationale, applicable criteria and signed review decision.
1 - Product owner + specialist	Define what each output means. Keep criterion conclusion, evidence status, condition index and action priority distinct.	Agreed labels, examples and escalation rules; no implied approval from No rule flag.
1 - Specialist + field/data team	Resolve isolation conflicts and establish asset / CP circuit identities.	Canonical crosswalk, dated observations, evidence attachments and unresolved-case owners.
2 - Engineering + software team	Replace availability-only eligibility with method/validity/dependency checks.	Approved decision contract and reviewed examples for every unavailable-measurement route.
2 - Software team	Correct input normalization and response-contract inconsistencies.	Regression evidence for both original workbook and normalized-file paths.
2 - Project / data owner	Confirm public-versus-team-restricted content boundaries.	An access test of article, PDF, ZIP and underlying source routes; documented sharing decision.
3 - Specialist + evaluation team	Validate triage usefulness on independently reviewed cases.	Held-out evaluation, missed-concern review, uncertainty behavior and version comparison.
3 - Reporting team	Release a revised branch/parent review.	Reconciled counts, stable IDs, defined cohorts and statistic labels, visible unresolved findings.

Specific software defects and gaps to include

The audit reproduced these issues in isolated cases; they were not all present in the current source file. [L6]

- **Normalization mismatch:** local isolation scoring trims whitespace; the sanitized Jev field does not. "Good " can be scored as Good locally and sent as Unknown.
- **Conflict gating:** the normalized-input path can carry conflicting inventory status without the inventory-row marker that activates the conflict signal.
- **Misleading flag text:** the "Grade below C" message is triggered by score <60, which includes C grades.
- **Boundary inconsistency:** a 99.96 mV shift can display as 100.0 and activate a sanitized signal while the raw deterministic threshold remains unmet.
- **No independent review gate:** source Fail, shorted isolation or high model urgency is not individually sufficient to set Review needed. Decide the policy first, then implement it consistently.

These are smaller than the measurement-validity issue, but they affect trust in the interface. Fix them without silently changing historical records. Publish a version comparison explaining which rows changed, which rule caused each change and whether the change came from corrected evidence, software behavior or a new approved policy.

The first release target should be a review workflow the specialist can verify quickly. A more elaborate composite score is useful only if it improves a defined decision without hiding uncertainty or critical cases.

13. Questions and acceptance evidence for the review meeting

Decisions to obtain from the responsible specialist

1. What assets and surfaces are actually being assessed, and which criterion and edition apply to each group?
2. What do the 2023 Off and 2024 Depol dates represent? Is there a valid measurement sequence behind the pairing?
3. What reference electrodes, interruption methods and relevant operating conditions were used?
4. What physical model justifies the required-current formula and the derived resistance? Are system totals being assigned to multiple records intentionally?
5. What is the engineering basis for the source formula's Depol <= -400 mV condition?
6. Which isolation view controls each conflict, and what evidence establishes the current state?
7. Which unavailable-measurement routes may be used, on which assets, and who approves each route?
8. Which findings must remain visible at branch and parent level regardless of a favorable average?
9. Does the screenshot's Avg mean midrange throughout, and how were unpaired and zero-filled entries handled?
10. What would constitute a useful and acceptable ranking result, and what missed concern would be unacceptable?

Proposed validation plan

Build independently reviewed cases. Select representative systems, survey dates, missing-data conditions, isolation conflicts and interference questions. Have the responsible reviewers state the criterion conclusion, evidence sufficiency and action priority separately, with reasons. Preserve disagreements for adjudication. Agreement with the existing spreadsheet formula alone is not independent validation.

Separate development from evaluation. Hold out whole systems or later survey events where practical, so repeated station/system data do not appear on both sides. Compare the current workflow, deterministic review rules and Jev-assisted workflow using the same cases. Treat numeric thresholds for success as team decisions before reviewing results.

Measure the intended work. Track missed important concerns, unjustified clearances, unnecessary escalations, correct deferrals, explanation traceability, ranking stability and reviewer time. Evaluate confidence against observed correctness before describing it as calibrated. Record performance by missingness and asset group, not only one overall average.

Test structural behavior. Include zero versus blank; impossible/invalid units; a missing Off; an approved alternate route; cross-date Depol; a shorted isolation; contradictory source views; a new peer group; extreme potentials; an unpaired survey; duplicate placeholders; and a rejected model answer. Adding irrelevant metadata should not change a decision. Removing evidence must not silently become a stronger protection conclusion.

Release with a record of responsibility. Keep the source, derived inputs, approved rules, model response, reviewer action and final report linked. A human decision should be attributable and reversible without erasing the earlier state. Stop automated acceptance when essential prerequisites are unresolved; retain a clear route to continue the review.

AMPP describes CP4 work as spanning theoretical and practical system design, field evaluation and interpretation. Reading references can improve this project's reasoning; responsible specialist judgment and site experience remain necessary for approving its engineering use. [S10]

14. Primary references: measurement and criteria

Sources checked online on 18 October 2026. These are selected references, not an exhaustive literature review. A source's presence here does not establish its contractual or regulatory applicability to this facility. No lengthy standards text is reproduced.

S1. National Physical Laboratory - Cathodic Protection, Guide to Good Practice in Corrosion Control No. 1. Sections 3-4 and 6-7; printed pages 1-12. Introductory explanation of the electrochemical process, design context, measurement and limitations. Repository lists 2019; PDF describes an updated earlier publication and carries a 2020 copyright. Educational guidance, not the project's selected acceptance standard.

NPL guide PDF | Repository record

S2. 49 CFR Part 192, Appendix D - Criteria for Cathodic Protection and Determination of Measurements. Sections I.A, II, III and IV. Specific ferrous-material criteria, voltage-drop interpretation, polarization measurement and reference-electrode basis. eCFR displayed Title 49 current through 15 October 2026, with last amendment 13 October. eCFR is authoritative but unofficial; applicability must be determined separately.

Current Appendix D

S3. 49 CFR 192.1 and 192.463. Scope and external-corrosion-control CP provisions. Relevant to determining whether a gas-pipeline regulatory reference belongs in an asset's decision basis; not a finding that this entire plant falls under Part 192. Section 192.463 also addresses avoiding damage to pipe or coating from CP.

192.1 Scope | 192.463 Cathodic protection

S4. UFM 3-570-06 - Operation and Maintenance: Cathodic Protection Systems, 13 March 2026. Foreword; sections 2-1.1.2-.5, 2-1.2.4.1-.4 and 2-1.2.8.1-.4; Table 2-1 and Table 3-2. Printed pages 5-15. Covers potential-measurement error, criteria limitations and interpretation. The foreword explicitly distinguishes this informational manual from mandatory UFC requirements. No numeric overprotection limit is adopted from this document in this review.

2026 DoD O&M manual

S5. PHMSA Stakeholder Communications - Close Interval Survey factsheet. Survey types, purpose and limitations; educational material with a 2011 revision notation. It supports context-sensitive discussion of survey options, not a new acceptance rule or proof that an alternative is valid at this site.

PHMSA CIS factsheet

S6. Historical UFC 3-570-01 - Cathodic Protection, 28 November 2016, Change 1 dated 14 January 2019. Section B-8.8, printed pages 101-102, describes coupon test arrangements. Used only as historical technical background. The current WBDG catalog lists July 2026 digital UFC/UFS design entries; their full text was not retrieved. Do not specify work using this historical edition without checking the applicable current requirements.

Historical design PDF | Current WBDG catalog

15. Primary references: interference, evaluation and limits

S7. 49 CFR 192.473 - External corrosion control: interference currents. Paragraphs (a)-(c), particularly the specified onshore gas-transmission requirements in (c). Supports separating interference assessment from a generic voltage score. Its applicability and triggers are not transferred to the plant by this review.

Current interference provision

S8. University of Akron / PHMSA-sponsored CAAP final report - Probabilistic Performance Evaluation of Cathodically Protected Pipeline Considering AC Corrosion. Report cover dated 30 May 2024; file identifies the 2024 final report. Chapter 4 and section 4.3.1, printed pages 23-24, report experimental CP/AC interactions. Primary research under defined laboratory conditions; not a consensus acceptance standard or a universal field limit.

PHMSA-hosted research report

S9. NIST - AI Risk Management Framework 1.0, 2023, and AI Risks and Trustworthiness resource. Section 3, especially valid/reliable, explainable and accountable use. Voluntary AI evaluation guidance, not CP engineering criteria. The current resource notes that revision work is in progress. This review's case selection and proposed acceptance tests are project recommendations, not NIST-prescribed numeric requirements.

NIST trustworthiness guidance | AI RMF 1.0 PDF

S10. AMPP - Cathodic Protection Specialist program. Official description of the scope of CP4 education and certification. Used to distinguish specialist practice from software scoring or a literature review, not to verify any individual's credential.

AMPP CP Specialist program

S11. AMPP official standard catalog metadata. NACE SP0169-2024 concerns external corrosion control of underground/submerged metallic piping. NACE TM0497-2022 concerns measurement techniques related to CP criteria. Official indexed metadata and scope were available; full current standards text was not. No numbered requirement from these editions is represented as having been reviewed. Obtain the applicable licensed text and company interpretation before claiming conformance.

SP0169-2024 catalog | TM0497-2022 catalog

What this review could not establish

The complete signed CP4 report, site drawings, instrument records, survey procedures, material schedule, detailed electrical topology, coating condition and corrosion-loss measurements were not available in the reviewed evidence. Their absence from this review is not an assertion that they do not exist. Reference research cannot supply those missing site facts.

The engineering critique combines these primary references with direct inspection of local code, formulas and outputs. Recommendations to change labels, introduce approval gates or revise scoring are the review's own design judgments. They remain proposals until the responsible team accepts and validates them.

16. Local evidence and reproducibility

Local source files remain in the project owner's archive; paths below identify project-relative evidence.

The original workbook and screenshot are identified by filename. `R005` means source worksheet row 5, not fifth-ranked station. The detailed audit JSON preserves the complete lists behind the counts in this report.

ID	File or evidence	Use in this review
L1	Test Station Location Breakdown 2.xlsx	Original formulas and cached values, especially Test Station Current Status, Test Station Breakdown and Facility Breakdown.
L2	ranking-tool/engine.py; run.py; knowledge.py; export.mjs	Current formulas, flags, normalization, model contract, cache and knowledge boundaries.
L3	processed-2026-10-18/ranked-results.json	Saved 155-record condition and Jev outputs.
L4	system-review-2026-10-18/system-summary.json; cp4-survey-comparison.json	Parent/branch aggregation and 2024/2026 descriptive statistics.
L5	technical-critique-2026-10-18/audit_data.py; data-audit.json; data-audit-notes.md	Independent workbook/formula checks, availability, identity, survey and sensitivity audit.
L6	technical-critique-2026-10-18/scoring-implementation-audit.py; scoring-implementation-audit.json	Reproduction of unchanged scoring and isolated synthetic cases.
L7	Screenshot 2026-10-18 163427.png	Limited excerpt of the CP4 example supplied by the user.
L8	processed-2026-10-18/processing-audit.json	Provenance of the latest export and preservation of prior per-record answers.
L9	cortex-knowledge-base/; especially data/measurement-availability-policies.json	Draft rule boundary, knowledge structure and executable-rule separation.
L10	Live wiki/report routes and publication verification from this project session	Public accessibility of the existing shared wiki and prior PDF.

Hashes and versions

Original workbook SHA-256:

`2fbb4fb0615bff2d463ab0fee0ede3f4e1edd4d243614ad5fec779294a894b5d`

Saved ranked-results SHA-256:

`20079609946e00f806c9a137c8f3c903e8b28f55bd956797cb8391c9eaec85db`

Engine: `cp-triage-1.1`; model label: `jev-1.13.0`. Active knowledge fingerprint:

`6d035bb2d71ac77efb918f192714ebc0c33227629c40d8bce38885dc9cea992d`

Both audit scripts are retained with the detailed audit evidence in the project owner's archive. They read the source and saved outputs, write audit artifacts in their own directory, and do not call Jev or modify production results. Reproduction requires the owner-held project files and matching source hashes. Review the scripts before rerunning against a different source snapshot.

Existing team links: Wiki | Earlier system/survey review | Earlier PDF

This critique is published for team discussion and is DRAFT - awaiting CP4 review. It is not approved for engineering acceptance, is not active scoring policy, and has not been promoted into the active scoring bundle.