

OPERATING CASE STUDY / DATA, ANALYTICS & AI

Self-Service Analytics & Decision Quality

Turn a working data tool into a dependable path from question to decision

An anonymized operating case on adopting AI-assisted self-service analytics. The register describes a working installation, adversarial data checks and documented defects, while evidence of submitted reports, accepted fixes and wider adoption remains incomplete.

PROJECT CONTEXT / ANALYTICS ADOPTION & DEFECT REMEDIATION

BLUF / BOTTOM LINE UP FRONT

Scale self-service on verified answers.
Close defects at the point of consumption.

ANALYTICAL FINDING

An exact-match check supports only about one-third of the documented vocabulary entries. Counting documented categories therefore overstates usable literal coverage threefold; population-weighted answer coverage remains unknown.

CONTRIBUTION IN THE SOURCE RECORD

Performed local adoption and adversarial validation; prepared defect records. Provisional registry hoop support; product and repair ownership remained external.

EVIDENCE STATUS

Derivative-record validation and design case; originals unrecovered. A working installation does not establish rollout, correct-answer rates or savings.

BLUF: analytical inference. Source contribution: historical record. Client and internal identities anonymized.

EXECUTIVE ABSTRACT / DECISION AND OWNERSHIP

Promote self-service after the answer and repair paths are verified

[AI] The proposed Validated Self-Service Contract (VSC) separates usable installation, independently valid answers, permitted decision use and verified consumption of corrections. Each stage has its own evidence and owner.

Abstract element	Decision-relevant synthesis
Context [HR]	A newly available analytical product exposed curated data definitions to an AI caller. The register describes successful local setup followed by schema, vocabulary and correction-path defects.
Objective [PD]	Let analysts obtain correct, appropriately scoped answers while preserving an accountable fallback when the self-service path cannot support a decision.
Approach [HR/PD]	Local installation checks and adversarial queries generated defect records. This case reconciles those observations and proposes an all-request evaluation with independent semantic adjudication.
Findings [RC/AI]	An unweighted exact-match check supports about one-third of documented vocabulary entries. Installation and unit-test completion do not establish valid-answer coverage; documentation of a defect does not establish submission or repair.
Decision value [PD]	Compare the full burden of valid answers, specialist fallback and incorrect releases across every assigned request. Assess correction readback and information value separately from customer benefits.
Implications [PD]	The product owner controls promotion, data owners approve meaning and fixes, and an independent validator tests what later consumers actually receive. Usage alone is insufficient.

EXECUTIVE ACTION AND THRESHOLD

[PD] Fund a bounded validation pilot. Promote only after the lower confidence bound on matched-horizon resource benefit exceeds incremental cost and separate answer-validity, authority and availability requirements pass. Start by authenticating the source artifacts, reference answers and defect receipts.

Basis: S1 selected project register. [HR] historical record; [RC] recalculated; [AI] analytical inference; [PD] proposed design. No measured enterprise adoption or historical benefit is established.

EXHIBIT 1 / EFFECTIVE USE

A working installation is one gate in a longer answer pathway

[HR] The register describes a local installation, passing unit tests and a completed live query. The same validation work documents unusable fields and misleading literals. [AI] These findings test different properties and can coexist.

EFFECTIVE USE REQUIRES MORE THAN A SUCCESSFUL INVOCATION

Proposed VSC / every gate retains its own evidence and denominator



Passing one gate does not certify the next.
Defects and specialist fallback remain part of the original request population.

Source: S1 M07–M08. Proposed operating gates; no historical adoption or answer-accuracy rate.

Observed or proposed gate	What must be established
Accessible and executable [HR]	The local client can invoke a permitted query path. This is a point-in-time capability observation, not an adoption rate.
Semantically valid [PD]	The result answers the declared question with correct population, grain, values and maturity. Successful execution cannot supply this oracle.
Fit for the decision [HR/PD]	Diagnostic access remains distinct from certified performance reporting. The record describes an explicit restriction on reconstructing official scorecards.
Sustained and corrected [PD]	Measure repeat use and whether independently accepted fixes reach new consumers. Neither endpoint follows from a first successful request.

[AI] A named list of blocked users is not an enrollment denominator, and the record does not establish that the successful installer belonged to that list. An administrator-only usage endpoint leaves uptake unmeasured from the adopter’s perspective.

OPERATING IMPLICATION

[PD] Track all eligible requests through access, execution, validation and use. Retain failures and fallbacks in the denominator; assign a reason and owner to each unresolved gate.

Source: S1 M07–M08, implementation phases and role record. VSC gates are proposed. Historical tests and query evidence remain derivative descriptions until originals are authenticated.


EXHIBIT 2 / VOCABULARY ASSURANCE

A complete-looking list can overstate usable literal coverage

[HR] The register provides both the documented category strings and an enumerated observed vocabulary. [RC] Under exact string matching, the compatible documented-entry count is about one-third of the documented inventory.

EXACT LITERAL COMPATIBILITY IS SMALLER THAN DOCUMENTATION BREADTH

Historical vocabulary reconciliation / documented distinct-entry count = 100



Unweighted string compatibility does not identify record-weighted answer coverage.

Source: S1 M06. Exact comparison; a validated alias or normalization rule requires a separate test.

[RC/AI] Let D be the documented distinct literals and O the observed distinct literals. The relevant exact-match count is $|D \cap O|$, not $|D|$. With the documented inventory indexed to 100, compatible entries index to approximately 33.3. A count-based claim that assumes every documented entry is usable overstates that count threefold.

Coverage question	Evidence and decision boundary
Literal compatibility	Compare the actual strings under the declared comparison rule. A validated normalization or alias map is a different contract and requires its own test.
Population coverage	Weight matched categories by their eligible records on a fixed snapshot. The register does not provide the full compatible frequency distribution needed here.
Answer validity	Test complete questions, filters, joins and aggregations against independent truth. Category inventory coverage is neither routing precision nor a valid-answer rate.

NEW ANALYTICAL FINDING

[AI] Documentation breadth and executable vocabulary are different quantities. Before expanding use, correct and retest the literal contract; do not convert the unweighted inventory check into a population or customer-impact claim.

Source: S1 M06 and E03/E05/E07. Protected index: documented distinct-entry count = 100. Exact matching is the declared analytical assumption; raw strings, counts and version dates remain private.


EXHIBIT 3 / EMPTY ANSWERS AND GRAIN

An empty result can reflect an invalid predicate rather than an empty population

[HR] A documented scoping flag was recorded as unpopulated on one table, while similar logic was usable elsewhere. The register also identifies spelling, grain and deduplication hazards. [AI] A familiar query can therefore fail silently on a new data family.

AN EMPTY RESULT HAS MORE THAN ONE POSSIBLE EXPLANATION

Analytical distinction / the same visible output can require different action



Source: S1 M03–M06. Determine which branch applies using independent source and query evidence.

Failure mode	Required diagnostic [PD]
Unknown filter field	Test nullness and eligibility independently. SQL comparison with NULL is unknown; no matching rows do not prove that no eligible events exist.
Literal or prefix mismatch	Enumerate supported values and inspect valid exceptions. Do not assume a prefix captures every conceptually related record.
Wrong counting unit	Compare records with validated entity keys and retained event lineage. A distinct identifier is not automatically a unique physical entity.
Unstable row selection	ROW_NUMBER assigns one rank per row; ties make the chosen row unstable, not multiple rank-one rows within one partition. Inspect the actual key, filter and downstream joins.

[AI] The source’s explanation of duplicate survivors cannot be accepted solely as a tie effect. Without the executed query and intermediate outputs, its root cause remains unresolved. A record-count correction also cannot be multiplied by a spelling-coverage factor from another population.

OPERATING IMPLICATION

[PD] Use discriminating fixtures containing nulls, valid exceptions, duplicate entities and conflicting revisions. Test the intended answer and its population before accepting a plausible aggregate.

Source: S1 M02–M06 and D03; SQL semantics independently checked against official documentation. The case does not reconstruct the unavailable original query or infer a current backend state.

EXHIBIT 4 / VALUE AND THE SERVICE CLOCK

Query runtime does not measure time to a trusted business answer

[HR] The launch business case cites a historical specialist-request turnaround, while local verification records a short query execution. [AI] The populations, clocks and completed work differ, so dividing those figures would manufacture a speedup.

ALIGN THE START, STOP AND COMPLETED WORK BEFORE COMPARING SPEED

Proposed measurement / diagram is schematic and has no common numeric time axis



Request-to-validated-answer clock: the whole pathway

Execution clock: one component

Source: S1 M01/M07. A historic request queue and a single query do not form a paired speedup.

Clock or value claim	Comparable measurement [PD]
Request-to-answer cycle	Start at accepted business demand; stop at an independently validated answer delivered to its owner. Include access waiting, clarification, reruns and specialist work.
Execution latency	Start and stop at the specified query events. Retain failed and abandoned jobs separately; a fast returned aggregate may still be wrong.
Usable capacity released	Measure incremental analyst and specialist effort on the same task mix. Avoided waiting is not automatically paid labor or cash savings.
Customer or operating benefit	Link an accepted answer to a delivered action and a mature outcome comparison. A diagnostic tool’s usage cannot establish that effect.

[AI] An unreconciled broad financial pull in the register is evidence of a reporting hold. It is not a concession total suitable for publication, a benefit pool or a basis for annualized savings. Its raw amount remains private.

OPERATING IMPLICATION

[PD] Freeze the clock, eligibility and answer-quality rule before comparing service paths. Price all-request resource burden and evaluate mature customer outcomes separately.

Source: S1 M01, M07 and M09. No historical speedup, labor saving or financial benefit is estimated. The prospective model uses explicitly hypothetical resources and matched request horizons.


EXHIBIT 5 / ADOPTION DEPENDENCIES

Removing one setup step can leave the same access dependency in place

[HR] The register describes a proposed distribution change that would remove repository cloning. It also describes a repository update at each application start. [AI] If both depend on the same unavailable credential, removing cloning leaves the startup failure path intact.

A DISTRIBUTION CHANGE CAN LEAVE A REPEATED FAILURE PATH

Conditional dependency analysis / installation and startup are different checks



Source: S1 M07/D01. Conditional on the recorded startup design; not a current access test.

Lifecycle check	Decision evidence [PD]
Initial installation	Reproduce the failure and correction on the supported platform, then verify local registration and a permitted request.
Restart and renewal	Test a fresh process after credential expiry, renewal and client upgrade. A successful first session cannot demonstrate these transitions.
Least-privilege operation	Distinguish a denied administrative operation from a broken user capability. Resolve required access through the authorized owner and scope.
Broad distribution	Require evidence across platform and dependency groups, plus rollback and owner support. Packaging changes alone do not establish reachable service.

[HR/AI] The record credits local troubleshooting to the adopter and product construction to another team. It does not establish that a security-review task was completed merely because that task appears in the project register.

OPERATING IMPLICATION

[PD] Test the full supported lifecycle before widening access. A proposed repair must remove the active failure path and preserve authorization, not merely change the first visible step.

Source: S1 M07, D01, implementation phases and E14 qualification. Dependency diagram is a conditional design interpretation; no current access or credential behavior was tested for this case.


EXHIBIT 6 / REMEDIATION AND LEARNING

A correction is not closed until a later consumer receives it

[HR] The source describes a path that can save shared knowledge while failing to read it, then using a bundled reference. It also identifies owners and a reporting route without preserving evidence of submission or receipt. [AI] These are separate gaps in the learning process.

LEARNING CLOSES AT THE NEXT CONSUMER, NOT AT THE WRITE ACKNOWLEDGMENT

Proposed remediation chain / every transition needs a distinct receipt



Source: S1 M08/M12/D04. Historical routing and closure are unconfirmed; this chain is proposed.

State	Evidence required to claim the next state [PD]
Defect documented	Reproducible question, expected answer, failing output, source version and affected decision. Discovery count is not a defect-prevalence estimate.
Owner accepted	Tracked submission plus receipt, scope, severity and an accountable repair owner. A routing plan does not prove any of these events occurred.
Correction approved	Accepted changed artifact, independent regression checks and the authority required to release it. A write acknowledgment does not prove readback.
Consumer verified	Fresh-client retrieval of the corrected version and successful answer replay. Check fallbacks and caches before marking the issue closed.

[HR/AI] A later retrieval repeats challenged documentation. That shows persistence in the retrieved text; it does not newly establish that the earlier backend failure still occurs, that every fix failed, or that the report reached its owner.

OPERATING IMPLICATION

[PD] Give every material defect a receipt and a consumer verification test. Track stage-specific age and recurrence, retaining unresolved records rather than treating planned escalation as closure.

Source: S1 M08, M12, D04 and implementation phase 3. Reporting, owner acceptance and verified closure remain unestablished. The proposed workflow does not authorize messages, grants or system changes.


EXHIBIT 7 / PROMOTION DECISION

Evaluate useful answers and harmful releases alongside adoption

[PD] A promotion decision needs an all-request view: independently valid self-service, accountable specialist fallback, incorrect release and unresolved demand. Installation counts or successful-query latency omit the work and risk that remain.

PROMOTION REQUIRES ECONOMIC VALUE AND INDEPENDENT ANSWER GUARDRAILS

Proposed decision logic / a gain in usage cannot offset a failed validity condition



Source: Luna Sol proposed VSC. The model is hypothetical; historical rollout and outcomes are unmeasured.

Decision measure	Accountable role and acceptance evidence [PD]
Validated-answer coverage	Independent evaluator; all originally eligible requests, fixed deadline, blinded answer adjudication and family-specific uncertainty.
Incorrect-release burden	Product risk owner; retain errors even when a later audit corrects them. Apply a separately specified acceptance margin.
Correction visibility	Data and product owners; accepted repair, versioned readback and a replay by a fresh consumer.
Resource and availability value	Business owner; matched-horizon effort, fallback and downtime, incremental cost, non-overlap and feasible recruitment.

[PD] The benchmark tests the full adoption-and-repair offer on independent dependency domains. A separate synthetic counterfactual increases installation while retaining the old correction path, revealing why higher use can coexist with greater wrong-answer exposure.

OPERATING IMPLICATION

[PD] Start with a blinded calibration pilot. Scale only when economic and validity gates both pass, the supported platform lifecycle works, and owners can demonstrate that corrections reach consumers.

Source: S1 adoption and remediation evidence gaps. Scorecard, benchmark and promotion rules are proposed. Historical answer accuracy, rollout uptake, mature impact and realized savings remain unmeasured.


TECHNICAL ANALYSIS

A correction creates value only after the consumer can use it

[AI/PD] The proposed Validated Self-Service Contract (VSC) separates usable access, correct output, permitted decision use and verified correction uptake. The graph identifies the links each claim must support.

VALIDATED SELF-SERVICE CONNECTS ACCESS TO CORRECTION CONSUMPTION

Proposed DAG / benchmark loss and mature customer outcomes are different endpoints



Paths beyond validated answers remain hypotheses. Readback verifies visibility; it does not itself certify truth.

[PD] Z assigns the verified-adoption offer. I is usable installation; W is an accepted correction write; R is the consumer reading that corrected version. Z affects I, W and R, including readback conditional on a write. I and R influence valid analytical output V, then decision D, delivered action A and mature customer outcome Y. Z also affects V through validation. Readback verifies visibility, not truth.

[PD] Independent reference J and request outcomes determine benchmark resource burden L. Q captures source quality, grain and maturity; it affects V and recorded customer measure O, while Y also affects O. C names platform/access, analyst experience, request difficulty and owner capacity; U covers unrecorded source changes and concurrent programs. These affect adoption, correction and downstream outcomes.

[AI] The registered installation and diagnostic work concern I, W, R and V; they do not identify the customer path. A write receipt cannot establish $W \rightarrow R$, and a query completing cannot establish $I \rightarrow V$. The proposed randomized benchmark would identify Z's effect on L within its defined setting. Residual differences may reflect unmeasured source, user or operating paths.

Basis: registered adoption and validation account [HR]; VSC causal graph and proposed benchmark [PD]. Arrows are mechanisms to test, not historical effects.

TECHNICAL ANALYSIS

Measure qualified use across the whole request journey

[HR/AI] A connected client, passing tests and a completed query are distinct achievements. The supplied record does not provide a complete adoption cohort or an end-to-end request-turnaround comparison.

Stage [PD]	Observable measure and denominator [PD]
Reach and access	Eligible users offered the tool; successful authorized installation under their intended platform and identity. Preserve blocked and unattempted users.
Repeated qualified use	Original eligible requests answered correctly for the permitted purpose, with user and request identifiers; exclude duplicate calls and test traffic from activity counts.
Service time	Elapsed time from original request intake to independently accepted output or justified specialist disposition; unresolved requests remain at the deadline.
Trust and continuity	Wrong releases, unsupported reporting use, specialist fallback and repeated valid outputs across later sessions; successful execution alone is insufficient.

[AI] The product’s historical request-resolution business case includes queueing, clarification, specialist work and acceptance. The reported execution time covers a single query after installation. Dividing those clocks would manufacture a speedup. Compare the same task population, intake point, completion rule and maturity before claiming a cycle-time effect.

[HR/AI] The register reports local installation and validation, but ordinary users could not inspect the administrative usage report. That visibility limit makes wider adoption unknown. It does not imply no adoption. Nor does a count of blocked users reveal a failure rate without the eligible pilot population and observation window.

[PD] Freeze eligibility before outreach. Track invitation, entitlement, attempted installation, usable access, first qualified answer and later qualified requests as separate transitions. Report transition-specific rates and the all-eligible result. Conditioning evaluation on connected users would select on a mechanism changed by the intervention.

[AI] Effective-use theory motivates this separation; it supplies no project effect estimate. An analytics service can gain usage while its answers remain unsuitable for an executive report. More activity is valuable only when the intended task and its quality boundary are satisfied.

Basis: registered install, test, query and visibility claims [HR]; measurement design [PD]; Burton-Jones and Grange; DeLone and McLean. No historical adoption or speedup estimate.

TECHNICAL ANALYSIS

Validate the answer against cases that can expose a silent error

[HR/AI] The tool supplied grounding information to a caller that authored queries. Incorrect documentation could therefore generate a valid query with an invalid answer, rather than a visible software error.

Test layer [PD]	What passing establishes—and does not [AI]
Installation and unit checks	The tested setup and specified assertions work; untested data semantics, current permissions and future sessions remain separate.
Live schema and vocabulary	The observed source contains the examined columns and literals in that scope and vintage; absence in a filtered sample does not prove global absence.
Discriminating query fixtures	The query handles specified nulls, aliases, duplicated items, boundaries and nonmatching rows; the reference meaning itself still needs validation.
Accepted analytical answer	Population, grain, aggregation, maturity and intended use match an independently established task; a clean result table alone cannot establish this.

[RC/AI] Exact-literal intersection is only about one-third of the advertised vocabulary inventory in the examined example, with that inventory indexed to 100. Comparing list sizes would overstate usable exact matches threefold. This is a label-set calculation, not request-weighted answer coverage or a population error rate. Similar spellings may encode different status paths.

[AI] A filter applied to an all-null field can yield no rows without establishing an empty real-world population. Ordinary SQL comparisons with null are unknown. The appropriate check is a source-value profile and a task-specific scope rule, not automatic replacement of unknown by zero or an invented category.

[PD] Build held-out fixtures containing values that distinguish the intended query from plausible wrong queries: included and excluded rows, genuine zero outcomes, missing keys, multiple items per entity and boundary dates. Freeze the gold task before inspecting candidate output. Equivalent SQL strings may differ; agreement on one database can also be accidental.

[AI] Test-suite evaluation and the oracle literature explain these two failure modes. They do not prove any local query correct. Independently adjudicate the reference and retain unresolved task ambiguity rather than forcing a false binary pass.

Basis: registered grounding, vocabulary and null findings [HR/RC]; Zhong, Yu and Klein; Barr et al.; PostgreSQL comparison semantics. No benchmark accuracy is transferred.

TECHNICAL ANALYSIS

Counting rules must preserve the entity, cohort and evidence clock

[AI] Correcting a row count does not automatically identify physical entities or mature outcomes. The source’s counting conventions need stronger premises than a successful DISTINCT expression.

Construct [PD]	Validity requirement [PD]
Entity count	Specify item, tracking identifier or physical parcel. Audit one-to-many identifier changes, collisions, omitted records and linkage errors before claiming a physical-entity bound.
Status or vocabulary coverage	Separate exact literals, semantic groups and row-weighted coverage. Preserve path distinctions and missing categories; do not merge labels solely because they look similar.
Mature rate	Freeze eligible cohort, numerator event, positive exposure, attribution rule, as-of timestamp and maturity horizon. Calendar recognition flow is not automatically cohort risk.
Diagnostic versus reporting use	Record the allowed purpose and authorized reference feed. A valid exploratory query need not reproduce the approved reporting numerator.

[HR/AI] The register describes a correction from item rows to distinct identifiers that changed a draft ranking. It supports a counting-rule correction within the stated snapshot. It does not establish reduced customer defects, savings, or a complete physical-parcel count. Replanning can issue multiple identifiers for the same entity, so the stated identifier count is not automatically a floor.

[AI] ROW_NUMBER assigns a different row number within each partition even when ordering values tie. A subsequent rn=1 filter retains one row per partition before downstream multiplication. Ties can make the selected row arbitrary; they do not themselves create multiple first-ranked rows. The missing original query prevents deciding whether a later join, partition choice or different ranking function produced the reported duplication.

[PD] Replay the full query pipeline on a frozen snapshot and inspect row counts after each join. Require a deterministic tie rule only when its business meaning is defensible. MAX/MIN status summaries are alternative definitions or sensitivity bounds under explicit assumptions, not an automatic truth oracle.

[AI] Later write-off recognition can change an as-of count without any new intervention. Preserve versioned cohort extracts and report pending maturity separately. The tool’s diagnostic-only boundary should remain visible when a result is exported or reused in a report.

Basis: registered grain, status-path and reporting restrictions [HR]; entity-resolution and cohort logic [AI/PD]; PostgreSQL window-function semantics. No missing query is reconstructed.

TECHNICAL ANALYSIS

Close the repair loop at the next consumer, not at the write receipt

[HR/AI] The register describes shared-knowledge writes with a failed read path and a bundled-content fallback. This combination can preserve apparent availability while blocking the intended correction pathway.

Repair state [PD]	Required independent evidence [PD]
Documented and submitted	Reproducible defect, source version, intended result, affected scope and a transmitted intake record; a proposed route is not a submission.
Accepted and changed	Receiving owner accepts responsibility; approved change identifies the source or artifact version. Administrative closure alone is insufficient.
Delivered and consumed	The intended consumer identity reads the corrected content through the normal path; origin, version and fallback status are captured.
Verified and sustained	A held-out affected request passes independent semantics and later sessions retain the repair. Reopening preserves the original defect identity.

[AI] Let W denote a stored correction and R its consumption. A write success estimates neither $P(R|W)$ nor answer correctness. The chain requires both a valid correction and consumer exposure. If R is absent, a local workaround may solve an immediate request while the shared defect remains available to other users.

[HR/AI] The available record identifies intended product, code and upstream-data owners but does not authenticate a sent report, receipt or accepted fix. Later documentation retrieval reportedly retained problematic text. That establishes persistence of the retrieved text, not a fresh backend execution failure, an unfixed state in every environment or a closure rate over the initial defect list.

[PD] Test a harmless approved correction in an isolated environment: write an identifiable version, read it under the consumer’s actual permissions, execute an affected fixture, then verify the next session. Capture rejected writes, denied reads and fallback origin. Permission changes still require the legitimate owner; the validation protocol grants no additional authority.

[AI] The organizational-learning distinction is precise here: diagnosing and routing a recurring problem is different from demonstrating system change. The proposed technical addition is to make consumer-visible repair an auditable outcome. Repeated documents about the same defect are not independent evidence that this loop closed.

Basis: registered shared-knowledge, routing and later-retrieval accounts [HR]; proposed write/readback protocol [PD]; Tucker and Edmondson; W3C PROV-DM.

TECHNICAL ANALYSIS

Audit the measurement process before ranking users or defects

[PD] Measure access, semantic correctness, permitted use and repair uptake on separate frozen frames. Neither a selected debugging session nor the longest defect list estimates ordinary deployment quality.

[PD] For future task adjudication, sample a random core of 200 request-context cases from the declared deployment frame. Independently code source validity, answer correctness, permissible hold and correction uptake with two blinded reviewers. Add rare null, alias, grain and readback failures as challenge cases; apply design weights before population estimates. Shared source families require clustered uncertainty.

[PD] Under independent cases, that core gives a worst-case raw-agreement normal-screen margin of about ± 6.9 percentage points at 95%. This is not kappa precision. Use pilot class prevalence and reviewer confusion to simulate kappa intervals and adjust sample size. Preserve both original labels and independently adjudicate disagreements; agreement cannot certify a mistaken reference.

Variation requires comparable observations

[RC/AI] The private audit retains ten two-cut screens: four as-of quotation pairs and six mechanical comparisons. With MR the absolute adjacent change, individual limits are $\text{mean}(x) \pm 2.66 \times \text{mean}(\text{MR})$, and the MR upper limit is $3.267 \times \text{mean}(\text{MR})$. All points and ranges fall within these arithmetic bands; the unchanged base pair is degenerate. None establishes operating stability or change.

[HR/AI] There is no commensurate operating panel, repeated eligible-user adoption cohort or fixed defect-closure series. As-of revisions concern measurement maturity, not treatment response. Mechanical comparisons mix artifacts or definitions. A Laney p-prime chart cannot be reconstructed without compatible repeated numerator/exposure counts. Within-limit movement is not established improvement or deterioration.

Selection and moderators

[HR/AI] Validation concentrated on a difficult analytical task and touched tables, not a randomly sampled estate. There is no untreated worst-user or worst-table follow-up panel. As a hypothetical normal-error sensitivity, selecting the worst tenth gives a mean 1.755 SD; assumed repeat reliability of 0.3, 0.6 or 0.9 implies untreated reversion of about 1.23, 0.70 or 0.18 SD. These are assumptions, not measured project effects.

[PD] Freeze focus/non-focus membership and compare pre-period task mix. Test platform, business-line, node-type, source-family and analyst-experience moderators. High missingness concentrated in one node type defeats automatic transport; broadening installation can expose more users to an unresolved defect.

Basis: proposed blinded audit and selection sensitivity [PD]; source coverage limits [HR]; Cohen; NIST SPC; Laney; Barnett et al. Private calculations disclose every eligible screen.

TECHNICAL ANALYSIS

Credit adoption and validation without claiming ownership of the product

[HR/AI] The register attributes installation work, diagnostic rules and errata to the project lead. Product construction, access authority, shared documentation and remediation belong to separate owners.

Contribution claim	Process-tracing assessment [HR/AI]
Directed local adoption and validation	Provisional hoop: derivative records identify direction and local work; original commands, logs and authorship still need authentication.
Diagnosed the installation blocker	Provisional hoop for the documented diagnosis and workaround. A contemporaneous authenticated execution trace could strengthen the bounded claim; broader platform recovery is not established.
Caught a counting error and revised a draft	Provisional hoop for analytical correction. Agent-executed queries and later narrative summaries do not establish sole human execution or downstream operating impact.
Reported and repaired shared defects	Documented findings and intended owners are supported at registry level. Submission, acceptance, deployment and consumer uptake remain unestablished.
Built the product or delivered financial benefit	Not supported. The originating product team owns the build and business case; no causal benefit or realized saving is supplied.

[AI] Repeated summaries are not independent corroboration. A first-person memory statement may describe assisted session work rather than a uniquely attributable human action. No smoking-gun or doubly decisive test is established for the full adoption-to-remediation chain.

Consolidated limits and reproduction

[HR/AI] Evidence is derivative and one-sided; original sources require recovery. Measurement has unresolved entity linkage, taxonomy, missingness, source vintage and reporting eligibility. Identification lacks a comparable untreated cohort and observed repair exposure. Transport is limited by selected tasks, environments and owner dependencies. Missing counterparty evidence does not establish that no work occurred elsewhere; current status needs a fresh authorized check.

[AI/PD] Economics remains hypothetical; corrected source totals are not savings. Reproduction requires source/version manifests, extraction rules, query-stage counts, coding decisions, task truth, consumer-origin traces and executable sensitivity calculations. Public anonymity keeps raw values, identifiers, dates and inversion anchors in the separate private annex. The benchmark cannot authenticate missing historical records.

Basis: registered roles, governance and evidence limitations [HR]; process-tracing grades [AI]. Case type: adoption/validation design and evidence reconciliation, not proven operating impact.

TECHNICAL ANALYSIS

Evaluate adoption together with correction uptake at fixed deadlines

[PD] Randomize independent analyst/source/owner domains to a verified-adoption package or usual rollout in a controlled benchmark. No real access grant, live client decision or customer intervention is executed.

[PD] Freeze 20 analytical requests per domain: 12 clean requests and eight affected by a shared defect. A hidden, independently checked source manifest defines the intended answer. Usable installation, accepted correction write and consumer readback are common domain states, so requests within a domain are correlated. Shared owners, knowledge stores and specialist capacity must be isolated in the benchmark or clustered together.

[PD] The response deadline is day 14 from original intake; blinded error ascertainment finishes by day 28 against the same task and source snapshot. Retain all assigned requests, including missing or late work. A wrong released answer remains wrong even if later corrected; a fallback counts as valid only if independently established. Missing or unresolved work receives the prespecified terminal penalty.

[PD] The primary ITT estimand is $\Delta = E[L(\text{usual})] - E[L(\text{VSC})]$, mean total resource burden per assigned domain over all original requests and the full horizon. Positive Δ is benefit. Installation and readback are mediators, not exclusions or randomized compliance. No effect of actual use is identified by comparing users with nonusers after assignment.

[PD] The planning law assumes reserved specialist capacity and no additional nonresponse; deviations require recalibration. The endpoint mixture is a synthetic effectiveness path, not an estimated dose response. Exact stochastic ordering must be verified on complete domain distributions before using a boundary-specific confidence inversion. The model pages compare detection and investment power with the same-horizon economic threshold.

[PD] Separately preregister wrong-answer, permitted-use and useful-coverage guardrails. Easier installation alone can lower aggregate burden while exposing more requests to uncorrected content. A favorable resource score cannot waive quality. Setup, owner coordination and production readback costs are distinct from request handling and research-only adjudication.

[PD] A later operational rollout could stagger adoption across independent clusters and estimate group-time ATT against not-yet-treated units. Require no anticipation, common support, pre-period and event-time diagnostics, case-mix adjustment and interference controls. Nonsignificant pretrends are not proof of comparability. Mature customer outcomes, variance and eligible volume must be collected for a separately powered field study; benchmark success is not customer impact.

Basis: proposed controlled VSC benchmark [PD]; model assumptions are synthetic. Later field design follows Callaway–Sant’Anna and remains unpowered; recruitment feasibility is not established.

TECHNICAL ANALYSIS

Four propositions make trustworthy self-service falsifiable

[AI/PD] VSC treats useful adoption and verified correction uptake as separate requirements. Its value is a future empirical question; the existence of a tool, an erratum or a proposed protocol is not evidence of its effect.

Proposition and mechanism [PD]	Disconfirming result or boundary [PD]
P1: discriminating source fixtures reduce confidently wrong answers at fixed request coverage.	A sufficiently precise failure to meet the preregistered quality improvement, or excessive valid-answer suppression, contradicts the proposed advantage.
P2: consumer readback of an approved correction reduces recurrence on affected held-out tasks.	Verified readback leaves the same semantic failures, or the recorded version is not the version used by the answering process.
P3: qualified adoption predicts useful output better than connection or test counts alone.	A frozen held-out cohort shows no material incremental prediction, or the relationship fails across supported platforms and request families.
P4: accepted-owner routing plus verified uptake reduces repeated exposure without excessive specialist burden.	Defects recur despite accepted delivery, or quality-adjusted resource burden cannot clear the economic threshold without violating guardrails.

[AI] Burton-Jones and Grange distinguish use from effective use; DeLone and McLean motivate separate quality and use measures. The case’s contribution is an observable consumer-read boundary within that broader distinction. It proposes a mechanism, not confirmation of either complete theory.

[AI] Text-to-SQL evaluation shows why a plausible result and a string match are imperfect correctness tests. The oracle literature adds that the test reference can itself be wrong. VSC therefore couples discriminating tasks with independently adjudicated source meaning; a larger passing-test count cannot compensate for missing semantic coverage.

[AI] Tucker and Edmondson distinguish local problem solving from organizational learning. This case narrows the distinction to a traceable chain from reported defect to approved correction to consumer uptake. It does not infer staff motives or an organization-wide learning culture from missing receipts.

[AI/PD] Deming requires comparable measurement before interpreting variation. Campbell and Goodhart motivate testing whether usage targets reward trusted output. Ittner cautions against equating measurement sophistication with performance. Composite design exposes compensability: usage cannot offset denied authority or an invalid answer.

[PD] Preregister material effects and guardrails. A lower confidence bound above break-even supports the modeled investment; an upper bound below contradicts it. A spanning interval is inconclusive. Statistical success cannot waive unresolved source or permission requirements.

Basis: proposed VSC propositions [PD]; Burton-Jones and Grange; DeLone and McLean; Zhong et al.; Barr et al.; Tucker and Edmondson; Deming; Campbell; Ittner et al.; OECD.

PROSPECTIVE MODEL / VALIDATED SELF-SERVICE

Evaluate the complete request, not the successful query

[PD] Randomize independent team/provider domains to a Validated Self-Service Contract (VSC) or usual rollout support. This is a controlled adoption-and-repair benchmark; every parameter below is hypothetical.

[PD] Freeze 20 analytical requests per domain: 12 unaffected and eight exposed to one known shared-content defect. Keep the agent, source snapshots and approved access fixed. Assess the original 14-day response deadline, then complete independent blinded semantic review by day 28. No live clients, grants or production releases are involved.

Terminal request state	Loss units	Assessment rule
Valid self-service answer	2	Correct for the frozen question and allowed use
Independently valid specialist fallback	6	Includes handoff and analyst effort
Wrong answer or unresolved request	18	Retained at original deadline; wrong answers remain wrong after repair

[PD] Primary ITT is usual-minus-VSC cumulative burden over every original request, averaged by domain. Do not condition on installation, query success or observed use. Track wrong answers and unresolved requests separately even when their planning penalties coincide.

THE DENOMINATOR DOES NOT SHRINK

[PD] Withdrawals, missing runs and holds remain eligible. The planning law reserves valid fallback capacity and assumes no additional nonresponse; violations require recalibration. Research adjudication is separate from operating effort.

[PD] Request correctness uses the original source snapshot, not later revised data. Freeze classification, time horizon and resource weights before assignment; no historical cycle-time reduction is inferred.

PROSPECTIVE MODEL / VALIDATED SELF-SERVICE

A written repair has no benefit until the consumer receives it

[PD] Installation, writing a correction and consuming the corrected content are distinct domain-level gates. All affected requests share those gates.

Conditional gate / verdict	Usual	VSC
Usable installation	0.60	0.85
Repair written, given installed eligible defect	0.50	0.80
Corrected content consumed, given written	0.45	0.90
Unaffected: valid / fallback	0.90 / 0.10	0.90 / 0.10
Affected, no consumed repair: wrong / fallback	0.50 / 0.50	0.50 / 0.50
Affected, repair consumed: valid / wrong / fallback	0.85 / 0.05 / 0.10	0.85 / 0.05 / 0.10

[PD] No installation sends every request to valid specialist fallback. Conditional on the shared gates, verdicts are independent. Isolate provider versions, support staff and repair ownership between randomized domains. Freeze independently reviewed correction content and verify what consumers actually receive; a successful write receipt is insufficient.

[PD] For power, $P(\lambda)$ mixes the usual and full-VSC domain distributions. λ represents modeled domain response, not direct gate interpolation, observed compliance or the causal effect among users. Exact rational probabilities verify every endpoint CDF inequality; the mixture therefore orders cumulative loss.

ADOPTION ALONE CAN EXPAND EXPOSURE

[RC/PD] Raising installation alone lowers modeled mean burden only 2.43 points, but increases wrong answers 41.67%. This separate counterfactual is not a randomized historical finding or part of the monotone response path.

[PD/RC] Synthetic fixed request mix. Primary economic power does not size wrong-answer, unresolved-work or access-compliance guardrails; preregister and power those separately.

PROSPECTIVE MODEL / VALIDATED SELF-SERVICE

Require detectable break-even value before funding expansion

[PD] Both gates must pass: 80% power against zero at break-even, and 80% probability that a one-sided 97.5% lower bound exceeds break-even at the full planning alternative.

[RC/PD] Usual burden is indexed to 100. Full modeled reduction is 28.56 points. Net resource value = reduction × avoidable share a × non-overlap h – incremental cost 12. Base a = h = 1 gives break-even 12; this is not a cash estimate.

Domains per arm / total	80% MDE	Economic power	Both gates
12 / 24	28.444	32.5%	Fail
67 / 134	11.931	93.6%	Conditional
140 / 280	8.158	99.9%	Conditional

[RC/PD] Minimum: 67 per arm. Lower consumer readback, 0.75 instead of 0.90, requires 115. Candidate 140 per arm clears that stress. Independent domains and recruitment remain unestablished; pilot calibration comes first.

Net at full effect / cost	a = 0.50; h = 0.80	a = 0.75; h = 0.90	a = h = 1.00
9	2.42	10.28	19.56
12	-0.58	7.28	16.56
15	-3.58	4.28	13.56

[PD] Cost 12 allocates installation support 5, shared-owner coordination 4 and production readback 3 over the same service pool. Request-specific effort is already in loss; do not duplicate shared repair costs. Uniform shares are assumed; reprice categories if realization or overlap differs.

[PD] Separate future field value = eligible treated volume × mature customer ATT reduction × avoidable cost/outcome × non-overlap – incremental cost. Match population and maturity; unavailable inputs leave this field claim unestimated and unpowered.

[PD/RC] Known finite-law two-sided alpha ≤0.05 and economic false clearance ≤2.5%; nuisance uncertainty, oracle error, altered fallback capacity or interference require recalibration.

PROSPECTIVE MODEL / VALIDATED SELF-SERVICE

Purchase evidence that distinguishes expansion from repair

[PD] This separate decision model asks whether additional evidence should redirect investment toward controlled expansion, the correction read path, or semantic assurance.

[PD] States are ready under declared tests, read-path-only fault, and semantic defect, with priors 0.40 / 0.35 / 0.25. Semantic defect includes combined failures and takes precedence. Its repair payoff includes any required read-path work. Other failure classes require a larger state space.

Planning action / state	Ready	Read fault	Semantic defect
Plan bounded expansion	15	-24	-30
Repair the consumer read path	5	17	-12
Fund semantic assurance and repair	-4	-5	20

[RC/PD] Prior-optimal read-path repair has utility 4.95; perfect information adds at most 12.00. A three-label signal is correct with probability 0.50, 0.75 or 0.90; wrong labels split the remainder equally. Choose the best plan after each signal.

Signal quality	EVSI	Evidence cost	Net EVSI
Low	1.1750	0.50	0.6750
Base	5.2313	2.00	3.2313
High	9.2925	4.00	5.2925

[PD] Net utilities already include subsequent planning-action costs; evidence acquisition is additional. These utilities do not use the trial benefit and must not be added to it. Favorable information value never waives source access, allowed-use limits or independent answer validation.

THE OPERATING DECISION

[AI/PD] Treat adoption and verified repair consumption as separate obligations. A successful install establishes a usable path; a recorded correction establishes work in progress; the consumer-visible validated answer is the release test.

[PD/RC] Hypothetical Bayesian information model. No historical remediation completion, program adoption rate, customer effect or realized financial benefit is estimated.

Reproduce the assurance case without confusing it with measured impact

[AI] This is an adoption-validation and operating-design case. The record supports analytical discovery and bounded personal contribution; enterprise uptake, accepted remediation and customer benefits require different evidence.

Claim class	Meaning and boundary
[HR] Historical record	An observation, described action or attributed contribution in the supplied register. Original operational artifacts remain a separate authentication requirement.
[RC] Recalculated	Set intersections, compatible source arithmetic, mechanical screens and hypothetical model outputs recomputed from specified inputs. Calculation does not certify provenance.
[AI] Analytical inference	A conclusion conditional on explicit premises: incomplete literal compatibility, incomparable service clocks or a missing correction-consumption path.
[PD] Proposed design	VSC, prospective experiments, independent truth checks, promotion gates and economic/information-value scenarios. These are not executed trials or realized savings.

Independent reproduction sequence

[PD] Verify the workbook hash and selected-cell extract. Recreate exact-string intersections and retain a separate semantic-alias analysis. Reconcile defect lists, source dates, actual query units and observed versus claimed remediation stages. Do not infer current backend behavior from later documentation alone.

[PD] Rebuild the domain-level mixture of access, correction and readback states; include every assigned request, fallback and wrong release. Recompute the planning tests at their economic boundary, the installation-only counterfactual and the separate information-value decision. Revalidate every displayed result against the frozen private data.

WHAT THE NEW ANALYSIS CHANGES

[AI] Documentation volume does not identify usable vocabulary. A working tool does not establish a valid-answer service. A correction reaches an operational endpoint only when a later consumer receives and successfully uses the approved version.

Labels: [HR] historical record; [RC] recalculated; [AI] analytical inference; [PD] proposed design. Source basis S1; see private annex.

RESEARCH REFERENCES / METHODS AND INTERPRETATION

Distinguish effective use from execution and test completion

Primary research and official methods inform the interpretation and proposed tests. They do not independently authenticate the historical project.

Burton-Jones, A., & Grange, C. (2013 issue; online 2012). From Use to Effective Use: A Representation Theory Perspective. *Information Systems Research*, 24(3), 632–658. Publisher abstract. Effective use depends on how the system represents the user's task; the case imports no empirical effect size.

DeLone, W. H., & McLean, E. R. (2003). The DeLone and McLean Model of Information Systems Success: A Ten-Year Update. *JMIS*, 19(4), 9–30. Journal abstract. System, information and service quality are distinguished from use and benefit. Installation is not treated as a complete success measure.

Zhong, R., Yu, T., & Klein, D. (2020). Semantic Evaluation for Text-to-SQL with Distilled Test Suites. *EMNLP*, 396–411. Full original. Agreement on a single database can miss a wrong query. The case proposes discriminating fixtures and independently justified reference answers without transferring benchmark scores.

Barr, E. T., Harman, M., McMinn, P., Shahbaz, M., & Yoo, S. (2015). The Oracle Problem in Software Testing: A Survey. *IEEE TSE*, 41(5), 507–525. Author manuscript. Test oracles have soundness and completeness limits; passing a suite does not establish every answer's semantic truth.

NIST (2023). Artificial Intelligence Risk Management Framework, AI RMF 1.0, NIST AI 100-1. Official full text. Deployment-relevant evaluation, documented limitations and monitoring inform the proposed assurance process. The case claims no certification or demonstrated regulatory conformity.

PostgreSQL Global Development Group. Documentation, §3.5: Window Functions. Official reference. ROW_NUMBER assigns distinct sequential numbers within a partition, with unspecified selection order under ties. That principle corrects the source explanation without reconstructing the unavailable original query.

PostgreSQL Global Development Group. Documentation, §9.2: Comparison Functions and Operators. Official reference. Ordinary comparisons with NULL yield unknown. An empty filtered output therefore does not, on its own, establish an empty real-world event population.

Access scope is recorded privately: fresh full-text or abstract review is distinguished from continuing primary references verified earlier in the series.

RESEARCH REFERENCES / METHODS AND INTERPRETATION

Track correction uptake and preserve causal interpretation

Primary research and official methods inform the interpretation and proposed tests. They do not independently authenticate the historical project.

Tucker, A. L., & Edmondson, A. C. (2003). Why Hospitals Don't Learn from Failures. *California Management Review*, 45(2), 55–72. Author manuscript. A local workaround differs from system-level learning. VSC adds a proposed observable consumer-readback endpoint without importing hospital findings into this project.

Moreau, L., & Missier, P., eds. (2013). PROV-DM: The PROV Data Model. W3C Recommendation. Distinct entities and activities inform tracing an approved correction, later read and executed query. A provenance trail records derivation; it does not prove semantic validity.

NIST/SEMATECH. e-Handbook of Statistical Methods, §6.3.2.2: Individuals Control Charts. Official method. Ordered observations and process comparability govern interpretation. This case keeps revision comparisons separate from an operating process and treats short limits as provisional.

Laney, D. B. (2002). Improved Control Charts for Attributes. *Quality Engineering*, 14(4), 531–537. Publisher abstract. Dispersion adjustment requires compatible subgroup counts and exposures. Large data volume or rounded rates do not provide a valid adoption or answer-quality panel.

Cohen, J. (1960). A Coefficient of Agreement for Nominal Scales. *Educational and Psychological Measurement*, 20, 37–46. Publisher record. Kappa supplements proposed blinded coding, raw agreement and confusion matrices; agreement is not a substitute for a correct reference.

Barnett, A. G., van der Pols, J. C., & Dobson, A. J. (2005). Regression to the Mean: What It Is and How to Deal with It. *IJE*, 34, 215–220. Primary abstract. Selecting severe failures does not establish expected rollout performance. Reversion sensitivity in this case is explicitly hypothetical.

Collier, D. (2011). Understanding Process Tracing. *PS: Political Science & Politics*, 44, 823–830. Primary publication. Contribution tests distinguish analytical direction, execution and outcome ownership. Derivative attribution does not authenticate a sent report, accepted fix or sustained product adoption.

Access scope is recorded privately: fresh full-text or abstract review is distinguished from continuing primary references verified earlier in the series.

RESEARCH REFERENCES / METHODS AND INTERPRETATION

Evaluate promotion against value, learning and hard constraints

Primary research and official methods inform the interpretation and proposed tests. They do not independently authenticate the historical project.

Callaway, B., & Sant'Anna, P. H. C. (2021). Difference-in-Differences with Multiple Time Periods. *Journal of Econometrics*, 225, 200–230. Author manuscript. Group-time effects inform the separate later field rollout; that ATT is distinct from the controlled benchmark's randomized ITT.

The W. Edwards Deming Institute. The Funnel Experiment. Official exposition. Reacting to routine variation can worsen a process. The present case first asks whether changing definitions and as-of maturity make the observations comparable; it identifies no historical tampering effect.

Campbell, D. T. (1976). Assessing the Impact of Planned Social Change. Original institutional record. Consequential indicators may distort behavior. VSC tests the risk of optimizing usage or ticket closure while unresolved requests and incorrect releases disappear from the denominator.

Goodhart, C. A. E. Problems of Monetary Management: The U.K. Experience. 1975 conference; revised volume listed as 1976 in the RBA archive. The measurement critique informs the proposed separation of activity targets from decision-quality conditions; no historical manipulation is alleged.

Ittner, C. D., Larcker, D. F., & Randall, T. (2003). Performance Implications of Strategic Performance Measurement in Financial Services Firms. *AOS*, 28, 715–741. Author manuscript. This case does not infer operating performance from richer documentation, tests or measurement architecture alone.

OECD / European Commission JRC (2008). Handbook on Constructing Composite Indicators. Official handbook. Weighting, normalization and compensation require explicit choices. Greater usage or lower resource burden cannot offset a failed authority or answer-validity requirement.

Strong, M., Oakley, J. E., Brennan, A., & Breeze, P. (2015). Estimating the Expected Value of Sample Information Using the Probabilistic Sensitivity Analysis Sample. *MDM*, 35, 570–583. Author repository. The separate finite-state information model is derived here; it does not use this article's regression estimator.

Access scope is recorded privately: fresh full-text or abstract review is distinguished from continuing primary references verified earlier in the series.

FINAL REFEREE REPORT / INTERNAL AGENT-ASSISTED REVIEW

Approve the validation case within its evidence and design boundaries

[AI] This author-led review uses separate evidence, methods and numerical reviewers. It is not external peer review, client approval or certification of a firm ranking.

Major issue	Disposition in the revised case
Adoption and attribution	Local installation and analytical validation are distinct from product construction, enterprise use and repair ownership. Incomplete filing and acceptance evidence constrains the remediation claim.
Semantic and source validity	The exact-literal finding is an unweighted set comparison. Population-weighted answer quality remains unknown. Source explanations of duplicate survivors and identifier-based bounds require correction or qualification.
Clocks and persistence	Request resolution and query execution are not divided into a speedup. Later repeated text is not treated as a newly executed backend failure or a closed-cohort defect rate.
Evaluation and economics	The controlled domain benchmark retains all assigned requests. Synthetic gate laws, cost weights, exact inference and independent guardrails remain distinct from empirical field benefit.
Reproducibility and disclosure	Private evidence preserves raw inputs and version conflicts. Public arithmetic, model outputs and actual rendered pages are reviewed; client and system identities remain private.

Partner critique: can a COO act on page one?

[AI] Yes: require independent answer validation before promotion, and close remediation only after a later consumer receives and correctly uses the accepted version. Assign the product owner the promotion decision and data owners the corrections they control.

Minor revisions and remaining limits

[AI] The case separates coverage counts from answer rates, proposal from accepted repair, and the randomized benchmark from a later treated-population field effect. Unrecovered originals and unmeasured rollout remain explicit. Recruitment, source truth, resource prices and guardrail margins need pilot calibration.

REVIEW DISPOSITION — PASS WITHIN STATED SCOPE

[AI] The case supports documented adoption-validation work, new source-reconciliation findings and a reproducible prospective design. It does not establish enterprise adoption, completed remediation, mature customer benefit or realized savings.

Labels: [HR] historical record; [RC] recalculated; [AI] analytical inference; [PD] proposed design. Source basis S1; see private annex.