

OPERATING CASE STUDY / RISK, CONTROLS & GOVERNANCE

Exception Controls & Pilot Governance

Permission design needs a measurable and executable control perimeter

An operating-project case on tiered permissions for bulk status changes, a limited pilot and its conflicting performance reports. The analysis links control specifications, actual exposure, service recovery and the evidence required for expansion.

BULK PERMISSIONS & PILOT INTEGRITY

BLUF / BOTTOM LINE UP FRONT

Reconcile the pilot population. Test every control path before crediting expansion value.

RECALCULATED FINDING

Analytical finding: the reported combined rate falls below both reported subgroup rates. Under a common outcome, period and exhaustive partition, no nonnegative weighting can reconcile them. The proposed explanation for the headline therefore remains unverified.

CONTRIBUTION IN THE SOURCE RECORD

Recorded contribution: coverage planning, ownership transfer, intake routing and management reporting. The program team owned design and pilot analysis.

EVIDENCE STATUS

Source-register reconstruction; historical originals remain unrecovered. Design approval and pilot activity do not establish causal benefit.

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

STRUCTURED ABSTRACT

Expansion needs an evidence contract

[AI] Retain the approved control rationale. Release each expansion decision only when the population, executed permission and physical outcome can be reconciled.

CONTEXT

[HR] A distributed service network used bulk virtual status changes alongside normal physical processing. The register describes routine use, incorrect responsibility coding and competing risks: permissive access could mask unresolved work; restrictive access could delay legitimate exceptions. A technical dependency change created additional pressure on the operating model.

OBJECTIVE

[PD] Establish whether tiered permissions can prevent unauthorized changes while preserving timely, legitimate recovery. The investment question concerns net mature service and resource consequences, rather than a reduction in submitted requests alone.

APPROACH

[RC] Reconcile the source register, permission versions, pilot membership, rate arithmetic, concentration measures and planning claims. [PD] The Permission-and-Exposure Reconciliation Contract (PERC) links a versioned permission to its executed population, fallback service and independently measured outcome. A separate finite operating model tests candidate evaluation designs.

FINDINGS

[RC] The reported total lies outside the interval defined by the subgroup rates under the stated common-basis interpretation. [AI] This blocks the claimed aggregate explanation before causal attribution is considered. A limited permission pilot also cannot validate untested approval and central-service components. Recorded handover, approval and implementation stages are distinct.

DECISION VALUE

[PD] The sponsor should require the metric owner to reconcile identities and rate definitions, the control owner to demonstrate rules on every route, and the service owner to accept overflow recovery. Release requires positive net resource value, improved physical outcomes and a prespecified legitimate-service safeguard. Unresolved evidence keeps the affected expansion or benefit claim on hold.

IMPLICATIONS

[AI] The case supports a controls-design and measurement-integrity mandate. It does not establish historical causal savings. Information value is modeled separately as improvement in a feasible decision, net of evidence cost; it is never added to the same operating benefit.

[HR] historical record; [RC] recalculated; [AI] analytical inference; [PD] proposed design. S1: private source and claim ledger; methods and references follow.

EXECUTIVE EXHIBIT 1 / RECONCILIATION

Reconcile the total before explaining its drivers

[RC] An aggregate below both exhaustive subgroup rates cannot be a positive exposure-weighted average of those rates on the same basis.

RECONCILIATION / A NECESSARY AGGREGATION TEST

Invented rates, one outcome and exhaustive positive-exposure groups



CHANGING POSITIVE WEIGHTS CANNOT CLOSE THIS GAP

A different numerator, denominator, population or version requires an explicit bridge.

The condition survives a common positive index scale; separately rebased group indexes cannot be pooled.

[HR] The register explains a nearly flat pilot headline as improvement in one business line offset by deterioration in another. It also supplies subgroup rates and shipped-volume denominators. [RC] The aggregate fails a necessary arithmetic condition before any choice of weights: it is below the lower subgroup rate in each reported week. Ordinary display rounding does not close that gap.

Required common basis	Reconciliation evidence
Outcome and grain	Same event definition, unit identity and counting rule.
Population and window	Same maturity cutoff; exhaustive, nonoverlapping groups.
Exposure and aggregation	Matched numerator and denominator; weights sum to one.

[AI] Different outcomes, scopes or extraction versions could explain the records. The result rejects their simultaneous interpretation as one coherent decomposition; it does not identify which record is wrong. Separately published pilot rates need an explicit bridge before comparison, and a network aggregate containing pilot sites is not an independent control.

DECISION / METRIC OWNER

[PD] Hold the driver waterfall and associated benefit calculation. Recover numerator, denominator, membership and query versions; recompute both publications from the same frozen extract. Preserve unreconciled alternatives until that bridge exists.

S1: pilot outcome and conflicting-publication records; private reconciliation calculations. The figure is an invented construction, not a transformed historical result.

EXECUTIVE EXHIBIT 2 / EXPOSURE

A limited pilot supports only the permissions it tested

[AI] Pilot participation is insufficient exposure information when the intervention contains distinct self-service, approval and central-processing components.

COMPONENT EXPOSURE / DISTINCT EVIDENCE REQUIREMENTS

Conceptual architecture; historical pilot coverage is narrower than the full design

LIMITED SELF-SERVICE
Pilot activity recorded
Outcome still unverified

APPROVAL WORKFLOW
Target and version records
Execution needs validation

CENTRAL FALLBACK
Required for exceptions
Capacity needs validation

PRECISION WITHIN ONE COMPONENT DOES NOT CREATE EVIDENCE FOR AN UNTESTED COMPONENT

Separate design approval, exposure, implementation fidelity and mature causal outcomes.

[HR] The architecture separates routine low-risk self-service, approval-dependent changes and centrally brokered exceptions. The pilot description emphasizes a limited self-service component; the native approval workflow appears with a later delivery target. Other records give different launch dates and site membership. These may describe different milestones, versions or populations.

Evidence level	What it can establish
Specification / approval	Intended permission and accepted design scope.
Deployment / execution	Version and route operating for each original event.
Mature outcome	Physical result under that exposure, with a valid comparison.

[AI] An approved design cannot supply missing runtime evidence. A reported launch date cannot establish that every component became active at once. Evidence from a narrow permission path has no direct support for unexposed paths, and a pilot selected for high usage needs additional transport assumptions before network-wide claims.

DECISION / CONTROL OWNER

[PD] Maintain a component-by-location exposure ledger with policy version, activation time, exception route and verified execution. Test each component before assigning it an effect. Expand only where authority, measurement and service capacity support the intended treatment.

S1: whitepaper, milestone, pilot-read, handover and expansion records. Treatment-version and external-validity research is applied in the technical layer.

EXECUTIVE EXHIBIT 3 / RULE SEMANTICS

A threshold needs its complete predicate

[AI] Raising a number can widen access, narrow access or affect a different processing route. The inequality and its scope determine the change.

PREDICATE DIRECTION / THE NUMBER ALONE IS INCOMPLETE

Conceptual rule comparison, holding the population and all other conditions fixed

UPPER PERMISSION CAP	Allow when quantity is at most cap Raising cap widens eligibility
MINIMUM ROUTING VOLUME	Route when quantity is at least minimum Raising minimum narrows this route
CENTRAL BATCH CAPACITY	How much work one task can process Capacity is separate from permission

Compare the same rule version, inequality, actor, event scope, clock and fallback before judging a threshold change.

[HR] The record contains per-request limits, cumulative limits, site-size bands, central batch capacity and minimum-volume language. A later threshold decision is criticized in the register because an earlier gate allegedly never fired. [AI] That criticism requires proof that the numbers govern the same predicate, population, time window and fallback path.

Predicate form	Effect of raising its threshold
Permit if quantity ≤ cap	More quantities become eligible, all else equal.
Route if quantity ≥ minimum	Fewer quantities enter that route, all else equal.
Central batch capacity	More work can fit a task; authorization is separate.

[PD] Test immediately below, at and above the boundary; repeated requests; overlapping actors; clock rollover; missing physical evidence; unavailable approvers; and failed execution. Record whether each condition changes permission, routing, workload or service. An absent trigger can mean weak coverage, inapplicable traffic or a broken implementation.

DECISION / POLICY AND PLATFORM OWNERS

[PD] Approve the complete actor–object–action–condition rule. Bind the tested predicate to the live version and log the enforcement result. Preserve authorized recovery when a request is outside the permitted scope.

S1: tier design and later threshold decisions. NIST SP 800-162 distinguishes access decisions, attributes and enforcement; the application here is an architectural analogy.


EXECUTIVE EXHIBIT 4 / EVENT INTEGRITY

Cumulative controls need a shared event ledger

[PD] Evaluate original work obligations and executed quantities so that splitting or relabeling requests cannot redefine the measurement population.

ORIGINAL EVENT / REQUEST PARTITIONS / CUMULATIVE LIMIT

Wholly hypothetical: 12 unique units, per-request cap 6 and window cap 10



[HR] The design combines per-instance and time-window limits, while reports emphasize submission counts and final status categories. [AI] Neither count is a complete physical workload measure. An underlying event can be partitioned into different requests, and different statuses can describe the same unresolved obligation.

Ledger	Required identity and state
Original work	Unit, originating event, promised outcome and eligibility.
Requested / authorized	Actor, quantity, reason, policy version and approver.
Executed / recovered	Unique changed units, original state, physical remedy and time.

[RC] In the invented example, each submitted quantity is below the per-request cap while their cumulative total exceeds the window cap. [AI] Shared accounting and reliable identity linkage are separate requirements from checking each request. This is a potential failure mode, not evidence that historical users split requests. Merging distinct events incorrectly can also deny legitimate work.

DECISION / OPERATIONS AND DATA OWNERS

[PD] Audit deduplication and event linkage on accepted and rejected work. Reconcile cumulative quantities across actors and routes before execution, with an authorized fallback for legitimate overflow. Retain unresolved and diverted work in the original outcome population.

S1: per-instance/cumulative permission design. The figure is wholly hypothetical. No historical request splitting or individual misconduct is asserted.

EXECUTIVE EXHIBIT 5 / RECOVERY CAPACITY

Design restriction and recovery together

[AI] A control can reduce an impermissible action while leaving legitimate service unresolved. The latter remains part of the operating result.

CONTROL AND SERVICE / FOLLOW EVERY ORIGINAL OBLIGATION

Proposed acceptance chain with an independent physical endpoint



BURST TEST / STAFFING / BACKUP / COMPLETE COST
Acknowledgement, permission, execution and recovery have separate completion clocks.

A restrictive rule can pass a permission test while its recovery service fails. Test and cost both.

[HR] The source identifies central workload, manual handling time, approver delay and large-event exceptions as constraints. It also describes self-service for selected low-risk work. [AI] Blocked work may transfer to approved central processing, return to normal physical flow, require rework or remain unresolved. It does not automatically become saved labor or central demand.

Separate clock	Acceptance evidence
Acknowledgement	Request received; no inference of completion.
Approval and execution	Authorized actor accepted and performed the permitted change.
Physical resolution	Original service obligation reached a verified outcome.

[PD] Stress burst size, cumulative-cap usage, simultaneous requests, missing evidence, absent approvers and central capacity. Cost all review, retry and handling work, including failed attempts. Average utilization leaves nominal headroom but cannot certify service during bursts; task-size and staffing distributions determine the bottleneck.

DECISION / ACCEPTING SERVICE OWNER

[PD] Pair restricted paths with staffed, authorized recovery and a deadline. The evaluation must pass a legitimate-service safeguard and an economic test. If the safeguard fails, change the operating package; more observations cannot repair its service design.

S1: workload, utilization, approval, contingency and regional design records. Proposed capacity and failure laws are disclosed with conditional assumptions in the model appendix.


EXECUTIVE EXHIBIT 6 / BENEFIT DISCIPLINE

Value the mature service contrast

[AI] A planning line, a process target and a reported pilot movement answer different questions. None supplies the missing causal or cost bridge.

BENEFIT CLAIM / FROM EXPOSURE TO NET INCREMENTAL VALUE

Proposed evidence sequence; no historical monetary benefit is validated here



[HR] The record carries multiple entitlement metrics and a later planning-year value, alongside unresolved overlap, incremental-year scope and accepting ownership. The pilot reporting basis remains unreconciled, and cost-avoidance follow-up was still recorded as open. A zero in a benefit tracker establishes that reporting entry; it does not identify a zero physical effect.

Value layer	Evidence required
Outcome contrast	Same mature original population under defined policy versions.
Resource consequence	Avoidable loss, signed service harm and full incremental work.
Net incremental value	Overlap removed; year and fixed/recurring cost boundaries explicit.

[PD] Apply the economic release threshold to the implemented package. Compare minimum detectable effect with break-even in the same units and horizon. Also estimate the chance of clearing the actual decision threshold: ability to detect a nonzero effect does not guarantee a decisive investment test.

DECISION / FINANCE AND OPERATING SPONSOR

[PD] Hold benefit credit until the metric bridge, causal contrast, cost basis and overlap ledger are accepted. Separately price information that could change that decision. Its value is net of evidence cost and cannot be added again to the same avoided outcome.

S1: entitlement, planning-row, overlap, tracker and open-action records. Model valuation and power numbers are hypothetical design sensitivities, not historical benefits.

EXECUTIVE EXHIBIT 7 / CONTRIBUTION

Transfer evidence obligations with delivery ownership

[HR] The recorded management contribution concerns continuity, accountability, intake and reporting. Design authorship and measured outcomes remain separate claims.

MANAGEMENT CONTRIBUTION / HANDOVER WITH ACCEPTANCE

Recorded roles remain bounded by derivative evidence and original-artifact recovery

PORTFOLIO LEADER Coverage + reassignment	PROGRAM TEAM Design + pilot analysis	OPERATING OWNERS Authority + recovery	METRIC / FINANCE Reconciliation + value
--	--	---	---

HANDOVER INCLUDES THE UNFINISHED EVIDENCE
Carry versions, open decisions, metric gaps, accepting owners and closure tests together.

Recorded management action is separate from implementation success, causal benefit and the new PERC analysis.

Recorded personal contribution	Evidence grade and boundary
Coverage / ownership transfer	Provisional registry-level hoop; supports a managerial action, not execution success.
Intake routing / reporting	Provisional registry-level hoop; original authorship and reporting artifacts need authentication.
Review audience / future ownership	Receipt supports involvement only; proposed future ownership is not confirmed delivery.

[HR] The program team authored the design, review materials and pilot analysis. The register describes a transfer including an unfinished measurement obligation. [AI] A handover can be real while analysis remains with the previous owner; an unassigned parallel intake item does not erase other governance or activity. Neither record establishes closure of all delivery obligations.

DECISION / PORTFOLIO LEADER

[PD] Transfer a signed inventory of policy versions, live routes, unresolved measurement bridges, accepting service owners and benefit commitments. Name each closure owner and acceptance evidence. Separate the leader’s recorded contribution from the PERC analysis developed for this case.


S1: management, intake, handover, reviewer, program-team and proposed-owner records. Grades assess derivative evidence internally; they do not claim external peer review.

TECHNICAL ANALYSIS

Separate permission conformance from physical service

[AI/PD] The Permission-and-Exposure Reconciliation Contract (PERC) defines the actual policy version, tests its execution and measures customer service on the original cohort. A valid digital transition is neither proof of physical custody nor proof of prevention.

PERC: A VALID PERMISSION DOES NOT CERTIFY A CUSTOMER OUTCOME



Context U: original need, workload, role, site and capacity affect L, F, D, H, W and Y.

Y -> O and J -> O; V -> R. Recording rules can change R without changing Y.

Z -> K; L / F -> W are summarized resource paths. Policy, version and authority are frozen before assignment.

[PD] Assignment Z changes event-linked enforcement L and authorized fallback F. Decision and execution D can alter digital state V and reported metric R; physical action H changes original-promise failure Y. Independent ascertainment O depends on Y and reference quality J. Independence refers to the measurement process, not statistical independence from truth.

[AI] Workload, genuine need, role and capacity affect both implementation and service. In an observational pilot they also influence selection and effective exposure. Conditioning on approved requests, completed transitions or successful fallback discards affected cases after intervention and can bias the comparison.

[PD] Actual review and execution work W is charged alongside physical harm; fixed package cost K is charged once. Randomize independent capacity domains under the same authorized rule. Repair known dangerous implementation faults before a live comparison; shadow replay can test those faults without deliberate unauthorized execution.

Basis: proposed PERC DAG [PD]; source interpretation [HR/AI]. Hu, Kuhn and Yaga (2017); NIST ABAC guidance. Arrows specify testable paths, not historical causal findings.

TECHNICAL ANALYSIS

A combined rate must first satisfy its own population identity

[RC/AI] Under one exhaustive, disjoint partition, common period, event definition and nonnegative exposure measure, the combined rate is an exposure-weighted mean of subgroup rates. It cannot lie below every subgroup rate.

Independent hypothetical rate	Value
Group A	4%
Group B	10%
Reported combined rate	3%
Any coherent weighted mean	4% through 10%

[AI] With positive exposures n_A and n_B , $r = (n_A \cdot r_A + n_B \cdot r_B) / (n_A + n_B)$. Thus $\min(r_A, r_B) \leq r \leq \max(r_A, r_B)$. Even rounding each displayed example by plus or minus 0.05 percentage points leaves a gap. No choice of mix repairs the example; the private rational solver checks all enumerated weights and the bound.

[HR/RC] The historical Total Lost table fails this necessary condition under its stated combined-versus-subgroup reading, including rounding tolerance. The defensible result is a conditional incompatibility, not a declaration that all source data are corrupt. Metric identity, membership, denominator, extraction vintage or transcription may differ. Another component can approximately reconcile without authenticating the whole table.

[HR/AI] The favorable mailer and detailed pilot read use different metric labels and potentially different shipment bases or site cuts. Their disagreement cannot be resolved by selecting one headline. Nor can the incompatible combined table establish that one subgroup erased another subgroup's causal improvement.

[PD] Rebuild numerator and exposure from linked entities, retain both source versions, and require additive closure before interpreting mix. Read all-cause physical outcomes separately from accounting labels. A network comparison supplies context; without comparable exposure and an identifying design it is not a counterfactual.

Basis: exact convex-hull proof [AI] and conditional source arithmetic [RC], counterexample.py/json and private evidence calculations. All displayed quantities are synthetic; no source rates are published.

TECHNICAL ANALYSIS

An observed tier cannot certify untested permission versions

[AI] Even an ideal randomized pilot cannot identify the effect of a different, unobserved implementation without transport assumptions. The historical selected-site pilot provides less identification than this synthetic example.

Same pilot evidence; different target worlds	World A	World B
Target share represented in pilot	40%	40%
Pilot failure reduction	5 pp	5 pp
Reduction in unsupported versions	10 pp	-20 pp
Target-average failure reduction	8 pp	-10 pp

[AI] Define benefit as usual minus offered failure risk. Then target benefit = $w \cdot \text{pilot benefit} + (1-w) \cdot \text{unsupported benefit}$. The two worlds have identical pilot risks of 20% versus 15%, yet opposite target decisions. With $w=0.4$, pilot benefit=0.05 and no further information on binary risks, the sharp target bound is [-0.58,0.62]. Those deliberately wide bounds concern missing support, not sampling error.

[HR/AI] The source describes a low-risk self-service pilot, interim central execution, a native approval workflow with delivery unconfirmed, and a later planning-row mechanism. They are not interchangeable versions of one treatment. Review in other regions does not demonstrate deployment or local suitability. A milestone marked complete does not establish every role, endpoint and fallback as exposed.

[PD] Build a role-by-operation-by-version support matrix with actual configuration and execution times. Distinguish planned launch, configured access, first eligible attempted use and sustained enforcement. Analyze comparable version cohorts, then state explicitly which populations and components remain outside support.

[PD] Transport requires measured effect modifiers, adequate overlap and defensible version consistency. Weighting cannot create observations for a missing tier. Extend testing to unsupported components or retain bounds; do not erase the gap with a more precise pilot estimate.

Basis: independent potential-outcome construction [AI]; VanderWeele and Hernan (2013), Tipton (2014). The illustrative pilot is hypothetical and stronger than the historical design.

TECHNICAL ANALYSIS

Write the permission predicate before interpreting the threshold

[HR/AI] A threshold number alone does not define what the system permits. The record mixes per-instance caps, cumulative limits, central-processing constraints and an incompletely specified later threshold change.

Rule type	If its threshold increases
Maximum allowable quantity: $q \leq c$	The allowable set expands, other predicates fixed.
Minimum eligibility floor: $q \geq c$	The eligible set contracts, other predicates fixed.
Cumulative event/window cap	Requires stable event identity, deduplication and atomic quantity accounting.
Authorized exception	Permits only the signed scope, operation, state and time window; unavailable fallback confers no authority.

[AI] In a synthetic example, one 12-unit event split into three requests of 4 satisfies an individual cap of 6 but exceeds its cumulative window cap of 10. This is not evidence that the historical tool allowed splitting. It shows why request-level success cannot certify a daily or event-level permission. Replays should cover splits, duplicates, concurrency, stale role membership and expired exceptions.

[HR/AI] The source infers that raising a later threshold makes a gate less likely to fire. Without the action predicate, that direction is undefined. Raising a minimum contracts the qualifying set; raising a maximum expands the allowable set, other predicates fixed. Confirm the operator, event unit, effective version and exposure distribution before calling the change contradictory. An untested channel may bypass an unchanged limit.

[PD] Store original event ID, package IDs, actor role, requested operation, current/target state, quantity, cumulative total, approval scope, channel and version; retain deny and unresolved results. Test allowed and forbidden transitions against an independently approved policy oracle. Unknown authority follows the existing protected path, never a guessed permission.

[AI] Hu and colleagues motivate model-to-implementation conformance testing. Park and Sandhu's continuing usage-control model motivates mutable quantities and obligations. Neither establishes physical custody, customer value or a historical implementation flaw in this case.

Basis: source predicate ambiguity [HR/AI]; exact nested-set and cumulative-count examples [AI]. NIST SP 800-192; Park and Sandhu (2004); NIST SP 800-162. All numeric examples are synthetic.

TECHNICAL ANALYSIS

Preserve original exposure and distinguish four estimands

[PD] A permission change can reduce executions, change status labels, redirect workload and alter physical receipt differently. The evaluation must keep those pathways visible under one fixed original booking cohort.

Estimand	Population and contrast
Implementation conformance	Allowed/forbidden decisions under the approved version, including attempted alternate routes and unresolved evidence.
Legitimate service harm	Offer minus usual promise-failure risk among independently adjudicated genuine original emergency opportunities.
Customer ITT benefit	Usual minus offer original-promise failure among all externally booked units, regardless of submission or execution.
Net resource benefit	Signed avoided physical loss minus added work and fixed offer cost, same cohort, horizon and valuation.

[PD] Randomize domains containing shared daily counters, staff and fallback capacity. Freeze externally scheduled original manifests and pretreatment strata before assignment. Keep unsubmitted, denied, split, resubmitted and alternate-path work in the denominator. Equal scheduled cohort size makes the modeled domain mean equal the package-weighted contrast; unequal field exposure needs prespecified weights and revised covariance.

[PD] Record event time, authorization time, write time, physical handoff, original promise and extraction vintage separately. Follow-up at original promise plus 28 days reconstructs whether the original promise was met; eventual receipt does not erase a missed deadline. Unknown physical outcomes receive worst-case bounds and maturity sensitivity rather than deletion.

[HR/AI] Reported dates and site lists may reflect a rollout ramp, separate components or source errors. They do not yet support a single treatment breakpoint. Changes in attribution or status capture can change concessions independently of physical outcomes. A source runbook description does not establish that its query was executed or that its output reproduces the disputed rate.

[PD] A later staggered observational study could target group-time ATT with credible parallel trends, no anticipation, comparable versions, overlap and modeled spillovers. Plot pretrends and use equivalence and sensitivity analyses; a nonsignificant pretrend test does not prove parallel trends. These additional assumptions are not granted retrospectively to the historical pilot.

Basis: proposed original-cohort ITT and measurement contract [PD]; source timing/metric limits [HR/AI]. Treatment-version, interference and group-time ATT methods applied with their stated assumptions.

TECHNICAL ANALYSIS

Treat selection, variation and effect modification separately

[HR/AI] The pilot selected high-volume sites. Volume selection affects representativeness and exposure weighting; it is not automatically selection on an unusually adverse outcome. A later focus on poor performers introduces a different selection problem.

[AI] For an explicitly assumed linear conditional-mean model, untreated follow-up is $m + \rho \cdot (x - m)$, where m is the population mean and ρ is persistence. Assumed ρ values 0.4, 0.7 and 0.9 imply 60%, 30% and 10% reversion of an outcome-selected excess. These are hypothetical sensitivities, not estimated fractions of this pilot's movement.

[HR/AI] Flat weekly exposure does not fix daily mix or capture. Bad days warrant investigation, not an automatic substitution claim. Removing the worst site after seeing outcomes changes the estimand. A falling share can coexist with a rising count.

[PD] Prespecify baseline cap pressure, emergency mix, configured version, available fallback capacity and calendar as modifiers. Realized exposure and readiness are mediators, not ITT conditioning strata. Test interactions; opposite weekly changes do not identify opposite treatment effects.

[PD] Stratify a probability sample across original approved, denied and unsubmitted cases, versions, nodes and emergency classes. Duplicate coders blind to arm and downstream outcome classify contemporaneous policy evidence; independent reviewers establish physical outcomes. A hypothetical simple random sample of 400 gives worst-case approximate 95% agreement precision of ± 4.9 percentage points, not kappa precision. Report weighted confusion tables, prevalence and kappa with domain-bootstrap uncertainty; cluster sampling changes precision. Adjudicate disagreement independently and test reference invariance across versions and baseline modifiers.

[RC] The private audit has 20 source/sensitivity screens, with 12 version audits, 10 design/placeholder audits and 8 cross-sectional/range profiles separate. All fitted source individual and moving-range points remain within tentative bands; no stable operating phase is established. Listed daily rates retain every supplied day; missing days are not fabricated. Zero-baseline placeholders have no valid baseline-100 index.

[PD] For a valid ordered phase, $MR_t = |x_t - x_{(t-1)}|$; individual limits are $\text{mean}(x)$ plus or minus $2.66 \cdot \text{mean}(MR)$, with MR upper limit $3.267 \cdot \text{mean}(MR)$. Sparse pairs are non-diagnostic. No repeated valid binary-count Laney panel exists here. Overdispersion adjustment cannot repair unverified exposure, taxonomy or maturity.

Basis: source selection limits [HR/AI], Barnett et al. (2005), Sun et al. (2010), NIST XmR and Laney (2002). No historical treatment effect or regression-to-the-mean fraction is estimated.

TECHNICAL ANALYSIS

Bound management contribution and the value claim

[HR/AI] Personal contribution is a provisional registry-level hoop: staffing coverage, transferring programme responsibility, specifying intake routing and reporting open measurement debt. Originals remain unrecovered. Those acts are distinct from authoring the whitepaper, setting thresholds or measuring a customer effect.

Claim	Evidentiary limit
Design and implementation	The registry names programme and technical owners. A manager's review or retained deck is not sole delivery ownership.
Managerial decisions	Coverage and transfer are attributed through derivative records; the content and execution need original authentication.
Later planning responsibility	Proposed ownership and an open Finance task do not establish accepted accountability or realized value.
Customer and financial outcomes	No identified historical physical treatment effect or validated incremental cash benefit is supplied.

[HR/AI] Repeated summaries can share one source. Calendar acceptance supports intended participation, not attendance or the substance of a decision. An unassigned intake does not prove every operating action was unapproved; a completed milestone does not prove all endpoints were live. Treat both as reconciliation needs.

[HR/AI] Workload scenarios mix task and package grains, alternative daily volumes and annualization assumptions. Capacity-equivalent hours are not payroll savings. Central demand need not equal removed self-service volume: legitimate needs can disappear, reroute, wait or require repeated review. Charge actual fallback and rework rather than imposing a conservation identity.

[PD] Use net = favorable avoided loss*a*h - added physical harm - incremental work - fixed cost, where a is realizable avoidability and h is the non-overlapping share. Discount favorable value only. Reconcile parent and child benefit claims using linked outcome cohorts and common prices; a difference between bookings is not a double-count amount.

[AI] Evidence gaps jointly limit attribution, transport and valuation. Pilot support, reference truth, fallback readiness and rule conformance are separate gates. Passing one cannot substitute for the others; an unpriced customer or capacity harm remains a decision limit.

Basis: bounded process tracing [HR/AI], Collier; source ownership and Finance limits; proposed signed valuation [PD]. Original authentication status will be bound to the final evidence audit.

TECHNICAL ANALYSIS

Test the whole approved offer, including emergency access

[PD] The proposed trial compares current approved implementation with stronger event-linked cumulative enforcement and documented authorized fallback. It tests a complete operating package, not the efficacy of forbidding a status label.

Design element	Frozen prospective contract
Assignment and exposure	Independent capacity domains; externally scheduled original bookings fixed before assignment, with all genuine and other opportunities retained.
Mechanism	Residual undetected wrong-scope acceptance versus stronger identity/counter checks; actual fallback readiness affects legitimate service.
Reference	Policy eligibility independently adjudicated; physical original-deadline outcomes ascertained by promise plus 28 days.
Joint decision	Net benefit beyond fixed cost, all-manifest physical improvement, and genuine-emergency harm below an approved margin.

[PD] The numerical prototype uses shared daily implementation and fallback states plus persistent domain heterogeneity. Conditional finite outcomes give exact means and covariance. Studentized-normal planning and finite-process simulation are conditional approximations, not universal confidence guarantees. The author tests selected allocations, actual zero/economic/safety boundaries and a stressed implementation; an independent solver reproduces the reported calculations.

[PD] Compare an economically material break-even effect with the two-sided 80%-power MDE, then check joint directional gates. High power to detect a large overall effect does not certify legitimate access. A hypothetical noninferiority margin needs Quality approval; passing it permits bounded harm and does not mean no harm.

[PD] Low fallback availability can make the mean economic result negative and emergency harm unacceptable. More observations cannot repair that operating failure. A staffed fallback is a costed treatment redesign with its own effect and verification, not a sample-size adjustment. The hypothetical probabilities are not fitted to historical pilot rates.

[PD] Before field recruitment, calibrate permission errors, fallback latency, reference disagreement, missing outcomes, covariance and operational capacity. Research ascertainment is separate from production verification cost. Preserve existing customer remedies and repair known dangerous faults in both arms. EVSI compares information actions only; its priors are disclosed assumptions.

Basis: new PERC prospective design [PD]; [model.py/results/pages](#) and independent audit. Eldridge et al. (2016) motivates staged feasibility and progression; it supplies no effect sizes or sample sizes.

TECHNICAL ANALYSIS

PERC yields four propositions that can fail

[AI/PD] A permission pilot supports scale only when its populations reconcile, its executed rules match the intended perimeter, its tested versions cover the target, and legitimate service survives the change.

Proposition	Prespecified falsification test
P1: Reporting is coherent	Fails when linked common-scope numerators and exposures cannot reproduce both subgroup and total rates within declared tolerance.
P2: Permission enforcement is complete	Fails if permitted/forbidden replay cases, split-event accounting or alternate endpoints fail the approved conformance target.
P3: Pilot evidence supports expansion	Fails as an empirical transport claim when target versions lack support and no justified bridging assumptions close the gap.
P4: The package creates net service value	Fails if the original-cohort economic or physical-benefit gate fails, or genuine-emergency harm exceeds its approved bound.

[AI] The new contribution is the combination of a convex-hull falsification and a treatment-version support proof. One rejects an internally incompatible mixture interpretation; the other shows why perfect pilot precision can still leave the target decision unidentified. The threshold proof adds a necessary operator-level question before criticizing a numeric change.

[AI] Deming cautions against changing permissions in response to unstable weekly swings. Campbell and Goodhart predict adaptation when submission or status counts become targets. Ittner, Larcker and Randall motivate testing the permission-to-service causal chain; Tucker and Edmondson motivate repairing the work that produces workarounds. Cost-weighted composites need visible signed physical and resource components, as OECD guidance emphasizes. Cohen agreement can audit reference consistency but cannot establish truth.

[PD] NIST conformance tests and usage-control theory make the rule executable; treatment-version and generalizability research bound expansion; subgroup and pilot-trial methods discipline follow-up. Keep those functions distinct. Reproducibility requires the source ledger, formula and proof code, model law, independent numerical review and final PDF/raster bindings. Authentication, feasible staffing and customer safety remain external decision gates.

Basis: PERC synthesis [AI/PD], exact proof and fresh primary methods. These propositions are proposed and falsifiable; none is labeled a historical implementation or outcome achievement.

Protect the whole event, including its overflow

[PD] Compare the current approved request implementation with event-linked cumulative controls and an authorized fallback. No arm knowingly permits an invalid transition. Replay permissions first; correct known critical faults symmetrically before considering live randomization.

[PD] Randomize independent shared-authority domains. Each contributes eight externally scheduled days of 320 original bookings: 80 genuine emergency-overflow units, 80 other over-cap units and 160 ordinary units daily. Each overflow class contains four original batches of 20. Class membership is independently established before assignment, never inferred from submission or treatment.

Wholly hypothetical parameter	Usual	Strengthened
Correct daily ledger probability	60%	95%
False scope acceptance: good / faulty ledger	2% / 50%	2% / 50%
Timely authorized genuine fallback	90%	85%
Genuine failure: timely / unavailable	3% / 18%	3% / 18%

[PD] Known absent authority never defaults to execution. False acceptance represents an undetected scope/linkage fault. Other over-cap units fail at 16% if wrongly executed versus 4% on the standard route; ordinary risk is 2%. Denial itself earns no benefit.

A DAILY CAP IS NOT A REQUEST COUNT

[PD] Freeze original batch IDs, role, transition, approved daily quantity and customer promise. Split or replayed requests cannot create new authority. Independently reconstruct original-promise receipt through promise + 28 days; late delivery remains a miss. Unsubmitted, denied, cancelled and unresolved units remain.

[PD] All probabilities are synthetic. An equiprobable persistent domain grade shifts ledger accuracy by +/-3 points and fallback by +/-4; daily states are conditionally independent. Eight days are not eight clusters.

Staffing can repair a mean failure that sampling cannot

[PD] Exact latent-state and conditional-binomial moments support approximate normal planning. The response path mixes whole-domain usual and strengthened laws; it is not observed compliance or an interpolation of classifier performance.

Design / per arm	MDE index	BE index	BE detection	Joint floor
Initial / 128	5.66	3.34	37.4%	41.8%
Base / 512	2.86	3.34	90.7%	95.3%
Weak fallback / 512	No positive path	3.34	Not attainable	<0.1%
Staffed version / 512	2.88	5.85	>99.9%	>99.9%

[PD] Require 80% detection at break-even and 80% joint success: resource benefit above fixed cost, original-promise improvement above zero, and genuine-case harm below an illustrative two-percentage-point margin. That margin needs approval and uses the genuine 80-unit daily denominator. Base harm is +0.75 points.

[PD] Each directional gate uses a nominal 97.5% normal bound. The MDE is 80% correct-direction rejection of the two-sided 5% zero test. Joint floor = max(0, sum of three marginal powers - 2); it is approximate, not an exact finite guarantee.

[PD] Fallback falling to 55% produces +5.25-point genuine harm and negative net value. More domains cannot repair that mean. A staffed version raises fallback to an assumed 94%, with additional reserved cost. It is a different treatment requiring fresh validation.

CALIBRATE BEFORE PROPOSING RECRUITMENT

[PD] The candidate requires 1,024 independent domains and 2,621,440 original bookings. Recruitment and repeat-day covariance are unestablished. Finite baseline joint success is 95.15% (Monte Carlo 95% interval: 94.71-95.55%). Separate boundary simulations do not certify field power.

[PD] Fixed base usual resource burden = 100; MDE and cost use index points. Four direct designs, no minimum claim. Ordinary-flow and severe-harm safeguards need separate calibration.

Price the legitimate work the gate must still serve

[PD] Failure costs 30 hypothetical resource units. Both arms pay 0.3 per batch assessment, 2 per genuine fallback attempt, 1 per timely genuine authorization, and 0.025 reservation per original booking. Attempts cost money even when fallback is unavailable.

[PD] Strengthened event linking adds 0.012 per original booking. Fixed base cost $K = 0.04$ allocates 0.015 to configuration, 0.010 to authority maintenance and 0.015 to recurring quality checks. Staffed $K = 0.07$ adds 0.03 reserved fallback capacity. Research-only reference adjudication is separate.

[PD] $Net_raw = a \times h \times S - A - W - K$. S is favorable physical-failure cost reduction; A is added harm; W is incremental work. Avoidability a and non-overlap h discount favorable benefit only. Added genuine harm and all extra work remain fully priced.

Scenario	$a \times h = 1.00$	0.75	0.50
Base	3.64	0.48	-2.68
Weak fallback	-24.24	-27.40	-30.56
Staffed	9.50	5.40	1.30

[PD] Base realization must exceed 0.712; staffed realization must exceed 0.420. Weak fallback cannot break even even at full realization. Reprice and re-power if costs, batch mix or response change; fewer requests and refunds withheld have no value term.

UNCERTAIN EVENT IDENTITY CAN ERASE THE MARGIN

[PD] With independently verified work, epsilon unknown or wrong physical labels per arm allow 60 epsilon raw resource-contrast error and 2 epsilon risk error. Fractions use each endpoint's original denominator. A proved affected-unit share from false event merges has the same bound. Unknown genuine membership needs reassignment bounds. Later evidence cannot rewrite the original promise.

[PD] $Net_index = 100 \times Net_raw / 1.19705$. Full original-unit pool, fixed cost allocation and baseline retained across versions. No historical prices or realized savings.

PROPOSED INFORMATION VALUE / 4 OF 4

Buy evidence about the constraint before scaling

[PD] Choose a bounded operating test, staff the fallback before testing, or reconcile authority and event identity. This separate information decision does not assume the operating package succeeds or confer permission to deploy.

Action / state	Ready	Weak fallback	Unresolved
Bounded operating test	13	-9	-16
Staff fallback, then test	7	12	-7
Reconcile authority / scope	2	3	6

[PD] Hypothetical priors are 0.40 / 0.35 / 0.25. Failed mandatory authority or lineage takes precedence as unresolved. Otherwise assess ready versus weak fallback. A signal identifies the state with probability q; each other label occurs with probability (1 - q)/2.

Signal q	Gross EVSI	Acquisition	Net EVSI
0.50	0.33750	0.50	-0.16250
0.75	2.39375	1.40	0.99375
0.95	4.99875	2.80	2.19875

[PD] All 27 contingent policies are enumerated. Prior-optimal preparation is worth 5.2500; perfect information adds at most 5.6500. Payoffs include subsequent preparation and the conditional test; acquisition cost is additional. Decision utility is not historical cash and is never added to operating benefit.

AUTHORITY AND CUSTOMER PROTECTION REMAIN CONSTRAINTS

[PD] Resolve the actual role, transition, channel and cap version first. Test split, concurrent and replayed requests before field use. A favorable evidence value cannot authorize a known-invalid transition or substitute a denied-count target for mature physical outcomes.

[PD] State probabilities, signal quality and utilities are assumptions. Existing customer remedy remains available. No benefit is assigned to an unauthorized or silently approved action.

REPRODUCIBILITY / READER PROTOCOL

Rebuild the event, permission and outcome

[AI] The case distinguishes the written permission, its actual execution, the original service population and the reported metric. None can silently replace another.

Review layer	Reproduction requirement
Source and versions	Compare every selected populated cell with the immutable workbook. Preserve missing originals, later derivative rewrites and distinct program scope. Repeated summaries are not independent corroboration.
Aggregation proof	Recompute the positive-weight interval, rounding envelope and matched-exposure arithmetic. A failed necessary condition blocks that interpretation; it does not select a replacement historical total.
Exposure support	Enumerate the tested components and target versions. Rebuild the hypothetical opposite target effects compatible with identical narrow-pilot evidence. Precision cannot supply missing support.
Operating model	Rebuild original units, shared rule state, batch correlation, authorized fallback, mature outcomes and full costs. Keep fault remediation and approval requirements explicit.
Decision and information	Compare the economic threshold with detectability and decision power. Enumerate signal-contingent decisions, subtract evidence cost and preserve the perfect-information bound.

Public conclusions and private evidence

[AI] Public numerical constructions are hypothetical and independent of historical quantities. The private package retains identities, raw arithmetic, alternative versions, source paths, analytical code and review bindings. It is not intended for external portfolio distribution. A shared positive index preserves aggregation inequalities; separately rebased group indexes cannot recover aggregate levels.

Future evidence applies to the tested version

[AI] PERC and its prospective test are new analysis. A future result from strengthened event accounting and staffed fallback would describe that particular operating package and population. It would not retrospectively validate the historical pilot headline or transfer authorship of the original program.

PRACTICAL HANDOFF

[PD] Recover originals, reconcile the pilot extract, document active permissions and verify every authorized route. Accept the service and metric owners before calibration. Publish causal benefit only for a supported comparison with mature outcomes and a reconciled cost basis.

The private package contains executable calculations and source mappings. Source-semantic, independent numerical and actual-page reviews bind the released artifacts; substantive changes require renewed review.

RESEARCH REFERENCES / PERMISSION AND EXPOSURE

Verify the policy and its tested population

Primary research supports the method and proposed design. It does not authenticate project artifacts, supply local parameters or establish this project's causal effects.

Hu, V. C., Kuhn, R., & Yaga, D. (2017). Verification and Test Methods for Access Control Policies/Models. NIST SP 800-192. Official full PDF, selected sections. Distinguish the intended policy, formal model and implementation. PERC tests both excess authority and blocked valid work; this source does not certify historical control compliance.

Park, J., & Sandhu, R. (2004). The UCONABC Usage Control Model. ACM Transactions on Information and System Security, 7(1), 128–174. Author-hosted full PDF introduction and publisher metadata. Mutable conditions motivate cumulative quantity and continuing-authority checks. The case borrows the conceptual distinction, not a measured operational effect.

VanderWeele, T. J., & Hernán, M. A. (2013). Causal Inference Under Multiple Versions of Treatment. Journal of Causal Inference, 1(1), 1–20. Primary article, selected full-text sections. Define the permission and fallback version underlying each potential outcome. Coarsened exposure can be valid only under its identification conditions.

Tipton, E. (2014). How Generalizable Is Your Experiment? An Index for Comparing Experimental Samples and Populations. Journal of Educational and Behavioral Statistics, 39(6), 478–501. Institutional full PDF, introduction and support discussion. Within-pilot identification does not establish support for target settings; no generalizability index is estimated here.

Eldridge, S. M., et al. (2016). CONSORT 2010 Statement: Extension to Randomised Pilot and Feasibility Trials. BMJ, 355, i5239. Official publisher-indexed text; direct PDF unavailable. Prespecified feasibility and progression criteria inform the proposed test. The historical operational pilot is not represented as a compliant randomized experiment.

Sun, X., Briel, M., Walter, S. D., & Guyatt, G. H. (2010). Is a Subgroup Effect Believable? Updating Criteria to Evaluate the Credibility of Subgroup Analyses. BMJ, 340, c117. Official publisher-indexed text; direct PDF unavailable. Prespecify moderators and test interactions. Opposite descriptive subgroup movements do not identify different causal effects.

Hu, V. C., et al. (2014; updated 2019). Guide to Attribute Based Access Control (ABAC) Definition and Considerations. NIST SP 800-162. Official full PDF, selected sections. Attributes, policy decisions and enforcement are separate functional responsibilities. PERC applies that distinction to operational permissions without asserting historical ABAC adoption or mandatory applicability.

Access scopes distinguish selected full text, abstracts, publisher records and publisher-indexed text. Applications are conceptual or methodological; local results require local evidence.

RESEARCH REFERENCES / EVALUATION AND UNCERTAINTY

Measure the complete operating package

Primary research supports the method and proposed design. It does not authenticate project artifacts, supply local parameters or establish this project's causal effects.

NIST/SEMATECH. e-Handbook of Statistical Methods, section 6.3.2.2: Individuals Control Charts. Official method. Compute tentative limits for commensurate ordered points; preserve sparse windows, rounding and missing phases. A within-limit movement is not demonstrated improvement, and an invalid aggregation is not repaired by a control chart.

Laney, D. B. (2002). Improved Control Charts for Attributes. *Quality Engineering*, 14(4), 531–537. Primary publisher abstract. Dispersion adjustment needs actual eligible counts and denominators across rational subgroups. Do not recover fictitious counts from rounded DPMO or use cumulative cross-site shares as a repeated binary process.

Cohen, J. (1960). A Coefficient of Agreement for Nominal Scales. *Educational and Psychological Measurement*, 20, 37–46. Primary publisher record. Blinded coding should separate policy eligibility, executed action and physical outcome. Report prevalence, disagreement, kappa and uncertainty; agreement between coders does not authenticate their common evidence.

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. *International Journal of Epidemiology*, 34, 215–220. Primary abstract. Selection on an adverse result can produce reversion without intervention. High-use selection is not automatically selection on high outcome risk; quantify the relevant ranking relationship or retain explicit sensitivity assumptions.

Hudgens, M. G., & Halloran, M. E. (2008). Toward Causal Inference with Interference. *Journal of the American Statistical Association*, 103, 832–842. Original manuscript. Shared operators, approvers and overflow resources create interference across sites. Assign independent capacity domains and retain ordinary work. Repeated days within a domain are not additional independent clusters.

Callaway, B., & Sant'Anna, P. H. C. (2021). Difference-in-Differences with Multiple Time Periods. *Journal of Econometrics*, 225, 200–230. Author manuscript. A later observational staggered expansion needs defined treatment versions, group-time effects and appropriate comparisons. Parallel-trend diagnostics and sensitivity cannot repair missing exposure or outcome definitions.

Strong, M., Oakley, J. E., Brennan, A., & Breeze, P. (2015). Estimating the Expected Value of Sample Information Using the Probabilistic Sensitivity Analysis Sample. *Medical Decision Making*, 35, 570–583. Author repository. Information value depends on a feasible decision under explicit priors, likelihoods and payoffs. Deduct evidence cost and enforce the perfect-information bound. Do not add it to the operating benefit it may help select.

Access scopes distinguish selected full text, abstracts, publisher records and publisher-indexed text. Applications are conceptual or methodological; local results require local evidence.

RESEARCH REFERENCES / CONTRIBUTION AND GOVERNANCE

Preserve the connection to physical service

Primary research supports the method and proposed design. It does not authenticate project artifacts, supply local parameters or establish this project's causal effects.

Collier, D. (2011). Understanding Process Tracing. *PS: Political Science & Politics*, 44, 823–830. Primary publication. Distinguish managerial coverage, handover, intake authorship, program design and outcome ownership. A derivative record supports provisional process-tracing claims; involvement and proposed ownership do not establish delivery.

Tucker, A. L., & Edmondson, A. C. (2003). Why Hospitals Don't Learn from Failures. *California Management Review*, 45(2), 55–72. Author manuscript. Workarounds can sustain immediate work while preventing system learning. PERC retains the physical obligation after a denied virtual action and assigns its recovery. Neither intent nor displacement is inferred from status counts alone.

The W. Edwards Deming Institute. The Funnel Experiment. Official exposition. Reactive changes to routine variation can worsen a process. Do not repeatedly retune permissions from sparse weekly swings; first establish stable definitions, common phases and actionable causes.

Campbell, D. T. (1976). Assessing the Impact of Planned Social Change. Original institutional record. Accountability measures can change both behavior and recording. Preserve original events when request partitions or status categories change; process counts cannot replace the ultimate service outcome.

Goodhart, C. A. E. (1992). The Objectives for, and Conduct of, Monetary Policy in the 1990s. Reserve Bank of Australia Annual Conference. Original author paper, section 1(b)(i). The relationship between an intermediate target and an ultimate objective can change with adaptation. Submission reduction needs independent service and workload evidence before it becomes an operating benefit.

Ittner, C. D., Larcker, D. F., & Randall, T. (2003). Performance Implications of Strategic Performance Measurement in Financial Services Firms. *Accounting, Organizations and Society*, 28, 715–741. Author manuscript. Performance measurement needs a defensible connection to strategy and implementation. A permission specification, a pilot movement and a finance entitlement require explicit bridges before forming one value claim.

OECD / European Commission JRC (2008). Handbook on Constructing Composite Indicators. Official handbook. Weights, denominators and compensation assumptions govern aggregation. A common positive scale preserves convex bounds, whereas separately rebased groups cannot be pooled as levels. Retain a legitimate-service safeguard beside cost-weighted value.

Access scopes distinguish selected full text, abstracts, publisher records and publisher-indexed text. Applications are conceptual or methodological; local results require local evidence.

Accept the reconciliation; condition the expansion

[AI] This author-led, agent-assisted report is not external peer review or client approval. The verdict concerns the source-bounded analysis and proposed design, not historical causal-impact certification.

Major issue	Disposition in the revised case
Headline explained by subgroup offsets	The common-basis aggregate fails the convex-hull test. Require numerator, denominator, cohort and version bridges before accepting the decomposition.
Limited pilot generalized to full policy	Separate tested permissions from approval and central-service paths. Component exposure and target support precede causal transport.
Threshold movement called tighter control	Retain cap, minimum and capacity interpretations. Test the full predicate, cumulative event scope and live enforcement.
Restriction treated as recovered service	Retain all original obligations. Cost authorized fallback and ordinary work; require a separate legitimate-service safeguard.
Planning value credited as delivered impact	Keep source alternatives, tentative limits, selection and concurrent programs visible. Require mature causal contrast, full cost and overlap reconciliation.

Partner critique: can a COO act on page one?

[AI] Yes: reconcile the published pilot before using it to support expansion value, verify each permission path and give legitimate overflow an accepting service owner. Metric, platform, operations and Finance owners have separate acceptance tests. Recorded management work receives bounded credit.

Remaining conditions and minor issues

[AI] Original authenticity, live component exposure, outcome and event linkage, regional support and implementation costs require validation. The finite trial law, covariance and service margins are planning assumptions. Internal numerical and visual checks do not establish field performance, ranking or independent endorsement.

REVIEW DISPOSITION / PASS WITHIN STATED SCOPE

[AI] Accept the logical corrections and proposed evaluation within their stated assumptions. Historical causal benefit, realized cash and verified full expansion remain unestablished. A failed permission or service condition constrains the proposed rollout.

Internal source, method, numerical, executive and actual-page review. Acceptance covers the released analytical case and its limitations; no external referee or firm-ranking claim is made.