

OPERATING CASE STUDY / DATA, ANALYTICS & AI

Data Access & Investment Readiness

Turn source discovery into a governed reuse, access and investment decision

An anonymized operating-design case on an analytics program blocked by source fitness and access. The investigation identified a candidate incumbent model, corrected the data roadmap and specified an application-identity access plan.

PROJECT CONTEXT / ANALYTICS INVESTMENT & DATA GOVERNANCE

BLUF / BOTTOM LINE UP FRONT

Validate reuse before funding a rebuild.
Prove access for the intended workload.

ANALYTICAL FINDING

Derived scope finding: consolidating a proposed request did not consolidate all required authority. Ownership attribution remained in a separate data domain, while catalogue visibility established neither permission for the application nor fitness of the underlying records.

CONTRIBUTION IN THE SOURCE RECORD

Directed discovery, source rejection, reuse decisions and request scoping. Provisional register-level hoop support; data, model and approval ownership remained external.

EVIDENCE STATUS

Registry-based decision and design case. Request submission, workload access, reuse implementation and customer or financial results are not established.

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

EXECUTIVE ABSTRACT / DECISION AND EVIDENCE

Fund the evidence needed to choose reuse, repair or build

[AI] The recorded contribution is a change in investment logic: verify the available data product and its governed access path before committing to another model.

Structured abstract	Executive assessment
Context [HR]	An operating analytics application exposed generic category proxies where detailed attribution was needed. Available reporting channels did not supply the required daily unit panel.
Objective [HR]	Identify fit-for-purpose numerator, exposure and attribution sources; establish an application-identity request path; and determine whether a candidate model already addressed the planned prediction need.
Approach [HR]	Inspect catalogue metadata and permission records, reject incompatible sources, group candidate requests by provider, and draft a source-specific query and access plan.
Findings [AI/RC]	Provider membership is narrower than end-to-end authority. A separate ownership domain expands the dependency map. Weekly aggregates cannot uniquely reconstruct daily observations; an active model catalogue entry does not establish usable predictions.
Decision value [HR/AI]	The register records a switch toward consume-and-reconcile and corrections to the request specification. Avoided expenditure, elapsed time saved and completed reuse are unmeasured.
Implications [PD]	Release funding for bounded source and incumbent validation. Hold downstream attribution and production use until the Access-to-Decision Contract (ADC) is satisfied for the actual workload.

[PD] The sponsor owns the reuse-versus-build decision. The platform owner proves the deployed identity; data owners approve the declared resources and use; an independent reviewer accepts the grain, clock, population and reference reconciliation. Every mandatory condition must pass before operational use.

DECISION NOW

[PD] Approve a controlled diligence stage with explicit exit criteria. Discovery can justify a validation budget; it cannot justify a production-readiness claim or a booked benefit.

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


EXHIBIT 1 / INVESTMENT SEQUENCE

A discovered incumbent changes the next investment, not the final answer

[HR/AI] Catalogue discovery changed the roadmap from a parallel predictor toward consuming and reconciling a candidate existing score. The record does not establish its performance or successful integration.

DISCOVERY CHANGES THE NEXT INVESTMENT GATE

Historical roadmap decision / proposed validation sequence



An inventory entry earns a validation test, not a production-readiness label.
Retain the build option until substitution is demonstrated.

Source: S1 M07. Cadence, recent records and comparative performance were not verified.

Decision gate	Evidence required before commitment [PD]
Candidate exists	Record the product, version, intended use, accountable owner and declared output. An active lifecycle label is only an inventory state.
Candidate fits	Verify labels, population, horizon, point-in-time availability, calibration, subgroup performance and applicable use restrictions. A date partition does not prove refresh cadence.
Candidate adds value	Compare consume-only, calibrated reuse and a new build on an untouched, mature evaluation set, including integration and ongoing service cost.

[HR] The source metadata identified a probability field and a partition date. The detailed metric record classifies daily refresh as inferred; no table rows were read in the registered investigation. The stronger narrative claim that a usable daily product already shipped is therefore not adopted.

CAPITAL ALLOCATION RULE

[PD] Stop funding duplicate development while testing substitutability. If the candidate fails the required scope, timing or incremental-value test, retain a justified build option. The discovery alone does not prove that reuse is cheaper.

Source: S1 M07, narrative and reuse decision. Evidence class HR means register-attested. Research: Mitchell et al., Model Cards for Model Reporting; primary sources listed in the appendix.


EXHIBIT 2 / APPROVAL SCOPE

A compact request can still leave the decision’s authority incomplete

[AI] Grouping resources by provider reduces administrative fragmentation only when scope, identity and policy permit it. The later ownership-domain correction changes the dependency map.

GROUPED REQUESTS DO NOT COLLAPSE RESOURCE AUTHORITY

Proposed request scope / later ownership-domain correction



Source: S1 M01–M02 and provider correction. No submitted or approved grant is established.

Boundary in the source [HR]	Interpretation and action [AI/PD]
Panel and behavior resources	A common provider supports a grouped request proposal. It does not prove a common steward, classification, approval count or subscription entitlement.
Ownership-attribution resources	A distinct data domain supplies the authoritative ownership layer. Preserve it as a separate dependency; the panel request cannot be assumed to unlock it.
Request-set versions	The nominal compression comparison and the longer candidate inventory do not enumerate the same set. Freeze a resource manifest before counting administrative effort.

[AI] The correct objective is the least burdensome authorized path for a stated decision, subject to required coverage. Minimizing ticket count alone can hide missing resources, broaden sensitive access unnecessarily or couple a simple request to a slower approval branch.

NEW SCOPE RESULT

[AI] The consolidation thesis can hold for the panel while failing for end-to-end attribution. Report readiness by use case and dependency, then test whether bundling or staged requests reduce elapsed time under the actual policy.

Source: S1 M01–M02, D02 and later provider correction. Distinct providers concern distinct purposes; neither version is discarded. No approved request, grant or time saving is established.


EXHIBIT 3 / SOURCE FITNESS

Weekly totals leave the daily operating path unidentified

[RC/AI] A live weekly report can agree perfectly with an operating total while remaining unable to support daily decisions. This is an information constraint, not a connection failure.

THE SAME WEEKLY TOTAL CAN HIDE DIFFERENT DAILY PATHS

Synthetic counterexample / each profile totals 100; not internal observations



[RC] The illustrative profiles above each sum to an indexed weekly total of 100. One allocates the total evenly; the other concentrates it on one day. For seven unknown daily amounts, a weekly sum imposes one independent equation, leaving six allocation degrees of freedom before additional information or assumptions.

Test	Consequence [AI/PD]
Temporal resolution	Repeating or allocating a weekly total creates estimates, not observed daily history. Require a daily source or disclose and validate a reconstruction model.
Entity resolution	A service-partner identifier does not become an individual identifier through a renamed field. Preserve the supported decision level.
Clock and maturity	Creation-day concessions, shipping-day volume and dispatch-day losses use different event clocks. Date alignment alone does not create a mature unit-cohort risk denominator.

[AI] The source’s rows-per-week arithmetic describes output density, not temporal grain. The grain rejection rests on the reported grouping definition; it requires original query verification before reuse. Extending the extraction window does not repair missing resolution.

Source: S1 M04, M08 and source-rejection decisions. Daily profiles and algebra are synthetic analytical examples, not internal metrics. Weekly total is normalized to 100; no operating values are disclosed.

EXHIBIT 4 / IDENTITY AND AUTHORITY

Authentication, authorization and usable data require separate evidence

[HR/AI] The register records a human permission deficit, a browser-consent barrier and an absent intended application identity. None can be resolved by relabeling another layer as complete.

PROVE EACH LAYER IN THE CONTEXT THAT WILL ACTUALLY USE IT

Historical unresolved states / proposed acceptance evidence



Catalogue visibility is not row-level entitlement.
A human session does not establish an application’s effective permissions.

Source: S1 M01 and authentication decisions. Readability and decision fitness are separate.

Condition	Accountable evidence [PD]
Requester session	The requester clears required consent and supplies an accurate business purpose. A browser session does not confer rights on an application.
Workload identity	The platform owner creates and proves the actual deployed principal. A proposed role cannot be tested as though it already exists.
Resource entitlement	The resource owner or delegated authority decides the exact principal, resource, action and scope. Request submission and permission propagation are different events.
Accepted use	The consumer verifies records, allowed fields, currentness and declared use. An authorized read can still deliver data unfit for the operating question.

[AI] A metadata response does not establish row-level authorization. A failed sign-in does not show that entitlement exists. The recorded human denial and the untested workload state must remain separate; repeated authentication attempts cannot substitute for the missing policy decision.

Source: S1 M01, authentication decisions and unresolved deployment identity. NIST SP 800-207 distinguishes resource-specific authentication and authorization. No real consent, grant or access operation was performed in preparing this case.

EXHIBIT 5 / EVIDENCE RECONCILIATION

Resolve catalogue ambiguity before promoting it into an operating fact

[AI] The register is strong on the questions asked and weaker on acceptance evidence. Several precise-sounding statements require narrower interpretations.

OBSERVED METADATA AND VERIFIED OPERATING FACTS DIFFER

Historical source conflicts / analytical interpretation

SEARCH PAGE Returned set	SCHEMA Declared fields	ACCEPTANCE Tested eligible records
-----------------------------	---------------------------	---------------------------------------

A full page does not prove a complete inventory.
A dated partition does not prove fresh, covered or permitted predictions.

Source: S1 M05–M08. Preserve unknowns until a discriminating test resolves them.

Registered statement [HR]	Selected interpretation [AI/RC]
Returned inventory equals search cap	The returned set is a lower bound unless completeness or pagination is established. Its active status does not date creation or prove use.
Model has dated partitions	The catalogue suggests a candidate score. Actual cadence, recent coverage and leakage-free availability require record-level validation.
Two extraction-window versions	The longer design spans about 50% more lookback than the shorter. Neither proves execution; preserve both and rerun maturity checks before selecting a window.
Repeat-volume view or abuse signal	Frequency is not an abuse classification. Consume the authorized adjudication signal and confirm its population, meaning and completeness.

[PD] Maintain a fact ledger with the observed object, source version, test result, unknown condition and next resolver. Replace intuitive confidence percentages with reproducible acceptance evidence. A correct schema field still needs populated records and a permitted operating use.

Source: S1 M05–M08 and specification corrections. Search-cap interpretation is newly derived; catalogue completeness was not independently authenticated. Lookback comparison is relative and does not disclose raw windows.

EXHIBIT 6 / OPERATING CONTROL

The Access-to-Decision Contract makes each unresolved condition actionable

[PD] ADC is a proposed acceptance record linking a decision to its source product, required authority, actual workload and independently checked evidence.

BIND THE DECISION TO ITS SOURCE, AUTHORITY AND EVIDENCE

Proposed Access-to-Decision Contract / no compensating score



Every mandatory condition must pass for the specified workload and use.
Failure or uncertainty changes the allowed action and names the next resolver.

Proposed ADC. Ticket closure, source count and schema breadth are not readiness evidence.

Contract section	Required content and owner [PD]
Decision and substitute	Sponsor: intended action, required resolution, candidate incumbent, reuse test and criterion for reopening a build.
Source and meaning	Data steward: version, grain, population, event clock, maturity, eligible fields, lineage and reconciliation evidence.
Authority and execution	Resource and platform owners: principal, scope, allowed action, consent dependency, approved subscription and tested read.
Acceptance and accountability	Independent reviewer: valid readiness verdict, unresolved-condition owner, evidence expiry and permitted fallback or hold.

[PD] Version the contract whenever the required outcome, resource set, workload identity or source definition changes. Do not close a use case because an upstream ticket closed. A wider request can improve coverage while increasing approval dependencies; record both changes.

RELEASE THRESHOLD

[PD] Every mandatory condition must pass for the specific decision and workload. A known failure holds that use; an unresolved condition remains unresolved. Catalogue visibility and sensitive-field breadth cannot compensate for absent authority.

ADC is a proposed framework derived from the registered decisions and the new scope/identifiability analyses. It is neither a deployed product nor an externally certified standard.

EXHIBIT 7 / IMPLEMENTATION AND VALUE

Prove a better diligence process before claiming downstream operating value

[PD] The next test evaluates time to independently verified readiness, then values it under explicit resource assumptions. It preserves resource-owner authority and separates subsequent production outcomes.

EVALUATE VALID READINESS BEFORE CUSTOMER IMPACT

Proposed diligence experiment / separate subsequent field-value study



Retain every enrolled request and the original observation horizon.
Invalid readiness does not stop the clock; legitimate owner authority is preserved.

Planning assumptions are hypothetical. The model does not issue approvals or claim savings.

Stage	Decision and exit criterion [PD]
Freeze the use case	Recover request versions and evidence. Define population, protected fields, source clocks and a candidate-product test before assessing readiness.
Run controlled diligence	Compare the contract-led intake with usual intake across independent work/approval domains. Keep eligibility, authority and acceptance standards identical.
Verify complete decisions	Audit every enrolled request, including unresolved and denied cases. Invalid ready verdicts never stop the clock; a correct hold preserves governance while contributing unready time.
Assess economics	Compare diligence-resource effects with its incremental cost. Evaluate later data delivery and mature operating effects separately; do not book avoided build spend without a counterfactual budget.


[AI] The portfolio contribution is the disciplined sequence: identify a substitute, reject wrong-grain sources, scope the required authority and make unresolved work explicit. Its success should be measured by valid decisions and lower rework, not by an inflated count of sources or closed requests.

Proposed implementation sequence. Historical request filing, approval, integration, operating outcome and realized cash benefit remain unestablished. The subsequent model is explicitly hypothetical planning analysis.

TECHNICAL ANALYSIS

Discovery changes the plan before it changes the outcome

[AI/PD] The proposed Access-to-Decision Contract (ADC) binds the required decision, source fitness, accountable authorization, intended workload and acceptance evidence before committing to reuse or build.



[PD] Z assigns early diligence; R is the source/request plan; V is independently verified readiness; L is restricted unready time. D is the investment or operating decision, A qualified delivery and use, and Y a mature customer outcome. H/I combines resource policy, consent and the intended workload identity. Q contains source grain, clock, scope and capture; O is the recorded outcome.

[PD] C names request complexity, shared staff and approval capacity; U represents unrecorded policy changes and concurrent programmes. The proposed benchmark would measure readiness and L under scripted authority conditions. It would not execute customer-facing actions or identify the A→Y effect.

[AI] Diligence can change reuse, escalation or investment directly, so Z is not automatically an instrument for data use. Q can change O without changing Y. Residual delays and outcome differences remain compatible with unmeasured approval, measurement and operating paths; the graph does not identify which path caused them.

Basis: registered discovery and request design [HR]; proposed ADC graph and benchmark [PD]. Access, deployment and mature customer effects remain distinct.

TECHNICAL ANALYSIS

A common provider is not a common authorization decision

[HR/AI] The registry supports a consolidated request proposal. It does not establish a filed request, accepted scope, shared approver set or successful read by the intended application.

Layer	Evidence required for the stated conclusion [PD]
Identity and session	Identify the actual subject and process; distinguish a human session from the application workload.
Resource authorization	Confirm the resource, permitted action, scope, conditions and approving authority for that subject.
Data path	Verify subscription, query execution and returned content under the intended deployment configuration.
Accepted use	Confirm required population, grain, clock, fields and refresh quality before exposing a decision.

[HR/AI] The registered human principal lacked effective permission. The intended workload identity did not yet exist, so its grantability was untested. Creating that identity would make a permission test possible; creation would not itself grant access. Successful catalogue reads and a blocked questionnaire also do not prove that all remaining barriers are authentication-only.

[AI] Provider membership can reduce coordination work, but authorization may still differ by dataset, field, purpose or steward. Aggregate grain and fewer columns do not establish an accepted privacy classification. The later ownership-attribution provider correction expands the source-to-owner map; it cannot be absorbed into the earlier single-provider promise.

[PD] Treat browser consent, requester attestation and resource-owner approval as separate decisions. A human bridge is conditional on actual entitlement and subscription. Give any bridge an accepted operator, expiry and failure route; do not convert an interactive session into a claim of durable unattended permission.

[AI] NIST distinguishes authentication from authorization, while OAuth separates client, owner, authorization server and resource server. These are analytical role distinctions here, not proof of the internal platform’s implementation or a security-compliance certificate.

Basis: registered principal, request and provider conflicts [HR]; NIST SP800-162; RFC6749 and RFC9700.

TECHNICAL ANALYSIS

Weekly totals cannot identify the missing daily decisions

[AI] Source fitness depends on retained information. A live query channel does not restore daily detail or individual attribution that its output has already aggregated away.

One weekly total leaves multiple daily explanations

[AI] Let W be the sum of seven daily amounts. For a positive total away from boundary constraints, the daily vector has six allocation degrees of freedom. Moving any feasible amount from one day to another preserves W . Daily interpolation therefore adds assumptions; it is not recovered observation. An aggregate lacking individual identifiers likewise cannot establish individual accountability.

[HR/AI] The register describes the available reporting extract as weekly and the proposed input as daily. Its rows-per-week approximation fits only one business branch. Row-count arithmetic alone does not establish keys or grain: inspect schema, grouping logic and unique-key multiplicity against actual rows.

Requirement	Acceptance test [PD]
Daily panel	Verify station-day keys, timezone, complete eligible dates, duplicate handling and the meaning of zero.
Comparable rate	Reconcile numerator population, exposure unit, clock, maturity and join multiplicity.
Individual attribution	Require an authorized identifier, ownership definition and event linkage; provider-level identity is insufficient.

[HR/AI] Recognition day, shipment day and dispatch day need not describe the same customer cohort. Their calendar-day ratio may be a valid accounting statistic but is not automatically a mature failure risk. The source leaves an exposure-column name and business-line crosswalk unresolved.

[RC/PD] The proposed SQL lookback is about 1.5 times the narrative lookback. Freeze both versions, then compare common-window coverage, join retention and maturity before choosing. Longer history cannot cure incompatible clocks. For causal evaluation, freeze eligible units upstream and count prespecified outcomes through a declared horizon from the original eligibility date.

Basis: registered schema, extraction-window and grain conflicts [HR/RC]; non-injectivity proof [AI]; Gray et al. (1995); Wang and Strong (1996).

TECHNICAL ANALYSIS

Discovering an incumbent creates a testable reuse option

[HR/AI] The catalogue identified an existing model output and a candidate ownership route. That justified validation before duplicate investment; it did not establish that reuse was already fit or cheaper.

Discovery evidence	What remains unverified [HR/AI]
Probability field and declared purpose	Actual records, target definition, calibration, discrimination and coverage for this decision.
Snapshot partition and related tables	Refresh cadence and publication lag; the source labelled the cadence an estimate.
Lifecycle marked active	Current production use, accepted service obligations, owner support and deployment fitness.
Catalogue page at its return limit	Full inventory size; the result is a returned subset until pagination or an exact count is verified.

[PD] Register reuse, extend, build and defer as alternatives. Compare them on the same task, eligible population, as-of information and acceptance loss. First establish permission and sample records; then evaluate a frozen validation set. Keep model prediction quality separate from the causal effect of actions taken from its scores.

[PD] The reuse contract includes target semantics, missing-data behavior, permitted purposes, refresh commitments, operating ownership, change notices, rollback and exit conditions. A named probability field is not a probability guarantee, and a missing literal field name does not establish absence of every possible equivalent identifier.

[AI] A staged investigation can preserve the option to avoid an unsuitable duplicate build. Its value depends on the possible decisions, evidence quality, investment reversibility and cost of delay. The case contains no authenticated cancelled budget or engineering counterfactual, so avoided-build savings remain unestimated.

[AI] The provider correction also changes who must participate in validation. Consumption of an incumbent score does not transfer ownership of its training process, its attribution rules or customer-abuse adjudication. A consumer may reject an unfit input while leaving those decision rights with the responsible owners.

Basis: registered incumbent schema and estimated cadence [HR]; NIST AI RMF1.0; McGrath (1997), option timing concept; no historical option valuation.

TECHNICAL ANALYSIS

Measure readiness against the original requirement

[PD] Discovery, approved access, usable data and decision use have different denominators. Count each against a frozen request-family population and preserve the reason a requirement remains unresolved.

Construct	Validity and invariance contract [PD]
Source fit	Required population, grain, clock and fields are independently tested; a schema description alone is provisional.
Workload permission	The intended subject can perform the required action on the named resource under accepted conditions.
Verified readiness	Permission, source fitness and the declared integration check all pass; ticket closure is not the event.
Coverage and use	Expected request families remain in the denominator, including denial, withdrawal, unresolved and unknown states.

[PD] Bind each judgment to requirement version, principal, provider, table version, approved purpose, evidence vintage and accepted owner. Preserve raw responses and independently verified row samples privately. Provenance records how a claim was formed; it does not authenticate that claim by itself.

[PD] Audit future request-context cases with two independent blinded coders using an owner-approved reference dossier. Retain their labels before adjudication. A random core of 200 contexts gives an iid worst-case agreement margin near ± 6.9 percentage points. Shared providers and reviewers require cluster-aware sizing; this is neither a historical sample nor a kappa precision claim.

[PD] Report confusion matrices, category prevalence, raw agreement and $\kappa = (p_o - p_e) / (1 - p_e)$, where p_o is observed agreement and p_e chance-expected agreement. Pilot the prevalence-specific interval precision. Oversample rare exceptions only with known sampling weights for population summaries.

[AI] Access-restricted documentation can be selectively missing for the most complex resources. Available metadata are therefore not a random validation sample. Report missingness by provider, purpose and request class; use worst-case bounds for unresolved acceptance rather than silently dropping it. Agreement between coders using the same wrong dossier does not prove validity.

Basis: proposed measurement audit [PD]; Wang and Strong (1996); Cohen (1960); W3C PROV-DM (2013).

TECHNICAL ANALYSIS

A catalogue snapshot is not a process improvement series

[HR/AI] The supplied record contains no eligible historical operating time series. Snapshot states and planned reductions cannot establish a change in delivered access or customer performance.

Evidence structure	Permitted interpretation [AI]
Permission, proxy and inventory snapshots	One observation per construct; no adjacent moving range or time trend is identified.
Planned request/provider reductions	Mechanical design-shape checks only; not achieved approval or delivery observations.
Totals covering several weeks	No weekly sequence without its component observations; do not invent evenly allocated values.
Schema breadth and alternative windows	Heterogeneous inventory or version scenarios, not repeated process measurements.

[PD] For a future commensurate ordered series, define $MR_t = |x_t - x_{(t-1)}|$ and tentative individual limits $\bar{x} \pm 2.66 \cdot MR$; the moving-range upper limit uses $3.268 \cdot MR$. Freeze process phase, spacing and rules. Short, autocorrelated or changing-definition sequences need explicit suitability limits; an SPC signal would still not identify intervention causality.

[AI] No repeated binary count/exposure panel supports Laney p' here. A planned decline or a zero permission snapshot is not a sequence of independent trials. Unknown follow-up cannot be carried forward as observed continued failure.

[PD] No fixed worst-delay or worst-N treatment cohort is observed. If future enrollment selects the worst tenth of normally distributed noisy delay scores, its mean is 1.755 standard deviations above the population. Under assumed repeat-score correlations of 0.3, 0.6 and 0.9, untreated reversion would be 1.23, 0.70 and 0.18 standard deviations. These are illustrative sensitivities, not adjustments to this case.

[PD] Stratify future evaluation by business line, node type, sensitivity class and shared approval domain. Early diagnosis should matter most when source-fit work is on the critical path. Where permission remains binding throughout follow-up, faster diagnosis need not change readiness by the deadline.

Basis: private evidence inventory and mechanical screens [HR/RC]; proposed normal-tail sensitivity [PD]; NIST XmR; Barnett et al. (2005).

TECHNICAL ANALYSIS

The contribution is access diligence with bounded ownership

[HR/AI] The registered contribution concerns scoping, source rejection, reuse direction and request design. Approval, live consumption and operating impact remain separate claims.

Personal contribution claim [HR]	Process-tracing assessment [AI]
Directed source discovery and requirements	Provisional hoop: the register describes direction and delegated agent execution; originals are not authenticated.
Chose incumbent validation and corrected source mapping	Provisional hoop: decisions are registered, but independent decision correspondence and adoption traces are absent.
Prepared request and daily extract design	Provisional hoop: artifact pointers support the claimed design role; they do not establish submission or execution.
Owned approvals, incumbent model or realized impact	Unestablished: these belong to other owners or lack outcome evidence. No smoking-gun or doubly decisive test passes.

Consolidated limits and their decision consequences

[AI] Provenance: source pointers remain registry-level evidence; original logs, requests and datasets were not recovered. Repeated summaries are not independent corroboration. The absence of a recovered completion trace establishes neither failure nor success outside the supplied corpus; current status is unknown.

[AI] Identification: there is no untreated comparison, executed access intervention panel or mature customer cohort. Schema availability cannot identify a permission effect, adoption effect or field benefit. The sandbox is a proposed test of diligence timing, not a reconstruction of history.

[AI] Measurement and transport: capped discovery, inaccessible documentation, different clocks and unresolved joins constrain fit claims. The prospective scripted permission law and case mix do not estimate the live approval environment. Genuine revocation or reopening would require a revised readiness endpoint.

[AI] Economics and privacy: no cancelled-build counterfactual, operating-cost panel or cash reconciliation is supplied. Hypothetical information value is not additive realized savings. Public conclusions retain anonymized sources and protected quantities; the private annex preserves exact locators, forks, formulas and review checks.

Basis: registered role, delegated discovery and artifact pointers [HR]; process-tracing grades and consolidated limitations [AI]; Collier (2011).

TECHNICAL ANALYSIS

Test earlier diagnosis without giving it approval authority

[PD] Randomize independent shared-work and approval domains to early ADC diligence or usual intake. Freeze request families and the original intake clock before assignment; test readiness without granting real access.

[PD] The proposed sandbox follows six request families per domain for eight weeks. Permission timing is scripted under the same law in both arms; source-fit work differs by assignment. No actual grant, data release or customer action occurs. Within-domain requests share a latent approval environment, so the domain—not each request—is the independent unit.

[PD] Let $T(z)$ be verified readiness time under assignment z . Primary benefit is $\Delta = E[\min(T(0), 8)] - E[\min(T(1), 8)]$, averaging all original families within domains. Positive Δ means less restricted unready time. A correct permission hold still counts as unready; that state is not labelled improper delay. Denials, withdrawals, merges and administrative closure cannot remove original families.

[PD] Independent weekly verification applies the same frozen acceptance rule in both arms. Readiness requires permission plus source fit and the specified integration check. False-ready declarations, scope omission and authority bypass are separate guardrails. Their margins and joint testing must be sized independently; a positive time contrast cannot authorize unsafe release.

[PD] The finite planning model tests the zero-effect and economic-boundary laws, incorporates shared-domain variation and compares minimum detectable benefit with same-horizon break-even. Its cluster requirements are conditional, not recruited capacity or achieved field power. A cost for unready time represents an assumed resource burden; it is not concession loss or automatic cash.

[AI/PD] Parallel assignment preserves a concurrent control and avoids protocol learning during crossover. It does not solve shared-staff spillovers: isolate work domains or redesign and repower. Field approval timing may itself respond to better scoping and must be measured. A later live evaluation needs authorized rollout, stable eligibility, maturity and independent customer outcomes; this benchmark does not power that trial.

Basis: proposed domain-randomized benchmark [PD]; restricted mean-time estimand adapted from Royston and Parmar (2013); detailed synthetic law and economics on model pages.

TECHNICAL ANALYSIS

Four propositions turn access diligence into a research programme

[AI/PD] ADC extends source discovery into a testable decision contract. The historical record motivates its mechanisms; it does not establish their average causal effects.

Proposition and mechanism [PD]	Disconfirming evidence and boundary [PD]
P1: earlier fit diagnosis reduces restricted unready time when fit is on the critical path.	No benefit beyond the prespecified investment threshold; fixed permission may leave completion unchanged.
P2: requirement-to-owner mapping reduces scope omissions across provider boundaries.	Lower request counts accompanied by worse valid coverage contradict the claimed consolidation benefit.
P3: incumbent-first validation reduces duplicate commitment when the incumbent meets the same requirements.	Independent task acceptance rejects the reused input, or total commitment rises without a justified gain.
P4: explicit grain and clock contracts reduce false-fit source acceptance.	Held-out aggregate/cohort counterexamples are still certified as suitable at the same or higher rate.

[AI] Wang and Strong distinguish accessibility from contextual fitness; this case illustrates why one cannot substitute for the other. The aggregation proof confirms an information constraint, not a local causal effect. NIST and OAuth role separation explain why consent, identity and resource permission require different evidence.

[AI] Tucker and Edmondson distinguish workarounds from system-level learning. The proposed expiring bridge and owner route apply that mechanism without importing hospital results. McGrath’s option logic supports staged investigation before irreversible commitment; the project’s option value remains unpriced.

[AI] Campbell’s measurement argument and Ittner and colleagues’ performance-measurement findings caution against treating fewer requests as value. ADC uses noncompensable fitness and authority gates: a weighted score cannot offset a failed mandatory permission. The testable contribution is the connection from verified prerequisites to a bounded investment decision.

[PD] Preregister effect thresholds and guardrail margins. A lower confidence bound above break-even supports the modeled investment gate; an upper bound below it contradicts that benefit claim. An interval spanning the threshold is inconclusive, not falsification. Archive versions, code, assignments and adjudication evidence for independent reproduction.

Basis: proposed ADC propositions [PD]; Wang and Strong; Gray et al.; NIST; Tucker and Edmondson; McGrath; Campbell; Ittner et al. Literature supports mechanisms, not historical efficacy.

PROSPECTIVE MODEL / VERIFIED READINESS

Measure time to verified readiness, not time to close a ticket

[PD] Every clock, probability, cost and sample size is hypothetical. The controlled experiment holds permission timing fixed to isolate early source and incumbent-model diligence.

[PD] Randomize independent work/approval domains to early diligence or usual intake. Freeze six eligible request families per domain at original intake; follow all through eight weeks, including denial, withdrawal, merge, wrong-fit discovery and hold. An independently verified readiness event requires both authority and source fitness. A closed ticket or unsupported approval does not stop the clock.

Readiness input	Usual intake	Early diligence
Source-fit time S	Weeks 2 / 4 / 7; probabilities 0.20 / 0.40 / 0.40	Weeks 1 / 3 / 5; probabilities 0.40 / 0.40 / 0.20
Permission time P	Week 2 / 5 / after horizon; probabilities 0.25 / 0.35 / 0.40	Same scripted law; no grant is accelerated
Independent verification	One week after both prerequisites	Same one-week verification

[PD] Verified readiness is $T = \max(P, S) + 1$. The primary endpoint is restricted unready time $\min(T, 8)$, from original intake. Legitimate authorization waiting remains in this endpoint; only the randomized contrast is proposed benefit. Truth is independently adjudicated, not inferred from self-reported readiness.

DISTINCT OUTCOMES

[RC/PD] Mean unready time falls from 7.03 to 6.30 weeks in the planning alternative, while readiness by week 8 remains 60%. Faster verified progress need not change the horizon completion share.

[PD/RC] Scripted eight-week readiness experiment; no actual request, grant or production read is performed. Withdrawal and missing verification never become ready outcomes.

PROSPECTIVE MODEL / DOMAIN DEPENDENCE

Permission constraints cap the benefit of faster source discovery

[PD] Six requests share a domain environment. Treating them as independent domains would overstate precision.

[PD] Domain types are equally likely. In a faster-permission type, week-2 probability rises by 0.10 and after-horizon probability falls by 0.10; the other type reverses both shifts. Week-5 probability stays 0.35. Source-fit and permission times are conditionally independent within type. Requests are conditionally independent; domains are independent.

[PD] Lambda mixes the usual and early source-fit distributions; it changes no permission probability. Compose each domain’s six-request distribution before averaging types. Compare independent arm sums. This is parallel randomization, not a paired test. The finite law preserves shared-domain dependence and recalculates variance along the effect path.

[RC/PD] Source-fit time under early diligence is stochastically no slower. Applying $\max(P, S) + 1$ and truncation preserves that order. Therefore greater lambda lowers unready time, and the break-even boundary is least favorable for smaller benefits within this known family.

Threat	Required control before a field trial
Permission interactions	Real request scoping may alter approval timing or quality. Estimate this joint pathway rather than imposing the scripted exclusion.
Shared queues or owners	Cluster common approval/support dependencies; nominal request count does not establish independent sample size.
False readiness	Blinded authority, grain, source and accepted-read checks; predeclare error margins and separate guardrail power.
Revocation or reopening	The model assumes readiness persists within the horizon. Field recurrence requires an explicit longitudinal endpoint.

SCOPE OF IDENTIFICATION

[PD] An untouched bank and independent adjudication can test the proposed mechanism. Catalog discovery, approval authority and usable delivery remain distinct. This model estimates no historical access or customer effect.

[PD] Hypothetical conditional dependence structure and restricted effect family. Nuisance uncertainty and provider/source complexity require blinded pilot calibration.

PROSPECTIVE MODEL / POWER AND RESERVATION VALUE

A readiness test must resolve the resource break-even threshold

[PD] Require 80% power against zero at break-even (two-sided alpha = 0.05) and 80% probability that the exact lower bound clears break-even at the full alternative.

[RC/PD] Normalize usual unready time to 100; full improvement is 10.38 points. If resource reservation per unready week is c, net value = time reduction × c × avoidable share a × non-overlap h minus added cost. With a = h = 1 and cost 5 per baseline reservation pool, break-even is 5.

Domains / arm; total	80% MDE	Economic power	Both gates
25 / 50	7.19	51.3%	Fail
51 / 102	4.99	80.1%	Conditional
64 / 128	4.45	88.0%	Conditional

[RC/PD] Minimum: 51 per arm. The 64-per-arm buffer fails doubled heterogeneity; 80 per arm clears that stress (83.4% economic power). With a = 0.75 and h = 0.80, break-even is 8.33 and 363 per arm are required; feasibility is unestablished. With a = 0.50 and h = 0.80, no sample expansion fixes unfavorable economics.

Full-effect net / a	h = 1.00	h = 0.90	h = 0.80
0.50	0.19	-0.33	-0.85
0.75	2.79	2.01	1.23
1.00	5.38	4.35	3.31

[PD] Exact monotone upper-tail inversion gives a one-sided 97.5% lower bound; it must exceed the applicable break-even. Known-law false clearance is at most 2.5%. Cost 5 covers added diligence 2, fit validation 1.5 and coordination 1.5; common research adjudication is separate. Reservation value is assumed, not saved cash.

[PD] Separate field value: eligible volume × mature customer ITT reduction × avoidable cost/outcome × non-overlap share minus incremental cost. Fix population and maturity. Inputs are unavailable; field impact is unestimated and unpowered.

[PD/RC] Same fixed service pool; sample size does not redefine it. Scenarios are not confidence intervals. Recalibrate unknown nuisance laws and unweighted safety guardrails.

PROSPECTIVE MODEL / REUSE INFORMATION VALUE

Inspect the incumbent before committing to repair or replacement

[PD] This separate information model values technical-fit evidence. Permission and accepted-delivery gates remain mandatory regardless of which architecture is preferred.

[PD] The incumbent either fits as-is, fits after defined repair, or requires new construction. Priors are 0.45 / 0.35 / 0.20. Choices prepare reuse, prepare repair or prepare a new build. These are planning decisions, not approvals to access data or serve an unvalidated model.

Planning action / true state	Fits	Repairable	New needed
Prepare reuse	20	-14	-25
Prepare repair	9	16	-12
Prepare new construction	-18	-10	18

[RC/PD] Prior-optimal repair planning is worth 7.25; perfect information adds at most 10.95. A three-label signal is correct with probability 0.45, 0.70 or 0.90; wrong labels split the remaining probability equally. Optimize the action after each signal.

Signal	EVSI	Evidence cost	Net EVSI
Low	0.0000	0.50	-0.5000
Base	2.5125	1.50	1.0125
High	8.1375	3.00	5.1375

[PD] Payoffs are hypothetical net planning utilities with action costs included. Evidence-acquisition costs are separate. None relies on the readiness-time effect, and information value is not added to operational benefit. Real evaluation must test task fit, calibration, integration and permissions before reuse.

DECISION BOUNDARY

[PD] Discovering an incumbent can change the proposed build budget without establishing its fitness or granting access. A positive evidence value supports diligence; it cannot certify the asset or authorize its use.

[PD/RC] Finite-state Bayesian decision model; no historical priors or avoided-build savings are claimed. EVSI is bounded by EVPI in the same utility model.

Reproduce the investment logic without exposing the source estate

[AI] This document is a decision and operating-design case. It records a proposed path to source fitness and authority; it does not establish successful access, a deployed substitute or a historical operating effect.

Claim class	Meaning and boundary
[HR] Historical record	A registered requirement, artifact, reported observation or decision. Original drafts, SQL and access responses have not been recovered and authenticated.
[RC] Recalculated	Source-fork comparisons, the synthetic aggregation counterexample and explicitly hypothetical planning calculations. Arithmetic does not authenticate the record.
[AI] Analytical inference	A bounded interpretation of scope, grain, identity, permission, evidence or attribution. An explanation is not a proven mechanism.
[PD] Proposed design	ADC, prospective tests, measurement contracts, release gates and value scenarios. These are proposed methods, not historical deployments.

Reproduction sequence

[PD] Verify the workbook hash and exact extraction. Reconcile the source set, query limit, horizon alternatives, proxy numerator and chronology. Recreate the normalized daily counterexample and the numerical planning law. Independently recompute sample size and information value. Rebuild both PDFs, inspect every page and bind the checks to the final files.

[PD] The private package contains source identities, cell locators, raw values, unresolved source versions, calculation code, research notes and review dispositions. The original workbook and unrecovered operational artifacts are not embedded. No production access is required to reproduce the planning model.

WHAT THE NEW ANALYSIS CHANGES

[AI] Consolidation must be evaluated against a fixed use-case scope. Catalogue discovery is not readiness, and weekly agreement cannot identify the daily path. These constraints change which evidence the sponsor should buy before funding integration or another model.

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

RESEARCH REFERENCES / METHODS AND INTERPRETATION

Tie each request to a subject, resource and intended use

Primary sources inform the proposed contract, measurement and tests. They do not independently corroborate the historical project.

Hu, V. C., et al. (2014; updated 2019). Guide to Attribute Based Access Control (ABAC) Definition and Considerations. NIST SP 800-162. Subject, object, action and environment distinguish request scopes. This case does not assert that the internal platform implements ABAC.

Rose, S., Borchert, O., Mitchell, S., & Connelly, S. (2020). Zero Trust Architecture. NIST SP 800-207. Authentication and authorization are distinct, resource-specific functions. A successful session is not general application entitlement.

Hardt, D., ed. (2012). The OAuth 2.0 Authorization Framework. RFC 6749. Roles, authorization, tokens and protected-resource access are different stages. Error strings in a derivative register do not diagnose every underlying failure.

Lodderstedt, T., Bradley, J., Labunets, A., & Fett, D. (2025). Best Current Practice for OAuth 2.0 Security. RFC 9700, BCP 240. Current security guidance informs scoped token and audience reasoning. It is not evidence of a historical grant.

Gray, J., Bosworth, A., Layman, A., & Pirahesh, H. (1995). Data Cube: A Relational Aggregation Operator Generalizing Group-By, Cross-Tab, and Sub-Totals. MSR-TR-95-22, revised October 18. Aggregation context; the daily-allocation non-identifiability proof is derived in this case.

Wang, R. Y., & Strong, D. M. (1996). Beyond Accuracy: What Data Quality Means to Data Consumers. JMIS, 12(4), 5–33. Accessibility and contextual fitness are distinct. A usable source must fit the operating decision as well as be readable.

Mitchell, M., et al. (2019). Model Cards for Model Reporting. FAT* '19. Original Google Research publication record. Intended use, evaluation conditions and performance reporting inform incumbent validation; a named score field supplies none of these by itself.

NIST (2023). Artificial Intelligence Risk Management Framework, AI RMF 1.0. NIST AI 100-1. Task scope and third-party component assessment inform adoption gates. The framework is voluntary; no compliance certification is claimed.

Project evidence: protected register S1. Research access scope and prior-series methods are documented privately. Literature supplies concepts and methods, not project certification.

RESEARCH REFERENCES / METHODS AND INTERPRETATION

Define the population, clock and evidential test before measuring success

Primary sources inform the proposed contract, measurement and tests. They do not independently corroborate the historical project.

Royston, P., & Parmar, M. K. B. (2013). Restricted Mean Survival Time: An Alternative to the Hazard Ratio for the Design and Analysis of Randomized Trials with a Time-to-Event Outcome. *BMC MRM*, 13, 152. The case adapts the restricted-time summary, using a separately derived finite planning law.

NIST/SEMATECH. e-Handbook of Statistical Methods: Individuals Control Charts. Moving-range estimation requires ordered observations and a suitable process phase. Planned counts, single snapshots and schema versions are not an operating time series.

Cohen, J. (1960). A Coefficient of Agreement for Nominal Scales. *Educational and Psychological Measurement*, 20, 37–46. Publisher metadata. Kappa and disagreement tables support a proposed blinded readiness audit; no observed reliability is claimed.

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. Prior-series primary abstract. Worst-delay recruitment may recover without intervention; the selection sensitivity is hypothetical.

Collier, D. (2011). Understanding Process Tracing. *PS: Political Science & Politics*, 44, 823–830. Primary publication. Diagnostic evidence grades constrain personal-contribution claims. Direction of agent work does not establish personal execution or control of approvals.

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 comparisons are relevant if a later field rollout is observational; they are not the randomized benchmark estimator.

Moreau, L., & Missier, P., eds. (2013). PROV-DM: The PROV Data Model. W3C Recommendation. Versioned entities, activities, agents and derivations inform the evidence ledger. Lineage represents a claim; it does not establish correctness or authority.

Project evidence: protected register S1. Research access scope and prior-series methods are documented privately. Literature supplies concepts and methods, not project certification.

RESEARCH REFERENCES / METHODS AND INTERPRETATION

Treat access diligence as an investment decision with explicit limits

Primary sources inform the proposed contract, measurement and tests. They do not independently corroborate 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 working paper. Workarounds and system-level correction differ. The case proposes owner escalation without transporting hospital effect sizes.

McGrath, R. G. (1997). A Real Options Logic for Initiating Technology Positioning Investments. *AMR*, 22(4), 974–996. Publisher abstract. Uncertainty and boundary conditions affect staged technology investment. A discovery can preserve a build option without proving a historical saving.

Campbell, D. T. (1976). Assessing the Impact of Planned Social Change. Occasional Paper 8, Dartmouth College; original reprint. Consequential indicators can distort behavior. Request closure and source count must not substitute for independently valid readiness; intentional gaming is not established.

Ittner, C. D., Larcker, D. F., & Randall, T. (2003). Performance Implications of Strategic Performance Measurement in Financial Services Firms. *AOS*, 28, 715–741. Primary abstract. Measurement architecture and economic performance require distinct evidence; a contract does not prove benefit.

OECD / European Commission JRC (2008). Handbook on Constructing Composite Indicators: Methodology and User Guide. Official handbook. Weights and normalization need sensitivity analysis. ADC uses mandatory gates rather than allowing strengths to offset absent authorization or unusable grain.

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. The case uses a separate finite-state decision calculation, not the paper's regression estimator.

The W. Edwards Deming Institute. The Funnel Experiment. Official explanation. Reacting to routine variation can worsen a stable process. No stable historical process or tampering effect is demonstrated here; proposed monitoring preserves that distinction.

Project evidence: protected register S1. Research access scope and prior-series methods are documented privately. Literature supplies concepts and methods, not project certification.

FINAL REFEREE REPORT / INTERNAL AGENT-ASSISTED REVIEW

Approve the diligence case within its evidence boundary

[AI] This report records author-led review with parallel evidence, methods and numerical agents. It is not journal peer review, client approval or a firm-ranking certification.

Major issue	Disposition in the revised case
Scope and authority	The candidate bundle, longer inventory and separate ownership domain are retained as different scopes. Provider grouping does not prove approval, access or elapsed-time reduction.
Source and incumbent fitness	Cadence remains inferred; rows were not queried. Weekly aggregation is formally non-injective. Event-day ratios do not become mature unit-cohort risks through date alignment.
Evidence and contribution	Original artifacts remain unrecovered. Directed discovery and specification decisions receive provisional register-level support; data, model, adjudication and approvals remain externally owned.
Identification and value	The historical effect is not identifiable. The proposed restricted-time experiment has a declared synthetic law. Waiting burden, avoidable resource value, decision information and customer outcomes remain separate.
Arithmetic and review	Source conflicts and displayed calculations are reproducible. Independent model checks and actual rendered-page review are bound to final files in the private package.

Partner critique: would a COO act on page one?

[AI] Yes: authorize a bounded incumbent-and-source validation stage, assign the workload and data-authority dependencies, and keep production use on hold until all mandatory conditions pass. The cover does not sell a catalogue discovery as a completed integration.

Minor revisions and limitations

[AI] Review narrowed catalogue-cardinality, model-cadence and approval-compression statements; corrected proxy and event-clock interpretations; and separated legitimate permission waiting from avoidable cost. The technical limitations govern transport to real operations.

REVIEW DISPOSITION — PASS WITHIN STATED SCOPE

[AI] The case supports a documented decision process, new identifiability analysis and a reproducible proposed test. It establishes neither completed workload access nor historical savings or customer impact.

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