

OPERATING CASE STUDY / STRATEGY & CAPITAL ALLOCATION

Workforce Capacity & Headcount Planning

A smaller allocation does not make the work smaller

An operating-project case on headcount and capacity planning. It examines how recurring work, project effort, availability and release sequence shape a staffing request when the funded envelope is fixed.

WORKLOAD, CAPACITY & DELIVERY COMMITMENTS

BLUF / BOTTOM LINE UP FRONT

Keep the work estimate intact. Match the delivery promise to available capacity.

ANALYTICAL FINDING

The reported staffing gap depends on unresolved units and denominators. Matching the funded total establishes allocation compliance; it does not establish productivity, completion dates or executable coverage.

CONTRIBUTION IN THE SOURCE RECORD

Registry-level hoop: reported model replication, allocation reconciliation and escalation. The planning lead did not own the sanctioned method or funding decisions.

EVIDENCE STATUS

Evidence-reconciliation and proposed-design case. Historical originals remain unauthenticated; staffing approval, execution and customer benefit are unverified.

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

STRUCTURED ABSTRACT

Make the staffing decision conditional on the work contract

[AI] A funding constraint and an effort estimate answer different questions. A credible plan must connect both to a dated, skill-specific delivery schedule.

Element	Decision-relevant conclusion
Context [HR]	A planning lead replicated a sanctioned calculator, reconciled competing supply records, allocated a fixed staffing envelope and escalated unresolved method choices. The register describes artifacts, not an accepted resource increase.
Objective [AI]	Decide which original obligations can be delivered with the committed capacity, which must be deferred, and which residual request is traceable to work that remains.
Approach [RC/AI]	Separate one-time effort, recurring event frequency, concurrency, productive hours and availability calendars. Preserve alternative model readings rather than average them.
Findings [AI]	A cached example tests one calculation path. It cannot validate all project types or units. Uniformly reducing funded hours leaves the estimated effort unchanged unless scope or the production method also changes.
Decision [PD]	The method owner resolves units and calculation rules; delivery owners accept scope and dates; the resource sponsor approves committed capacity. Keep unresolved commitments conditional.
Implication [PD]	Use the proposed Effort-to-Delivery Capacity Contract (EDCC). Compare the full planning rule with a simple release sequence before buying additional diagnostic work.

[PD] Require an auditable workload identity, available capacity by period and skill, and an executable release schedule. No unexplained funding residual is a productivity result; no unvalidated programme value is automatically zero.

[HR] historical record; [RC] recalculated; [AI] analytical inference; [PD] proposed design. Protected source S1; original recovery, arithmetic and identity mappings remain private.

EXHIBIT 1 / WORKLOAD UNITS

Charge recurring work by its actual event and frequency

[AI] A weekly rate, a one-time task and a shared meeting are different quantities. Adding their printed numbers before reconciling units can change the staffing decision.

Work component	Required quantity	Decision risk
One-time delivery	Hours per completed project or phase	Applying the rate again every week repeats the same work.
Recurring project work	Hours per event × events per active week × active weeks	Charging once understates repetition; extending duration can add more work.
Shared portfolio work	Person-hours per distinct shared event, including participants; allocated once	Charging every project duplicates a genuinely shared meeting or report.
Operating support	Observed recurring tasks plus bounded interruption demand	A role count or empty financial-value cell is not a labor-hour estimate.

[HR] The record contains alternative readings of recurring rows, a task inventory without complete time estimates and per-project assumptions not supplied by the original plan. These support a method-resolution question; they do not authenticate either reading as the true demand.

A UNIT TEST BEFORE A STAFFING ASK

[PD] For every material row, record the event, duration, frequency, active period, sharing rule and source. Keep an unknown rate visible; do not replace it with zero or copy a rate from a different work unit.

[AI] Workload identity is dimensional accounting, not an empirical productivity model. Protected source S1; no historical quantities are displayed.

EXHIBIT 2 / CAPACITY DENOMINATOR

An effort allowance and a capacity reserve are not interchangeable

[AI] The same percentage label can encode different operating assumptions. State which side of the workload equation it changes.

Invented illustration	Calculation	Required capacity
Fixed work before adjustment	100 hours	100 hours
Add a 25% effort allowance	100×1.25	125 hours
Reserve 25% of gross capacity	$100 / (1 - 0.25)$	133.33 hours

[PD] These are hypothetical hours, unrelated to the historical figures. The allowance means additional work equal to a fraction of planned work. The reserve means only the remaining fraction of gross time is available. Applying both for the same interruption would count it twice.

[HR/AI] The source uses competing productive-hour conventions and a separate unplanned-work factor. It also mixes annual staffing statements with temporary availability. A reconciled headcount must distinguish funded seats, committed people, deployable skills and productive time in the relevant period.

DO NOT LABEL ONE READING A UNIVERSAL FLOOR

[AI] A lower bound requires all admissible changes to move the estimate upward. A credible alternative denominator can lower it, while missing recurring work can raise it. Preserve the unresolved range without calling its endpoints empirical confidence limits.

[PD] Numerical example is invented. Availability is reported only in aggregate; personal circumstances, dates and locations are excluded.

EXHIBIT 3 / MODEL CREDIBILITY

Reproducing one cached answer does not validate the calculator

[AI] Fidelity to an existing file is useful evidence about implementation. It is not evidence that every input path represents the work correctly.

Test layer	What a pass establishes	What remains open
Worked-example replication	The reproduced calculation matches that input and version	Whether other categories, durations and settings follow the intended rule.
Input-path verification	Each actual dropdown value reaches its intended branch	Whether the intended multiplier is empirically appropriate.
Dimensional review	Event units and frequency transformations are coherent	Whether actual time, sharing and interruption demand match assumptions.
Operating validation	Predictions are compared with independent task and calendar observations	Transport to another portfolio, season or staffing mix.

[HR] The record reports a label-to-formula mismatch and inconsistent accounts of which project types are affected. Original validation lists, formulas and cached values must be inspected together. A larger corrected estimate is not evidence that the correction is right.

RESOLVE THE RULE WITHOUT REWRITING HISTORY

[PD] The instrument owner should approve the intended mapping and retain a versioned before/after test. The planning lead can document a defect and an alternative calculation without claiming to have authored or authorized the underlying method.

[AI] Verification and validation are separate claims. Literature and source limitations are developed in the technical layer.

EXHIBIT 4 / ALLOCATION VERSUS DEMAND

Rationing funded hours changes coverage, not measured productivity

[AI] A numerical allocation can satisfy the funding rule while the remaining work still exceeds the resources assigned to it.

Invented planning ledger	Hours	Interpretation
Original estimated work	150	A demand estimate, conditional on its assumptions.
Explicitly deferred scope	30	Removed from this delivery commitment; still tracked.
Remaining estimated work	120	The effort estimate after the scope decision.
Funded allocation	90	Coverage of three quarters of that estimate.
Unresolved effort difference	30	Requires scope, timing, resources or a validated method change.

[HR/AI] The source retains an unreduced demand estimate while scaling the surviving allocation to the funded constraint. That is transparent rationing if reported as such. It cannot simultaneously establish that the funded rows retain their original deliverables, quality and dates.

WRITE THE TRADEOFF AS A DECISION


[PD] Identify what is removed, delayed or delivered differently; obtain acceptance from the delivery owner. An equal percentage reduction across rows is an allocation rule, not evidence that every project can absorb the same reduction.

[PD] Ledger quantities are invented. Deferral is not realized savings, and allocation percentages are not measured productivity gains.

EXHIBIT 5 / RELEASE SEQUENCE

Recurring overhead can make more concurrent starts deliver later

[AI] The recurrence unit changes the benefit of parallel work. Counting active projects without their overhead can mistake congestion for a staffing requirement.



[PD] Independent illustration: two projects each require 24 fixed hours; the shared team has 10 productive hours per week. Each active project requires 4 recurring hours per week unless the event is genuinely shared. Work is continuously divisible; there are no other bottlenecks or setup costs.

[PD] Concurrent release with per-project overhead leaves only 2 hours weekly for 48 fixed hours: both finish at week 24. Serial release leaves 6 productive hours: finishes are weeks 4 and 8. If the recurring 4 hours is genuinely shared, parallel completion is week 8. The figure compares final completion time, not average project delay.

TEST THE RELEASE RULE BEFORE EXTRAPOLATING HEADCOUNT

[AI] These are conditional scheduling identities, not field effects. Dependencies, project-specific throughput limits, skill restrictions or different overhead-sharing rules can change the answer. A calendar model must encode them explicitly.

[PD] Source-independent constructive proof, independently recomputed. Historical work rates and staffing quantities are not used.

EXHIBIT 6 / OBSERVABLE WORKLOAD

Count the work once and locate it on the calendar

[AI] A role description establishes responsibility, a task record establishes demand, and an availability record constrains execution. None is a substitute for the other two.

Evidence to collect	Required distinction	Owner of the decision
Task and event sample	Active human effort versus automated runtime; shared versus duplicate task	Operating owner validates work unit and sharing.
Recurring calendar	Frequency, duration, active weeks and deadlines; interrupted or canceled events retained	Method owner approves conversion into hours.
Skill and availability calendar	Committed availability versus nominal seats or undated hiring plans	Resource owner confirms deployable capacity.
Delivery acceptance	Original scope, quality and dates; queued and deferred work visible	Delivery owner accepts the executable commitment.

[HR] The record inventories recurring deliverables and identifies missing estimates and ownership proposals. Automated runs do not each imply a full human work session; an absent producer field does not prove no work occurred. Conversely, a financial value of zero does not imply that the task consumes no capacity.

INDEPENDENT CHECKS MUST COVER MISSING WORK

[PD] Sample from the original obligation frame and reconcile diaries, event logs and accepted outputs. Do not validate the model only on tasks that successfully appeared in its own table. Keep unresolved overlap and unrecorded work as explicit sensitivity cases.

[AI/PD] Source supports an unresolved measurement requirement, not a completed time study or accepted workforce change.

EXHIBIT 7 / DECISION AND ACCOUNTABILITY

Prefer the simplest release rule that meets the actual obligation

[PD] The proposed Effort-to-Delivery Capacity Contract binds work units, recurrence, usable capacity, sequence and accepted scope before a staffing commitment is treated as executable.

Decision gate	Release requirement
Size the work	Unit and frequency rules resolve; material recurring work and shared events are reconciled.
Place it on the calendar	An explicit feasible schedule covers the original obligations under the relevant skills and availability.
Compare practical options	Retain simple sequencing, scope reduction and explicit deferral; include their waiting and opportunity costs.
Test added value	The prospective benchmark evaluates scheduling loss and audit cost. It cannot establish historical staffing productivity or customer benefit.

[PD] In the homogeneous toy workload, always-sequential release is a strong comparator: a fallible recurrence audit can improve on usual concurrent starts yet still lose to this simple rule. Bounded precision screening may also fail to distinguish the assumed audit price from noise. Neither result should be engineered away.

ACCOUNTABLE NEXT STEP

[PD] The planning sponsor should commission a bounded unit-and-calendar audit, then ask the method owner for a ruling and delivery owners for explicit scope acceptance. Retain the funded envelope as a constraint, not as evidence that the estimated work has disappeared.


[HR] Contribution is a provisional registry-level hoop for reported replication, reconciliation and escalation. Method authorship, resource approval, implementation and benefit remain separate.

TECHNICAL ANALYSIS

Translate effort into a delivery claim

[PD] The Effort-to-Delivery Capacity Contract (EDCC) links typed work and recurrence to dated, qualified capacity, an authorized release sequence and independently observed completion.

EDCC: TEST THE RELEASE POLICY ON THE ORIGINAL WORK COHORT



Policy also incurs audit work in L. C affects evidence, feasible sequencing and completion.

Queued work stays in the denominator. O measures Y with possible error; it does not cause Y.

Time-unrolled proposed DAG: recurring work consumes capacity during active periods.

[PD] Z assigns a planning policy; E is fallible evidence, R the admitted work and sequence, Y completion and delay, and L full resource and commitment loss. C contains actual scope, skills, availability and recurring obligations. Independent O reconstructs completion. The graph is a proposed causal structure, not a historical effect estimate. An execution model must unroll recurring work over time rather than treating the feedback between duration and workload as an acyclic one-period equation.

[HR/AI] The record describes a replicated calculator, alternative readings, a constrained allocation and escalated method questions. An exact funded total establishes arithmetic closure. It does not establish qualified hours at each date, accepted ownership or preserved benefit. The smallest reported demand is not a global lower bound when plausible denominators, scope and recurrence conventions differ.

[AI] Employment FTE, productive FTE and staffing available inside a delivery window require separate definitions. OMB A-11 illustrates why an employment reporting convention can include paid leave while workload estimates exclude non-work time. Use that distinction as a methodological analogy; the organization must approve its own basis. A roster relationship is not proof of spare time, required skills or future availability.

Basis: proposed EDCC DAG [PD]; registry scope [HR/AI]. OMB A-11 (2025), sections 85.2-85.3; Sargent (2011). All historical coordinates remain private.

TECHNICAL ANALYSIS

Recurring work changes the release decision

[AI] Exact independent hypothetical example. Each project needs 24 fixed production hours; the pool supplies 10 hours per week. Active governance consumes 4 hours per week under the stated sharing rule.

Release / sharing rule	Completion time	Total work
One project; governance recurs	4 weeks	40 hours
Two parallel; governance per project	24 weeks	240 hours
Two sequential; governance per project	8 weeks	80 hours
Two parallel; one shared governance event	8 weeks	80 hours

[AI] For one project, $T = H/(p - g)$ when $p > g$. Charging governance once gives $(24 + 4)/10 = 2.8$ weeks, while recurring work requires $24/(10 - 4) = 4$. For two parallel projects, each carrying its own governance, $T = 48/(10 - 8) = 24$. Sequential release or a genuinely shared event changes the active burden. If governance consumes all capacity, this fluid model has no finite completion.

[AI] Assumptions: divisible work, constant productivity, no learning, no new arrivals and no overhead while queued. Governance stops when its covered work finishes. All original projects remain tracked, including the queued project. This counterexample proves that recurrence and sharing can change the decision; it does not prove parallel work is generally inferior or identify the historical overhead.

[AI] Total hours also miss eligibility. Three one-hour jobs fit three one-hour resource slots, but if two jobs can use only the same specialist slot, no complete assignment exists. Hall’s subset condition detects the bottleneck. A real model adds calendars, precedence, indivisibility and handoff constraints. Annual pooled capacity alone is insufficient.

[PD] ADMISSION TEST

Record whether each recurring event is per project, shared or already deducted from availability. Then test the original commitments against skill and time constraints.

Basis: independent exact recurrence, conservation and matching constructions [AI]; NASA CEH Appendix J (2015); Hall (1935). Later decision-law parameters form a separate hypothetical fixture.

TECHNICAL ANALYSIS

Verify the formula and the business meaning

[PD] Replication asks whether code follows a specified model. Validation asks whether that model answers the intended staffing and delivery question with adequate accuracy.

Validation layer	Required evidence or test
Input domain	Exact permitted values, normalized labels, units and version; exhaust all category branches, including default and missing inputs.
Formula semantics	Separate fixed hours, event hours and hours per active period; test zero scope, duration changes and shared-event deduplication.
Availability	Dated qualified hours after declared exclusions; identify which meetings, support and leave are already deducted.
Observed completion	Original due date, acceptance rule, completed output, rework and residual queue; validate forecasts on held-out work.

[HR/AI] The reported dropdown and exact-string predicates appear incompatible, including the new-work category. Until the original workbook is recovered, this remains a reported domain/predicate conflict. One matching cache can validate one executed path while every allowed input takes a fallback. A faithful copy of an erroneous formula is still a faithful copy; its agreement does not validate the intended categories.

[AI] Inflate demand by factor $(1 + b)$ and reserve fraction b of supply are different operations: $H(1 + b)/p$ versus $H/[p(1 - b)]$. Choose a documented convention; do not silently combine both. A productive-time haircut can overlap explicit governance rows. Repeated role allowances also require unique event and participant-time records before addition. A missing budget-dollar line does not prove missing hours.

[PD] Blind duplicate coders to assignment and outcomes when classifying recurrence, sharing and accepted completion. Sample original queued, active and completed work across skills and versions; report weighted confusion tables, prevalence and cluster-bootstrap kappa. An illustrative simple random sample of 400 gives worst-case agreement precision near ± 4.9 percentage points at 95%; this is not kappa power and ignores clustering. Agreement is not reference validity. Direct observation, calendars and traces must resolve uncertain labels without personal performance judgments.

Basis: proposed model and measurement validation [PD]; Sargent (2011), sections 4-8; NASA-STD-7009B (2024); Cohen (1960). No source formula is claimed executed in this review.

TECHNICAL ANALYSIS

Test admission on the original workload

[PD] Begin with a bounded shadow scheduling replay. A later field experiment tests a release policy on actual work; it needs separate consent, operational authority and calibrated measurement.

[PD] Freeze original project scope, due dates, skill requirements and calendar before team-domain assignment. Compare usual release, evidence-guided release and a strong simple serial rule with equal information rights. Audit recurrence before the decision, never reveal later completion to the policy. Keep all queued or deferred projects in the original cohort; a smaller admitted denominator cannot manufacture success.

[PD] Define $\Delta L = E[L(\text{usual}) - L(\text{EDCC})]$ per original team portfolio over a fixed horizon. L prices production, recurring work, rework, waiting and unmet commitments without double counting paid capacity. Charge audit cost k on that same denominator. Economic clearance requires a lower confidence bound above k and a prespecified deadline-harm guardrail. These are decision-resource effects, not customer or headcount savings.

[PD] Randomize the resource-sharing domain to contain interference. Prespecify baseline modifiers: skill scarcity, current backlog, recurrence structure, shared-event scope and dated availability. Realized admission, staffing substitutions and completed workload are mediators, not adjustment variables for the total policy ITT. A later observational rollout needs defensible comparison trends, sensitivity analysis and no selection on completion.

[AI/PD] Choosing the most overloaded teams invites regression to the mean. Under an assumed linear untreated conditional-mean model, expected follow-up is $m + \rho \times (x - m)$, where x is the selected baseline, m the population mean and ρ assumed persistence. Values 0.4, 0.7 and 0.9 imply 60%, 30% and 10% reversion of selected excess. These are sensitivities, not estimated historical fractions.

[PD] Pilot domain variances, reference errors and delay tails. A rough equal-arm 80% screen is $MDE \approx 2.80 \sqrt{2 \sigma_{\text{domain}}^2 / n}$, in the same units as k . For a compact finite replay law, enumerate attainable outcomes and actual null/economic boundaries instead of relying on a normal approximation at tiny n . More observations cannot make an unnecessarily costly policy beat the simple comparator.

Basis: proposed estimand, measurement and power [PD]; Hudgens and Halloran (2008); Barnett et al. (2005); Callaway and Sant'Anna (2021). No historical delivery effect or achieved field power asserted.

TECHNICAL ANALYSIS

Prespecify rejection, hold and attribution rules

[PD] EDCC remains an untested framework. A missed funding threshold can reflect imprecision; it is not proof that a scheduling effect is absent.

Proposition	Evidence that rejects or holds the claim
P1: semantic validity	Reject the specified model if an allowed category takes an unintended branch or a recurring unit is dimensionally inconsistent.
P2: recurrence-aware feasibility	Reject the schedule if verified active work exceeds eligible skill-time; require observed calendar and sharing evidence before pricing.
P3: improved original-cohort delivery	Hold unless the policy clears cost and deadline gates against usual and the prespecified strong simple comparator.
P4: transportable improvement	Hold extension when errors or effects change materially across prespecified skills, calendars or recurrence regimes.

[AI] NASA cost/schedule guidance motivates duration-dependent effort; Hall exposes skill bottlenecks; Sargent separates code from model validity. Rubinstein and colleagues motivate measuring context-sensitive switching, not a universal overhead rate. Deming cautions against reaction to sparse noise. Campbell/Goodhart warn that exact allocation targets can decouple from delivery; Ittner/Larcker motivates testing the full measure-to-action chain. Tucker/Edmondson motivates finding hidden support work. Composite weights cannot erase a binding resource constraint.

[HR/RC] The frozen inventory has 40 screens: one weak two-point temporary-bench comparison, 17 version audits, 11 design audits and 11 profiles. The source pair lies inside its fitted XmR limits; that does not establish stability. No stable operating phase or eligible Laney panel exists; Laney is not computed. No repeated completed-work panel is supplied. Model alternatives and designed allocations are not operating outcomes.

Limitations and attribution

[HR/AI] The original calculation files, accepted ruling and operational completion panel are not authenticated. A recovered planning artifact has unresolved version and provenance; it does not authenticate the registered calculator. The record supports a provisional registry-level hoop for assembling, reconciling and escalating a model, not sole invention, accepted hiring or delivered benefit. Recurrence, event sharing, calendar capacity, switching and loss prices require measurement. Simplified replay results cannot resolve unmeasured interference, recruitment feasibility or indefinite backlog costs.

Basis: proposed falsifiable claims [PD]; source and process-tracing limits [HR/RC/AI]. Fresh literature and common methodological references are applied to this capacity case; source measurements remain private.

PROPOSED CAPACITY TEST / 1 OF 4

Test the schedule on the original work, not the allocated percentage

[PD] The Effort-to-Delivery Capacity Contract (EDCC) is tested here as a synthetic scheduling replay. No live staffing, permission or customer promise changes.

[PD] Each independent team-domain starts with two original projects, each needing 12 productive hours and limited to 6 productive hours per week. Both are due at week 4. Capacity is 10 hours per week in base and 8 under stress, known equally before assignment. Hidden recurring work is either 2 shared person-hours per week (probability 0.60) or 3 per active project (0.40).

Context / true unit	Concurrent finishes	Serial finishes
Base / shared	3, 3 weeks	2, 4 weeks
Base / per project	6, 6 weeks	2, 4 weeks
Stress / shared	4, 4 weeks	2, 4 weeks
Stress / per project	12, 12 weeks	2.4, 4.8 weeks

[PD] Randomize whole teams to usual simultaneous release or a rule that releases in parallel only when a fallible predecision audit says the overhead is shared. Accuracy is 0.85 in base, 0.65 in stress and 0.95 with costed corroboration. These are invented, symmetric signal likelihoods; the acting rule never sees true sharing.

KEEP QUEUED AND LATE PROJECTS IN VIEW

[PD] All original projects remain. A hidden week-12 reference reconstructs completion against the original deadline; every toy case completes by then. Waiting starts at original intake, including queue time. Real unfinished work would require follow-up and explicit residual cost, not removal.

[PD] This fixture differs from the executive 24-hour illustration. Work is divisible, skills identical and no other disruptions occur. Actual throughput, recurrence sharing and reference completeness require calibration.

Large assumed gains do not establish precision at break-even

[PD] Primary ITT is terminal scheduling-loss reduction per original two-project team. A separate gate requires fewer original projects completed late. Added audit cost sets the investment threshold.

Domains per arm	80% MDE	Audit price	Zero power at break-even
Small base / 8	Beyond full gain	1.00	4.4%
Base / 64	3.282	1.00	12.3%
Low availability / 64	5.606	1.00	6.7%
Corroborated expansion / 128	4.007	1.50	16.8%

[PD] Exact team-level moments include both projects sharing the same work unit. Fixed variance-envelope normal radii use nominal one-sided 97.5% bounds. Investment requires the loss-reduction lower bound above audit cost; late-project reduction must also clear zero. MDE tests benefit against zero, not clearance above the price.

[PD] All four bounded designs remain inadequate at break-even. Corroboration changes the treatment and costs more; doubling the stress sample still leaves its MDE above price. No minimum sample or recruitable field panel is asserted. The 512 original projects in the expanded replay are hypothetical.

DIRECT FINITE CHECKS LIMIT THE APPROXIMATION

[PD] Integer-PMF convolution gives base-zero rejection 2.88687%, versus nominal 2.5%; expanded zero is 2.38890%. At full assumptions, exact marginal joint floors are 92.03% and 92.48%, but break-even detection is only 13.73% and 16.75%. A zero lower floor is uninformative, not zero joint power.

[PD] Direct finite tails are conditional on the invented law, not field guarantees; no Monte Carlo uncertainty is involved. Base bounds are not nominally calibrated for confirmation. Preregister a valid inference rule and comparison before any confirmatory use.

PROPOSED CAPACITY TEST / 3 OF 4

The simple sequence wins even when the audit beats current practice

[PD] Prices are judgments: one unit per project-week waiting, four per late project, and 0.10 per reserved staff-hour. Reserve all team capacity through week 12, even if idle. Actual work uses that prepaid capacity and is not charged twice.

Policy / total loss	Base	Stress	Corroborated
Usual concurrent	23.600	27.200	27.200
Audited release	19.840	22.372	20.736
Always serial	18.000	17.680	17.680
Perfect unit, same rule	18.000	18.880	18.880

[PD] Audit costs 1.00 per assigned team in base/stress and 1.50 with corroboration, including unsuccessful checks. Always-serial has no extra audit. It weakly improves waiting and lateness in every modeled state; its advantage survives any nonnegative prices on those categories. Even a perfect unit label can choose the wrong schedule under limited capacity.

[PD] Base net versus usual = $a \cdot h \cdot 4.76 - 1.00$. Here a is realizable scheduling-loss value and h is non-overlap. Only favorable category changes are discounted; all additional work and any adverse changes remain fully charged.

Net units / team	$h = 0.50$	$h = 0.75$	$h = 1.00$
$a = 0.25$	-0.405	-0.108	0.190
$a = 0.50$	0.190	0.785	1.380
$a = 1.00$	1.380	2.570	3.760

NO AUDIT PURCHASE CASE AT THESE ASSUMPTIONS

[PD] A gain against concurrent release does not defeat the simpler comparator. Earlier completion releases no payroll in this prepaid-horizon model. Neither this result nor an unresolved historical estimate establishes cash savings.

[PD] Unknown team outcomes on epsilon of original teams in each arm permit loss-contrast error up to $64 \cdot \epsilon$ and late-count error up to $4 \cdot \epsilon$, holding verified costs fixed. Uncertain work eligibility needs separate bounds.

PROPOSED CAPACITY TEST / 4 OF 4

Buy information only if it changes a worthwhile decision

[PD] A separate one-time research choice concerns a recurring-unit problem, an availability problem or neither. Its information value is not added to scheduling gains.

[PD] Prior probabilities are 0.45 / 0.25 / 0.30. Choose a recurrence study, a calendar study or defer research. Payoff vectors are (5,0,-1), (0,6,-1), (0,0,0), including downstream action cost. Signal accuracy and acquisition cost vary by scenario; each wrong signal is equally split across the other states.

Evidence scenario	Correct signal	EVSI	Cost	Net EVSI
Low	45.0%	0.05625	0.25000	-0.19375
Base	65.0%	0.58125	0.75000	-0.16875
High	85.0%	1.19250	1.50000	-0.30750

[PD] All 27 signal-to-action policies are enumerated per scenario. No-information value is 1.95; EVPI is 1.80. All three information purchases have negative net value at the declared prices. Low/base/high are alternative judgments, not confidence intervals.

THE NEXT INVESTMENT IS A BOUNDED MEASUREMENT DECISION

[PD] Reconcile person-hours, recurrence sharing and dated availability first. Retain a simple sequence as the comparator; measure whether dependencies or differing project throughput make it inadequate. Do not assume a more accurate label must justify a more expensive planning rule.

[PD] This is independent research utility, not a valuation of historical programmes. No accepted headcount request, implementation, actual avoided cost or customer benefit is claimed.

REPRODUCIBILITY / READER PROTOCOL

Rebuild the work before reproducing the staffing number

[AI] EDCC links the units and recurrence of effort to available skill-time and accepted obligations. A faithful spreadsheet copy can reproduce an incorrect or incomplete representation of that work.

Layer	Required reproduction
Source integrity	Verify the exact selected rows, source hash and original-recovery status. A readable planning artifact with unresolved historical provenance cannot authenticate the registered calculator or submission.
Model semantics	Map allowed input categories to every formula branch. Separate once-per-project effort, per-period governance, shared work, productive-time deductions and demand uplifts.
Population and supply	Reconcile included projects, formal deferrals, funded rows and roster definitions. Distinguish nominal seats, accepted resources and dated skill availability.
Decision-changing proof	Reproduce the independent duration-feedback and skill/feasibility constructions. These reveal missing constraints; they do not estimate the historical staffing requirement.
Evaluation and economics	Retain every original obligation, waiting time, unfinished work and implementation cost. Compare the audited policy with a simple release rule and preserve results unfavorable to the more complex approach.

[PD] A replication is complete only when another analyst can recover the same conditional staffing statement from the same work inventory, calendar and recurrence convention. If the result changes under a credible alternative, record that decision fork and its accountable resolver. Do not replace the fork with a falsely precise point estimate.

[AI] Public constructions use independent hypothetical quantities. The private annex retains source coordinates, personnel-sensitive context, unresolved variants and executable arithmetic. No indexed historical number is published where companion cases could permit inversion.

Protected evidence S1; source audit, model-semantics checks and exact constructive analyses. Work-package calculation is distinct from demonstrated productive capacity.

REFERENCES / 1

Primary sources and their role in this case

[AI] These sources support specific methods and distinctions. They do not authenticate the historical workbook or supply an estimated effect for this team.

NIST/SEMATECH. Individuals Control Charts, e-Handbook section 6.3.2.2. Report tentative limits only for eligible ordered observations; alternative model readings and roster definitions are not operating time-series signals.

Laney, D. B. (2002). Improved Control Charts for Attributes. *Quality Engineering*, 14(4), 531–537. Large denominators cannot repair incompatible outcome definitions. Repeated authentic binary counts and exposure are required before dispersion adjustment.

Cohen, J. (1960). A Coefficient of Agreement for Nominal Scales. *Educational and Psychological Measurement*, 20, 37–46. Independently classify effort frequency, shared work and owner acceptance. Report agreement, confusion and kappa separately from validation against actual work.

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. A trial that selects the most overloaded teams needs a frozen frame and an untreated comparison; apparent recovery can reflect selection alone.

Callaway, B., & Sant'Anna, P. H. C. (2021). Difference-in-Differences with Multiple Time Periods. *Journal of Econometrics*, 225, 200–230. For later observational deployment, separate adoption cohorts, diagnose pre-trends and estimate group-time effects rather than a pooled coefficient under heterogeneous timing.

Collier, D. (2011). Understanding Process Tracing. *PS: Political Science & Politics*, 44, 823–830. Attribute replication, demand modelling, allocation and escalation separately from calculator ownership, hiring authority and executed delivery.

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. Measure the value of learning which capacity assumption changes the action. Compare with an admissible simple policy and subtract evidence cost.

Original author, publisher, government and official-method sources. Access scope and case application are recorded privately.

REFERENCES / 2

Primary sources and their role in this case

[AI] These sources support specific methods and distinctions. They do not authenticate the historical workbook or supply an estimated effect for this team.

The W. Edwards Deming Institute. The Funnel Experiment. Do not retune staffing on routine noise or interpret alternative model versions as process improvement.

Campbell, D. T. (1976). Assessing the Impact of Planned Social Change. Rewarding a balanced funded column can conceal unserved demand. Preserve the demand ledger and the original obligations.

Goodhart, C. A. E. (1992). The Objectives for, and Conduct of, Monetary Policy in the 1990s. Reserve Bank of Australia Annual Conference, section 1(b)(i). A staffing ratio can lose meaning when it becomes the target. Budget compliance does not validate workload measurement.

Ittner, C. D., & Larcker, D. F. (2003). Coming Up Short on Nonfinancial Performance Measurement. Harvard Business Review, 81(11), 88–95, 139. Validate the chain from effort estimates to available productive time, completed obligations and mature operating outcomes.

Tucker, A. L., & Edmondson, A. C. (2003). Why Hospitals Don't Learn from Failures: Organizational and Psychological Dynamics That Inhibit System Change. California Management Review, 45(2), 55–72. Recurring workarounds consume capacity and can hide a persistent ownership or system constraint. Count the work rather than treating its unpriced status as absence.

OECD / European Commission JRC (2008). Handbook on Constructing Composite Indicators: Methodology and User Guide. Keep effort, funding, availability and readiness as separate dimensions. A weighted completeness score cannot substitute for a feasibility test.

National Aeronautics and Space Administration (2015). NASA Cost Estimating Handbook, Version 4.0, Appendix J: Joint Cost and Schedule Confidence Level Analysis, February 2015. Separate fixed scope effort from recurring duration-dependent work and connect both to a resource-loaded schedule.

Original author, publisher, government and official-method sources. Access scope and case application are recorded privately.

REFERENCES / 3

Primary sources and their role in this case

[AI] These sources support specific methods and distinctions. They do not authenticate the historical workbook or supply an estimated effect for this team.

Sargent, R. G. (2011). Verification and validation of simulation models. Proceedings of the 2011 Winter Simulation Conference, 183-198. Test implementation fidelity, conceptual assumptions, data and operational prediction as distinct validity questions.

Hall, P. (1935). On Representatives of Subsets. Journal of the London Mathematical Society, s1-10(1), 26-30. doi:10.1112/jlms/s1-10.37.26. Use a skill-eligibility subset check to show why total hours alone cannot establish a feasible assignment.

Rubinstein, J. S., Meyer, D. E., and Evans, J. E. (2001). Executive control of cognitive processes in task switching. Journal of Experimental Psychology: Human Perception and Performance, 27(4), 763-797. doi:10.1037/0096-1523.27.4.763. Treat switching burden as a task- and context-dependent hypothesis requiring local measurement, not a universal productivity haircut.

Office of Management and Budget (2025). Circular No. A-11, Section 85: Estimating Employment Levels and the Employment Summary (Schedule Q), sections 85.2-85.3. Distinguish an employment-reporting FTE from locally available productive workhours and document exclusions on the correct basis.

NASA (2024). Standard for Models and Simulations, NASA-STD-7009B, sections 4.2.3–4.2.6. Distinguish implementation and solution verification from conceptual and empirical validation; record the domain tested and aspects still unvalidated. Used as a model-credibility analogy, not a claim of NASA compliance.

Hudgens, M. G., & Halloran, M. E. (2008). Toward Causal Inference with Interference. Journal of the American Statistical Association, 103, 832-842. Original manuscript. Randomize at the resource-sharing team or portfolio domain; admitted and queued projects can interfere through capacity.

Original author, publisher, government and official-method sources. Access scope and case application are recorded privately.

Accept the capacity diagnosis; validate the workload contract

[AI] Author-led, agent-assisted review; not external peer review. The supported contribution is a reconciled capacity model and a falsifiable evaluation design.

Major issue	Revision and disposition
Worked-example overclaim	Treat a reproduced cached value as a bounded verification check. Test allowed categories and recurrence semantics before claiming model validity.
Denominator and scope forks	Keep productive-time conventions, committed supply and funded allocation distinct. Reconcile formal deferrals before sizing a residual staffing request.
Allocation versus productivity	Describe proportional allocation as rationing coverage. It does not reduce work content or demonstrate greater productivity.
Endogenous recurring work	Show how recurring governance depends on how long work remains open. State sharing and concurrency assumptions; do not transfer a synthetic duration to the historical portfolio.
Decision and benefit claims	Retain the simple release comparator, incremental audit cost and conditional power limits. No historic hiring, completed delivery, customer benefit or realized cash is established.

Partner critique: can a COO act on page one?

[AI] Yes. Require a governed work inventory and an explicit capacity convention; preserve the funding constraint while exposing unserved demand. Resolve the model owner’s recurrence and formula questions before converting a staffing ratio into a hiring commitment or performance promise.

Remaining limits and minor issues

[AI] Historical originals, observed task durations, unique recurring-work hours, accepted owner availability and mature delivery outcomes remain incomplete. The technical analysis consolidates these limits. Independent calculations can verify the conditional model; they cannot manufacture the missing operating evidence.

REVIEW DISPOSITION / PASS WITHIN STATED SCOPE

[AI] Publish the source reconciliation and testable capacity contract. Preserve unverified staffing, implementation and outcome claims as unresolved.

Final publication requires source, methods, independent numerical, partner and actual-page reviews recorded privately. This internal disposition is not external endorsement.