

Where the next engagement gets its capacity

01. The decision

A completed fictional Firm Capacity Blueprint. Northbank is an invented specialist advisory firm. All company records, hours, targets, demand and financial amounts in this case are synthetic teaching inputs dated 26 September 2026. No client supplied them. No system was installed or workflow acceptance test passed. Currency: AUD; amounts exclude tax. Version 1.0 · example design · prepared by Build n Bloom.

Northbank sells a fixed-fee operating-model diagnostic: evidence gathering, analysis, a reviewed findings report and an agreed action handover. More signed work is loading preparation and correction onto the same senior reviewers. The important revenue path runs from a suitable enquiry to approved scope, delivery, acceptance, collection and permitted reuse.

Recommendation: first repair the evidence-intake method, then compare a separately scoped review-ready report capability against that simpler method. Design all three connected opportunities now; install only what survives the comparison. Use existing approval and accounting tools for close-out. A custom AI close-out agent is not justified in this case.

The three designs are: W1, an accepted evidence brief; W2, a review-ready report and professional decision; W3, an accepted close-out record for finance and permitted reuse. W2 ranks first for a possible installation because its assumed senior reduction is largest. W1's common record fields and source permissions are prerequisites; its full automation is not.

What could change the recommendation: better intake alone could eliminate most report correction. If W2 cannot outperform that method at unchanged quality, stop the AI installation. Additional fees depend on qualified demand and accepted, collected work. With no additional demand or real avoidable expense, released hours remain unpriced capacity.

Open prepared-report.html and report-inputs.html for a substantial source-linked output and its originals. This shows the design in concrete work; no runtime or professional acceptance is claimed.

Decision today: use this example to assess the completeness of our design work. It authorises no purchase, professional advice, data access or production change. A real buyer requires its own evidence, quote and decision owners.

02. One revenue path; three connected opportunities

Established service: a diagnostic for an existing business unit, excluding staffing recommendations, litigation, regulated assurance, investment recommendations and changes to client systems. Fictional mix: 40 comparable completed engagements per year. Each has an agreed scope, client-authorised evidence and one professional reviewer.

Stage	Current handoff and return	Accountable role	System of record
Enquiry and scope	Fit recorded; partner approves fee, limits and requested evidence. Out-of-scope request returns to scope.	Managing partner	CRM and signed scope
W1: evidence brief	Analyst searches folders and notes; reviewer resolves missing or contradictory input. Accepted brief passes to W2.	Engagement lead	Case folder and source register
W2: report and review	Analyst prepares report; senior reviewer repairs unsupported findings. Returned findings reopen their source record.	Professional reviewer	Versioned report and decision register
W3: close and reuse	Operations reconstructs completion evidence; finance checks billing authority; practice lead approves any reusable method.	Operations lead; finance and practice lead decide	Completion register, accounting record, permitted method library

Connection contract: W1 emits case ID, approved scope version, source IDs/versions/rights, gaps, owner and acceptance decision. W2 references those IDs and adds finding IDs, original passages, reviewer disposition, report version and outstanding actions. W3 accepts only the professionally approved report version and customer acceptance evidence, with the relevant billing milestone and reuse disposition.

Reject an incomplete handoff to its originating owner with a reason and version. Never advance a case because a model says it is complete. A source change invalidates affected findings and downstream completion status until a reviewer reaccepts them. Collection is a separate finance-confirmed event, not an inferred consequence of sending an invoice.

Current-state map: enquiry → partner-approved scope → W1 accepted brief → W2 professionally approved report → customer acceptance → W3 authorised finance handoff → invoice/collection recorded by finance → rights-approved reuse. Returns go to the owner of the disputed source, finding or completion record.

Target map: same professional and commercial decisions; less reconstruction between them. Rules check completeness, AI prepares bounded source-linked work, and named people accept each consequential handoff.

03. Count the whole case

Synthetic ledger L-01 represents active preparation, review, correction, exceptions and coordination per comparable case. Each hour appears once. Waiting time is excluded. These are teaching assumptions, not observations or a forecast.

Stage	Senior h	Analyst h	Operations h	Total active h
W1 accepted brief	4	6	1	11
W2 approved report	12	28	1	41
W3 close-out	2	2	2	6
Complete existing case	18	36	4	58

At 40 cases, the synthetic annual baseline is 720 senior, 1,440 analyst and 160 operations hours: 2,320 active hours. W1 includes allocated selling/scoping effort of 1 senior, 1 analyst and 0.2 operations hours per completed case, including its share of unsuccessful pursuits. That manual effort stays unchanged in the targets. Its remaining 3 senior, 5 analyst and 0.8 operations hours cover evidence preparation and acceptance. This is a synthetic full-path allocation; actual pursuit effort and conversion must be measured. W2 contains 8 senior preparation/review hours and 4 senior correction hours; its 28 analyst hours include 20 preparation and 8 correction. The first comparison must count all of them, plus any new monitoring or exception work.

First-install hypothesis: W2 becomes 9 senior + 22 analyst + 1.5 operations hours = 32.5 hours, compared with 41. Gross release: 3 senior + 6 analyst - 0.5 operations = 8.5 hours per existing case. W1 and W3 remain at their baseline hours in this scenario.

Other design hypotheses, not included in first-install economics: W1 target 3 senior + 5 analyst + 1.2 operations = 9.2 hours. W3 target 1.5 senior + 1.5 analyst + 1.5 operations = 4.5 hours. Their separate gross releases would be 1.8 and 1.5 hours per case. Combined targets cannot be claimed before a fresh dependency and whole-case measurement.

For a real firm, log role/case/stage/activity; reconcile time entries to accepted versions; record returns and abandoned work. Sample ordinary, difficult and incomplete cases. Count a model's quicker draft only inside the complete accepted case.

Sample plan: establish 10 consecutive eligible baseline cases and comparable complexity bands. Compare at least 10 eligible proposed-method cases; retain all failures and learning effort. A reviewer checks comparability and decides whether the sample supports a conclusion. Ten is this case's proposed evaluation sample, not statistical certainty.

04. W1: the accepted evidence brief

Boundary: qualified enquiry/discovery record → partner-approved scope and authorised sources → engagement lead accepts a versioned brief. The partner decides pursuit, fee and terms before evidence preparation proceeds. Exclude fee negotiation, scope changes, legal interpretation and unapproved cross-client reuse. Status: complete example design; authority proposed; all runtime tests not run.

Current work: discovery/scoping and partner approval → folder search → interview notes → draft summary → senior reconciliation → accepted brief. Missing versions return to the analyst; ambiguous rights return to the source owner. Synthetic baseline: 4 senior, 6 analyst, 1 operations hour per case. Target hypothesis: 3, 5, 1.2; gross release 1.8 hours/case before shared operation and setup.

Options: unchanged process retains reconstruction; method repair introduces a required intake register and interview checklist; incumbent folder/template configuration may solve version control; rules validate mandatory fields and expiry; bounded AI can extract candidate facts from permitted documents. Prefer method repair and rules first. Defer AI extraction until field coverage and rights are reliable.

Proposed sequence: scope locked → source permission checked → source registered → candidate facts prepared → gaps routed → engagement lead accepts or returns. Operations owns the queue; engagement lead is the decision owner; managing partner is the deputy. A missing owner pauses advancement.

Inputs and permissions: approved scope, client-authorised interview records and documents, approved method checklist. Read only the allowlisted case folder. Write a draft brief and source register within that case; never edit the original. Service identity is case-limited, with authentication and revocation owned by the client administrator. Actual APIs, model provider, hosting region and retention remain client-specific checks before commissioning.

Candidate role contract: evidence-preparation role; one job: prepare cited facts and gaps. No external send or unrestricted search. Output: case ID, scope version, source/version/locator, extracted fact, ambiguity and missing item. Limit: one scheduled run plus one authorised retry; no more than 20 allowed documents and five minutes per run in this design. Enforce bounds in the controller, not prompt text. Model choice and token ceiling must be benchmarked and written into the commissioning scope.

Human decision: engagement lead accepts source sufficiency and the brief. Contradictory sources remain visible; the role cannot decide which authority wins. Handoff to W2 requires a receipt signed by that lead for the exact brief version, not a generated approval label.

Failure and recovery: missing/stale source blocks the affected item; forbidden folder denies access; interview contradiction records both passages; timeout leaves a draft; partial write uses case/version/run ID to prevent duplicate records. After one retry, operations escalates and uses the manual template. Revoke the service identity to stop future runs.

Proposed tests: W1-T01 ordinary case preserves every original locator; T02 missing source produces a gap and no accepted handoff; T03 contradictory versions show both; T04 cross-case access is denied without disclosed content; T05 outage/replay creates one draft and resumes only after owner decision. Engagement lead and client administrator inspect traces and receiving receipts. W1-T06: attempt a 21st source, second retry and over-time run; controller refuses excess actions, preserves partial output as draft and records escalation. All not run.

Recommendation: implement the intake method and rules as W2 prerequisites. Full W1 automation ranks second and needs separately demonstrated value. Stop custom extraction if the incumbent template provides the same accepted result with less burden.

05. W2: review-ready work and an explicit decision

Boundary: accepted W1 brief → professionally approved report and action handover. Exclude autonomous advice, customer issue, new scope and client acceptance. Status: complete example design; authority proposed; all runtime tests not run.

Current work: analyst rebuilds context, prepares analysis/report, senior repairs evidence gaps, analyst corrects, senior rechecks. Synthetic baseline: 12 senior, 28 analyst, 1 operations hour per case. Target hypothesis: 9, 22, 1.5; gross release 8.5 hours/case, including increased operations time. This does not remove the professional decision.

Options: unchanged work retains correction; method repair uses the W1 brief, standard finding fields and a reviewer rubric; incumbent search/templates may reduce reconstruction; deterministic checks catch missing citations; bounded AI can assemble comparisons and draft finding candidates. Compare the AI capability against the repaired method, not only the weak baseline. Choose AI only if full-case effort improves at the accepted quality standard.

Proposed sequence: brief accepted → source versions confirmed → candidate findings assembled → mechanical checks → reviewer queue → accept/return finding → report approved → authorised person issues report. Operations owns the queue; professional reviewer decides; another qualified reviewer is the named deputy. No available reviewer means no issue.

Sources and tools: only W1-approved sources and the versioned diagnostic method; no arbitrary web search, client-to-client retrieval or hidden method invention. Read approved case records; write draft findings and report versions; record decisions separately. Client administrator grants/revokes case-limited identity. Provider capability, data terms, permitted geography and live tool permissions must be verified for the actual client before use.

Candidate role contract: analysis-preparation role and report-preparation role share a case controller. Each has one job, typed inputs, source locators, method version and explicit unknowns. No fee, acceptance, send or invoice tool. Maximum two preparation steps and one authorised retry per step; five minutes per step; no more than 20 allowed sources. Controller logs run/version/action/result. Benchmark-specific model/token ceilings are unresolved procurement checks, never unlimited defaults.

Professional gate: reviewer sees originals, the proposed finding, the rule/check results and missing evidence together. Every material conclusion has a recorded accept, revise, reject or unresolved disposition. Only the approved report version passes to W3; a reviewer return invalidates the affected downstream draft.

Failure and recovery: missing citation blocks a material finding; contradictory evidence cannot be smoothed away; document instructions cannot grant tool authority; source change marks affected findings stale; permission refusal stops that source; timeout/partial output stays draft. Idempotency uses case/report-version/run ID. Manual report template and original sources remain usable. A revoked identity stops preparation; the human owner reconciles any completed draft before re-entry.

Proposed tests: W2-T01 ordinary finding traces to its passage and method; T02 fabricated inference is rejected/marked unsupported; T03 revised source invalidates affected findings; T04 embedded instruction cannot broaden access or issue work; T05 permission failure/outage/replay produces no issue or duplicate accepted version. Qualified reviewer checks output; client administrator checks execution controls and traces. W2-T06: force source-count, preparation-step, retry and time-limit breaches; controller refuses further actions, keeps partial output draft and records escalation. All not run.

Recommendation: first candidate installation after W1 method repair and permission readiness. Reverse if AI causes more total review/correction, obscures a source, cannot enforce permissions, or fails to outperform the simpler comparison at unchanged quality.

06. W3: close the work without losing its context

Boundary: professionally approved report and documented customer acceptance → finance accepts the completion evidence; practice lead approves permitted method reuse. Invoice issue and collection remain finance actions.

Status: complete example design; authority proposed; all runtime tests not run.

Current work: operations finds acceptance emails and versions, analyst confirms actions, senior clarifies completion, finance reconstructs billing milestone. Synthetic baseline: 2 senior, 2 analyst, 2 operations hours per case. Target hypothesis: 1.5, 1.5, 1.5; gross release 1.5 hours/case before shared costs. No collection acceleration is assumed or monetised.

Options: unchanged process retains missing completion evidence; a checklist plus existing CRM/accounting approval fields can carry the job; deterministic rules can match report version, acceptance record and milestone; AI summarisation is optional but adds uncertainty to a mostly structured check. Recommendation: no custom AI build. Configure existing tools if their verified capability supports the required receipt; otherwise use the controlled manual register.

Proposed sequence: report approval received → customer acceptance checked → milestone evidence matched → finance accepts/returns completion record → finance independently issues invoice and records collection → practice lead approves a separate de-identified method record. Operations owns the queue; finance owns billing/collection decisions; practice lead owns reuse rights. Named deputies are finance delegate and qualified practice deputy, each within its written delegation.

Sources and permissions: approved report version, customer acceptance record, signed scope/milestone, invoice/collection record and source-rights register. Operations reads case completion and prepares a finance handoff; no payment or invoice-issue permission. Finance retains accounting access. Practice lead can approve a method extract only after rights and confidentiality review. Actual incumbent approval/export features require a live capability check.

Work allocation: humans decide acceptance, billing and reuse; rules check fields, versions, duplicate key and required approvals. Candidate AI role: none. Reuse is an allowlisted method object stripped of names, client facts and confidential figures; redaction alone does not establish a right to reuse. Default is no reuse when rights are unknown.

Handoff: case ID, scope/milestone version, approved report ID, customer acceptance locator, outstanding conditions, authorised finance decision and receipt ID. Reuse record has its own rights source, method version and approving practice lead. Collection status changes only on finance's recorded receipt. One case/milestone key prevents duplicate completion entries.

Failure and recovery: conditional acceptance remains pending; missing milestone returns to partner/finance; contradictory versions reopen approval; duplicate event returns the existing receipt; accounting outage queues the handoff without issue; wrong permission denies access. Operations reconciles the queue on recovery, finance rechecks the exact milestone, and the practice lead reviews any changed reuse object.

Proposed tests: W3-T01 ordinary accepted milestone creates one prepared handoff; T02 conditional acceptance cannot complete it; T03 duplicate event creates no duplicate record; T04 absent rights blocks reuse; T05 accounting outage/replay preserves the pending state and cannot issue an invoice. Finance and practice lead inspect receipts; client administrator verifies access. All not run.

Recommendation: implement method/checklist with verified incumbent features; ranks third for a separate configuration decision. Do not justify it by invented payment speed or include its savings in W2 economics.

07. The shared operating boundary

The three designs share records, not unrestricted authority. Every item carries case ID, source version, rights state, owner, state, decision time and related receipt. Defaults are draft, no external issue and no cross-client reuse. A controller checks identity, tenant/case and permissions before each tool action.

Responsibility	Proposed holder	Allowed action	Boundary
Commercial scope	Managing partner	Approve scope/fee/change	No model commits the firm
Intake	Engagement lead	Accept brief or return with reason	Original source preserved
Professional work	Qualified reviewer	Decide finding/report disposition	AI prepares; reviewer owns advice
Coordination	Operations lead	Route, queue, reconcile returns	Cannot override approval
Billing/collection	Finance lead	Approve invoice and record receipts	No inferred collection
Permitted learning	Practice lead	Approve rights-cleared method object	No rights means no reuse
Access and recovery	Client administrator	Grant/revoke access, inspect logs	Scoped identity and receipt required

Data-flow specimen: authorised source → case-only register → draft brief → draft findings/report → professional decision → accepted completion evidence → finance system. Separately: rights-reviewed method extract → approved library. Client originals stay in the client-approved source environment. Any transfer to a model provider, storage location, log or backup needs its own documented permission, retention and deletion treatment.

Proposed data treatment: no confidential production data in evaluation until actual agreements and permissions are checked. For an installation quote, the client chooses permitted regions, provider data-use terms, retention windows for drafts/logs/backups, deletion/recovery obligations and incident contacts. These are open requirements, not claims about an installed environment. Store source locators instead of duplicating text where feasible; treat logs as confidential when they contain client facts.

Return and revocation: rejection records the source/finding/version and reason; correction creates a new version; downstream acceptance cannot silently persist. Revoke identities and pause the controller before a suspected access incident. Named client owner decides fallback and re-entry after reconciliation. Proposed retry ceiling: one authorised retry, then manual queue. Scheduled waits and owner absence are explicit states, not model guesses.

Ownership specimen: client receives the design, schemas, calculations, permitted source references, test cases, configuration documentation and operating runbook under its actual agreement. Existing platform licences, provider terms and pre-existing IP must be listed separately. No promise that every third-party component becomes client-owned. Another competent implementer can use this Blueprint; its own contract and access approval are required.

08. Economics finance can recompute

First-install scenario only: W2. W1 and W3 remain at baseline. Under synthetic ledger L-01 and target A-01, the complete case becomes 15 senior, 30 analyst and 4.5 operations hours = 49.5 hours. On 40 existing cases, W2 releases 120 senior and 240 analyst hours while adding 20 operations hours. Reserve 8 senior, 12 analyst and 12 operations hours of annual fixed oversight outside the case targets. Net steady-state capacity is 112 senior, 228 analyst and minus 32 operations hours. Operations needs real spare capacity; its burden has increased.

Synthetic A-02 assumes 80 available operations hours outside current work and qualified demand for six comparable additional cases. Each added case needs 15 senior, 30 analyst and 4.5 operations hours. A-03 adds first-year setup/adoption: 18 senior, 32 analyst and 12 operations hours. First-year capacity caps are $\text{floor}((120-18-8)/15) = 6$ senior; $\text{floor}((240-32-12)/30) = 6$ analyst; $\text{floor}((80-20-12-12)/4.5) = 8$ operations; demand cap = 6. The limiting cap permits six additional cases in this illustration.

At 46 cases, including setup and oversight, total first-year effort is 716 senior, 1,424 analyst and 231 operations hours. Available capacity is assumed to be 720, 1,440 and 240 respectively. Total labour increases because output increases; no payroll saving is claimed. The illustration assumes a full operating year: delayed installation, ramp, changed case mix or failed adoption reduces it. Fewer than six accepted and collected cases changes the financial result.

Synthetic A-04: 18,000 collected fee per added case minus 8,000 genuinely incremental delivery/sales/servicing cash cost = 10,000 contribution before system costs and existing fixed overhead. Annual incremental non-labour system cost is assumed to be 6,000. K is actual one-time cash investment, unknown: Blueprint, separate installation and any incremental setup expenditure counted once. Client setup hours are already in the capacity bridge; any paid overtime or contractor setup cash belongs in K. A buyer should assess the opportunity cost of those hours separately, without calling fixed payroll an avoided cash expense.

Added accepted and collected cases	Contribution before system costs	After annual 6,000 run cost	First-year result
0: no additional demand	0	-6,000	-6,000 minus K; capacity unpriced
3: lower realised demand	30,000	24,000	24,000 minus K
6: central conditional model	60,000	54,000	54,000 minus K

All amounts are AUD excluding tax and synthetic assumptions, not rates, quotes or a forecast. Central first-year mathematical break-even is $K = 54,000$ under these inputs, not a recommended price or acceptable risk threshold. First-year required cases are $\text{ceiling}((6,000 + K)/10,000)$, limited by capacity and actual demand. A real buyer may require substantially more surplus. Recurring break-even is one added collected case under the same cash assumptions.

Strong contrary case: W2 still needs 10.5 senior hours. Annual senior release becomes 60 hours; after 18 setup and 8 oversight hours, only 34 remain. Each additional full case requires 16.5 senior hours. $\text{floor}(34/16.5) = \text{two cases}$; senior is the limiting role. Contribution after annual run cost is 14,000 before K. Better intake alone could produce the improvement without AI. Test that method first.

The six-case/54,000 model compares the proposed package with the unrepaired baseline. Incremental AI value against the required repaired-method comparison is unmeasured. Do not use this package-level result to justify an AI installation price.

Do not add charge-out-rate labour value to contribution earned using the same hours. No actual demand, accepted quality, collected fee, marginal cost, operations headroom, net effort or adoption has been measured. Hourly billing can reduce fees when hours fall. The arithmetic can be right while the investment decision is wrong.

09. What an implementer can commission

Rank 1: W2 preparation-to-review capability, after the W1 intake method and rights register are usable. Start: accepted, permission-checked evidence brief. End: source-linked report version and a real professional-review decision; customer issue remains manual/authorised. It may contain two bounded AI preparation steps, rules, queues and human returns. It does not install the entire revenue path.

Required deliverables: source/rights register; typed brief/finding/report/decision schemas; case-limited identities; controller and action ceilings; original-locator display; stale-source invalidation; duplicate/replay handling; human queue; exportable runbook, test cases and traces; operator training; revocation/manual fallback. Actual vendor choice, access, security terms, price, schedule, change control, correction obligations and stabilisation period belong in the installation agreement.

Proposed client participation: engagement lead selects representative cases and method; qualified reviewer defines the rubric and reviews comparisons; operations reconciles effort/returns; finance verifies demand/contribution; administrator verifies permissions and provider terms. Example setup effort is 18 senior, 32 analyst and 12 operations hours, not a commitment to the buyer. Availability and exact effort must be quoted after discovery.

Functional acceptance specimen: on the agreed representative case set, no unauthorised issue or cross-case access; material findings have source locators and explicit reviewer dispositions; a changed source invalidates the affected finding; replay produces no duplicate accepted record; revocation prevents new actions; forced source/action/retry/time-limit breaches stop further actions and preserve draft output with escalation; outage permits documented manual recovery. Require traces and receiving receipts for ordinary and failure cases. These are proposed acceptance requirements; every execution result is not run.

Business evaluation after acceptance: compare complete case effort and quality against the repaired method, by role and case mix, including learning, oversight and exceptions. Record accepted throughput and later collections separately. A technically accepted installation can fail the business case; stop expansion if evidence is inconclusive or negative.

Rank 2: W1 extraction capability, only if the repaired intake method leaves enough eligible preparation to justify it. Separate scope, benchmark and price. Rank 3: W3 incumbent configuration, conditional on actual approval/export features and completion burden. No custom AI recommendation.

Stabilisation and exit: name the operator, reviewer, administrator, incident contact and source-refresh owner. Quote the in-scope correction period and response duties; there is no default 30-day promise. Choose priced continuity only if justified, or hand over configuration, permissions, logs, tests, cost register, known limitations and recovery instructions. Revoke supplier access when the agreed handover is complete.

10. Inspect the evidence and challenge the case

ID	Artifact / label	What it supports	Limit and responsible role
L-01	Included synthetic effort ledger	Reproducible teaching baseline	Not measured; finance/operations must replace
A-01	Included synthetic design targets	First-install comparison hypothesis	Untested; professional reviewer validates
A-02	Demand and operations headroom assumptions	Added-case constraint arithmetic	No actual pipeline; growth/operations verify
A-03	Setup and oversight assumptions	First-year capacity bridge	No actual adoption record; operations verifies
A-04	Financial model inputs	Cost and contribution arithmetic	No quote or client results; finance replaces
X-01	Included fictional challenged finding and decision	How a professional boundary is documented	Authored example; no live AI execution
P-01	Separate real public-source review linked in the walkthrough	Inspectable source interpretation and correction	Historical public data; no advisory-client outcome

Open decisions for a real buyer: verify source rights/provider terms before model use; finance confirms real demand and marginal costs before any money claim; operations confirms its headroom; qualified reviewer confirms method and case comparability; managing partner approves the quote-specific purchase. If one remains unresolved, defer the relevant action. No fictional buyer signature or permission is asserted here.

Completed fictional reviewer disposition X-01: the draft says, 'The client can safely reduce its leadership team by two roles.' Source S-02 supports a handoff delay, not a staffing conclusion. The fictional professional reviewer rejects the inference as unsupported and outside the agreed staffing exclusion, retains the supported observation and returns any staffing request to the managing partner for a separate scope decision. The revised finding becomes: 'The sampled approval records show an unresolved handoff; leadership structure cannot be concluded from this evidence.' This is a designed example of judgement, not advice to a real firm.

What this sample proves: the structure, depth and inspectability of a Blueprint deliverable. What it does not prove: installation performance, professional endorsement, actual client time savings, revenue, margin, security certification or independent assurance.

Next useful action: open the annotated walkthrough, recompute the ledger, or discuss a comparable constraint in the free 25-minute Firm Capacity Fit Call. No confidential client files are required. If appropriate, a paid Blueprint maps your revenue path and fully designs three connected opportunities. Installation has a separate scope and price.