A COMPANION TO

THE BOTTLENECK IS NEVER THE STACK

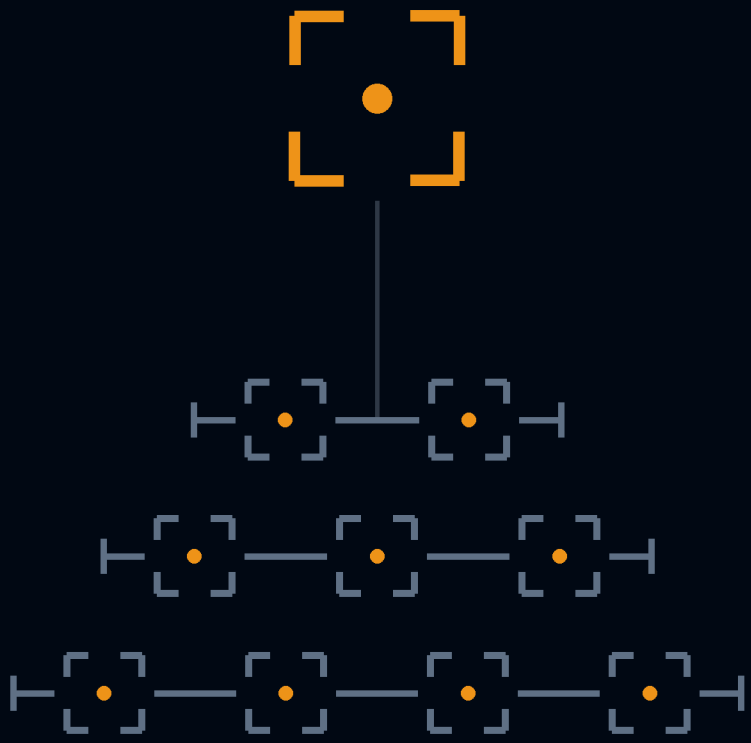

Putting the Book to *Work*

A Facilitator Workbook

VINNY CARPENTER

Putting the Book to Work

A Facilitator Workbook for The Bottleneck Is Never the Stack

Vinny Carpenter

About this workbook

A facilitator workbook for The Bottleneck Is Never the Stack. The book makes the argument; this workbook helps a leadership team run Platform Pilot Zero, redesign the work, charter a bounded product pilot, set controls, and make an evidence-based scale decision.

Contents

INTRODUCTION

Start Here: The Working Session

One prepared room, two linked charters, and one day to make the operating decisions that a memo cannot make.

> *The book makes the argument. This workbook puts a cross-functional leadership team around the same table and asks it to leave with a platform path it has tested and a product pilot it can govern, measure, stop, and learn from.*

This is a workbook, not a shorter version of the book.

Read Part IV of *The Bottleneck Is Never the Stack* before using it. The book owns the argument; this workbook owns the canonical artifact forms, fields, and completed examples, and it works without the book open beside you. The pages that follow are working surfaces: prompts, blank forms, filled illustrations, and decision gates.

The workshop produces two linked charters: Platform Pilot Zero, the path owner's own end-to-end journey, and one bounded product pilot. The aim is not to approve an enterprise transformation in a day. The aim is to authorize Platform Pilot Zero and design a credible 90-day product test that can begin only after the shared path earns that first outside adopter.

Choose a working mode

The workbook supports five working modes. Every mode ends with a named decision and an owner. No mode authorizes more than the decisions listed for it.

The 30-minute diagnostic. Participants: the sponsor, one engineering leader, and a facilitator. Pre-work: only the sponsor's one-sentence decision. Artifacts completed: the boundary screen, the readiness screen scored from what the room knows, and a next-decision note. Decisions permitted: whether to schedule a design session, and which evidence to gather first. Not permitted: authorizing any pilot, naming teams publicly, or setting dates.

The half-day design session. Participants: sponsor, outcome owner, engineering and platform leads, a risk partner, facilitator, and scribe. Pre-work: the source pack, at minimum recent work items and review-capacity evidence. Artifacts completed: readiness worksheet, both work-redesign passes, a risk-classification draft, and a draft pilot charter. Decisions permitted: proceed to a full workshop, prepare first, or choose different work. Not permitted: authorizing production use or scale.

The one-day workshop. The complete sequence in this workbook, run to its published timeboxes, with the full room. Pre-work: the complete preparation checklist. Artifacts completed: the full packet through the gate record. Decisions permitted: authorize Platform Pilot Zero; approve, conditionally

approve, or decline the 90-day product pilot. Not permitted: approving expansion beyond the chartered pilot.

The two-week sequence. The same sessions as the one-day workshop, run as shorter cross-functional sessions inside two weeks. Each session names a between-session owner who carries the evidence gaps to the next one. Artifacts and permitted decisions match the one-day workshop. Do not stretch past two weeks; evidence and assumptions decay while the calendar moves.

The monthly review or phase gate. Participants: sponsor, operating owner, outcome owner, control and people partners, and the named decision owner. Pre-work: the one-page pre-read per product line, the current scorecard, the capacity worksheet, and the exception register. Artifacts completed: a monthly operating review record or a signed phase-gate decision record. Decisions permitted: advance, hold, stop, or a time-bound exception. Not permitted: redefining measures mid-window, or treating time passed as approval.

The working surface

The PDF editions are print worksheets, not interactive PDF forms. For a digital session, copy the Markdown forms into an approved collaboration document or repository and type there. For an in-person session, print the workbook and use its writing rows, then transcribe the signed decisions into the versioned packet. In either mode, the shared copy is the working surface; the PDF is not a fillable system of record.

The decision the sponsor brings into the room

Write the decision in one sentence before invitations go out.

```
We are here to decide whether the platform owner may run Pilot Zero on
_____________________________________________________________
and whether ___________________________________________
should then run a 90-day product pilot of _____________________
to improve _____________________________________________,
subject to the Platform Pilot Zero gate and _____________________________ .
```

If the sponsor cannot complete that sentence, schedule a framing session first. A workshop with no decision is a long meeting with stationery.

Who belongs in the room

Invite the smallest group that can make the decision and own its consequences. People may hold more than one role, especially in a smaller organization.

Seat	What this person must be able to decide
Executive sponsor	Whether to fund, protect, hold, or stop the pilot

Seat	What this person must be able to decide
Outcome owner	Whether the business result is worth pursuing and whether it occurred
Product or domain lead	Scope, customer effect, and work priority
Engineering lead	Technical boundary, delivery ownership, and staffing
Platform lead	Paved path, pipeline, agent configuration, and shared capability
Spec owner	Whether intent is explicit enough to build and evaluate
Evaluator	Whether evidence is sufficient to accept generated work
Security, risk, privacy, or legal partner	Risk tier, obligations, and required controls
People leader	Role changes, capacity protection, mobility, and apprenticeship
Finance or operations partner	Baseline, value measure, and benefit validation
Facilitator	Process, timeboxes, unresolved questions, and decision readback
Scribe	Live artifact capture, decision log, action register, and version control

The sponsor should not facilitate. The person with the most positional authority cannot both shape the discussion and observe whether the room is reaching a real decision.

Preparation checklist

Complete this at least three working days before the session.

```
DECISION AND PEOPLE
[ ] Sponsor approved the one-sentence decision.
[ ] Outcome owner can attend the decision gates.
[ ] Control partners were invited as designers, not final reviewers.
[ ] Facilitator and scribe are named.

SOURCE PACK
[ ] Five to ten recent work items from the candidate team.
[ ] Current workflow from request through production and feedback.
[ ] Pipeline coverage, routine elapsed time, and break-glass evidence.
[ ] Review volume, active review time, re-review loops, and queue time.
[ ] Written standards and current agent or automation inventory.
[ ] Candidate outcome baseline, source, owner, and measurement window.
[ ] Known regulatory, data, customer, money, identity, and vendor constraints.
[ ] Current incident, exception, and production-change evidence.

ROOM AND ARTIFACTS
[ ] One editable copy of the Markdown forms or an approved collaboration copy
    is ready for live editing, or printed worksheets are ready for handwriting.
[ ] The workbook's forms are ready as the canonical field definitions.
[ ] The decision log, parking lot, and action register are visible.
[ ] A repository or other versioned home for final artifacts is named.
```

Do not substitute slideware for the source pack. A team that cannot retrieve its evidence has learned something important before the workshop begins.

One-day run of show

Time	Session	Decision or artifact
8:30	Open, scope, and working rules	Sponsor's decision and non-goals
9:00	Readiness and work suitability	Proceed, prepare, or choose different work
9:45	Current-state work redesign	Bottlenecks, handoffs, rework, and control points
10:30	Break	
10:45	Future-state work redesign	Paved path, human decisions, agent actions, feedback
11:45	Operating model and ownership	Canvas, exclusions, and org-design hypothesis
12:30	Lunch	Scribe resolves only formatting, never decisions
1:15	Capabilities, roles, and capacity	Named hats and a staffed evaluation plan
2:15	Risk and control profile	Required controls, stop path, and rollback path

Time	Session	Decision or artifact
3:00	Break	
3:15	Evidence and scorecard	Metrics dictionary and decision thresholds
4:00	90-day pilot and phase gate	Plan, gate owner, evidence window, outcomes
4:45	Readback and commitments	Decision, dissent, actions, owners, and dates
5:00	Close	

When the source pack or decision authorities cannot support one full day, use the two-week sequence mode: the same sessions as three shorter workshops, diagnose the work, design the operating hypothesis, and commit to evidence, each with a named between-session owner.

Working rules for the facilitator

1. Ask for evidence before opinion. Mark unknown facts as unknown.
2. Separate the current state from the desired state. Do not draw aspiration as if it already exists.
3. Put controls in the future-state flow, not in a review box at the end.
4. Name a person for every decision, artifact, agent, measure, and action.
5. Preserve dissent in the decision log. Forced consensus hides risk.
6. Treat thresholds as directional decision rules, not universal benchmarks.
7. Stop org-chart discussion until the work redesign is visible.
8. Read every gate aloud: advance, hold, or stop; then name the fallback and recovery disposition.

Definition of done

The workshop is complete when the sponsor can first read back the Platform Pilot Zero prerequisite and then the product-pilot decision without interpretation:

```
Platform Pilot Zero work class: _____
Pilot-zero completion and recovery evidence: _____
Pilot-zero decision owner / gate: _____
Pilot outcome: _____
Bounded work and non-goals: _____
Operating owner: _____
Risk tier and control owner: _____
Evaluator demand / staffed capacity: _____
Evidence window and decision thresholds: _____
Advance condition: _____
Hold condition: _____
Stop condition: _____
Fallback and recovery disposition: _____
First gate date: _____
```

The next chapter tests whether the team and the proposed work deserve to enter that sentence.

PART I: CHOOSE THE WORK

Choose the Team and the Work

Readiness is two decisions: can this team run the loop, and is this work worth learning from?

> *A high-scoring team can still choose a useless pilot. A valuable problem can still overwhelm a team with no review capacity. The room has to test both.*

Timebox: 45 minutes

Decision: proceed, prepare, or choose different work

Outputs: team-readiness worksheet, two work-suitability screens, evidence gaps

Use Chapter 10, *Where You Actually Are*, for the argument behind this session. This session owns the working forms.

Facilitation sequence

1. Score from the source pack, not from reputation.
2. Require one piece of evidence for every score of 2.
3. Resolve large scoring differences before adding the total.
4. Screen the proposed work separately.
5. Record missing evidence as an action, not as a generous score.

Worksheet A: team readiness

Score each condition **0 absent**, **1 partial**, or **2 systematic**, using this rubric. Preserve the five individual scores. The total cannot hide a zero in a load-bearing condition.

Condition	0, absent	1, partial	2, systematic
Spec discipline	Tickets name outcomes vaguely; acceptance lives in heads	Templates exist but use and review vary	Spines or equivalents are versioned and reviewed before build
Pipeline coverage	Manual or ticket-driven production change	Most routine changes use the pipeline; exceptions vary	Supported changes use the pipeline; break-glass is audited; speed meets the local service level
Evaluation capacity	Review queues behind one senior; rubber stamps under load	Reviewers are named but capacity is unprotected or single-threaded	Reviewers are named by domain, protected by hours, backed up, and latency stays inside plan

Condition	0, absent	1, partial	2, systematic
Written standards	Tribal knowledge or stale wiki	Useful standards exist but inheritance or incident updates are manual	Versioned standards are inherited by tools and updated after findings
Ownership	Unknown agents and scripts in the wild	Inventory exists but owners, telemetry, or stop paths are incomplete	Register is current: owner, spec, telemetry, evaluation evidence, and stop path per agent

Condition	Score	Evidence seen	Gap owner and date
Spec discipline			
Pipeline coverage			
Evaluation capacity			
Written standards			
Agent and workflow ownership			
Total out of 10			

Initial decision rule: proceed at 7 or higher with no zero in pipeline coverage, evaluation capacity, or ownership. Treat this as a screening rule, not a maturity grade. Local policy may be stricter. If it differs, record the approved local rule and its owner.

```
Local readiness rule, if different: _____________________________
Approved by: _____________________  Review date: _______________
Readiness decision: [proceed] [prepare first] [do not proceed]
Reason: ____________________________________________________
```

Boundary screen

Before scoring suitability, place the candidate work in the boundary table from the book's introduction. The posture column sets what the charter must contain.

Work characteristic	Likely constraint	Required posture
Bounded, reversible, observable	Framing and evaluation increasingly dominate	Use the full loop and measure it
Legacy or integration-heavy	Implementation and system knowledge may remain binding	Narrow scope and preserve specialist review
Regulated or safety-critical	Authority, evidence, and evaluation capacity dominate alongside implementation	Apply the approved control profile
Novel, ambiguous, or customer-discovery work	Product judgment remains upstream of specification	Do not automate uncertainty into false precision
Difficult to reverse or weakly observable	Verification and recovery may be insufficient	Hold, redesign, or exclude

```
Candidate work: _____________________________________________
Row that fits best, and why: _____________________________
Posture the charter must reflect: _____________________
```

Worksheet B: work suitability

Score **0 poor fit**, **1 uncertain**, or **2 strong fit**.

Test	Score	Evidence or unresolved question
A measurable result matters and has a credible baseline		
Scope is bounded enough for a 90-day decision		
Dependencies are known and influenceable		

Test	Score	Evidence or unresolved question
Risk and obligations are understood well enough to classify		
A rollback or disable path is feasible		
Learning should transfer to at least one other team or workflow		
An outcome owner will judge whether value occurred		
Total out of 14		

A zero in value, baseline, risk understanding, rollback feasibility, or outcome ownership blocks the charter. Resolve it or choose different work. A high total does not compensate for one of those zeros.

Evidence-gap register

Unknown that changes the decision	How it will be answered	Owner	Due	Effect if unanswered

The facilitator asks: *Would the room make a different decision if the answer went the other way?* If not, the item belongs in the parking lot. If yes, it belongs here.

Screen two linked candidates

Run Worksheet B twice. The first screen is **Platform Pilot Zero**: one repeated class of the platform team's own work that will exercise the path before a product team inherits it. The second screen is the proposed product pilot. They need separate baselines, owners, risks, exclusions, and charters.

Platform Pilot Zero is a prerequisite, not a warm-up task hidden inside the product pilot. The product candidate may be designed in parallel, but it does not enter production until Platform Pilot Zero has completed one end-to-end working loop, recorded the friction as platform backlog, and exercised the required stop and recovery path.

Harborline illustration

Harborline is a composite. The scores below illustrate the method; they are not reported results.

Platform Pilot Zero

Candidate team: platform engineering

Candidate work: standardize one repeated service-account request class

Work test	Score	Evidence
Value and baseline	2	Identity operations owns request elapsed time and rejected-request rate
Bounded scope	2	One pre-approved service-account class; no production-role redesign
Dependencies	1	Identity and security must agree on permissions and evidence fields
Risk understood	2	Privileged access, expiry, and traceability obligations are known
Rollback or disable feasible	2	Disable the paved route and return requests to the verified manual path
Transferable learning	2	The same templates, register, gates, and telemetry will support product pilots
Outcome owner	2	Platform operations lead accepts the internal-service outcome

Pilot-zero decision: proceed through one completed service-account journey, conditional on a security-approved control profile, named reviewer capacity, and a tested return to the verified manual path. Its charter must close before the payments pilot enters production.

Product pilot

These are the authoritative Harborline readiness values.

Candidate team: payments product team

Candidate work: consolidate payment-retry behavior behind one reviewed, gated path

Team condition	Score	Evidence
Spec discipline	2	Payment invariants are versioned; acceptance review already occurs
Pipeline coverage	2	Supported release paths are automated; evidence still needs consolidation
Evaluation capacity	1	Ledger reviewer is named; backup capacity is not yet protected
Written standards	1	Ledger and retry rules exist in separate repositories
Ownership	1	Inventory finds one temporary retry helper that is not registered
Total	**7**	**Eligible for a bounded pilot, not automatic expansion**

Work test	Score	Evidence
Value and baseline	2	Finance operations owns monthly failed-retry cost and source
Bounded scope	2	Retry policy only; no message-schema or settlement redesign
Dependencies	1	Ledger and gateway teams must agree on the invariant
Risk understood	2	Money movement and duplicate-entry risk require the highest local tier
Rollback feasible	2	Route can return to last verified retry policy
Transferable learning	2	Pattern applies to other payment workflows
Outcome owner	2	Finance operations lead accepts the role

Product-pilot decision: proceed to design, conditional on registering or retiring the temporary helper, protecting evaluator backup capacity, naming the standards keeper, getting the ledger owner to approve the duplicate-entry invariant, and closing the Platform Pilot Zero gate.

Negative screen: settlement reconciliation

The room also screens the candidate everyone expected to win, and it fails for durable reasons.

Candidate work: automate reconciliation on the vendor-managed settlement appliance

Boundary screen: difficult to reverse and weakly observable; the required posture is hold, redesign, or exclude.

Work test	Score	Evidence
Value and baseline	2	Manual reconciliation effort is visible, recurring, and owned
Bounded scope	1	Documented workflow, but exception handling crosses the vendor boundary
Dependencies	0	Vendor-managed appliance; no influence over its change path
Risk understood	1	Settlement obligations are known; failure modes inside the appliance are not
Rollback or disable feasible	0	No verified rollback on the appliance; recovery is the vendor procedure
Transferable learning	1	Little transfers while the appliance owns the state
Outcome owner	2	The settlement owner would accept the outcome

Decision: exclude, and record the exclusion. The work stays on the vendor change procedure with dual human approval and the recovery drill. Risk approves the boundary for six months. A vendor API release, a change-volume increase, a control failure on the manual path, or a missed review date reopens the decision. The complete entry lives in the exclusion register.

Gate 1: choose

Read the decision into the log.

```
[ ] PROCEED. Team screen passes, work screen passes, and conditions have owners.
[ ] PREPARE. Candidate remains viable, but a readiness condition needs work.
[ ] CHOOSE DIFFERENT WORK. Team can learn, but this candidate fails the work screen.
[ ] STOP. No credible owner, baseline, control path, or capacity exists.

Candidate: [Platform Pilot Zero] [product pilot]
Decision: _____  Owner: _____
Conditions: _____
Evidence due: _____  Revisit date: _____
```

Do not announce the pilot yet. First make the current work visible, including the parts everyone has learned to work around.

PART II: REDESIGN THE SYSTEM

Redesign the Work

Map the current flow honestly, then design a paved path that makes the preferred behavior the easiest behavior.

> *The org chart is a tempting place to start because it is visible. The work is harder to see and more important to redesign first.*

Timebox: 105 minutes, with a break between current and future state

Decision: whether the proposed path removes a real bottleneck without hiding a new one

Outputs: current-state canvas, future-state canvas, exclusions, first-user plan

Use Chapter 11, *Redesign the Work at Enterprise Scale*, for the argument behind this session. This session owns the working forms.

First pass: current-state canvas

Choose one representative work unit and follow it from request to production and feedback. Draw what happens, not what policy says should happen.

```
WORK UNIT: _____________________________________________________
OUTCOME THIS WORK SERVES: ______________________________________
START EVENT: _____________________  END EVENT: ________________

Step / handoff        Who or what acts     Input / artifact     Decision made
1. ________________ ________________ ________________ ________________
2. ________________ ________________ ________________ ________________
3. ________________ ________________ ________________ ________________
4. ________________ ________________ ________________ ________________
5. ________________ ________________ ________________ ________________
6. ________________ ________________ ________________ ________________

Where work waits: ______________________________________________
Where intent is re-explained: __________________________________
Where rework enters: ___________________________________________
Where a human judgment is essential: __________________________
Where a control is bypassed or added late: _____________________
Where feedback reaches the wrong owner or arrives too late: ___
```

Add observed evidence beside the map: elapsed time, active time, queue time, review loops, exceptions, and incident-linked defects. Precision is useful; invented precision is not.

Bottleneck and burden inventory

Observation	Type: clarity, capacity, policy, handoff, tooling	Evidence	Who bears the burden today

Prompt the room with four questions:

- Which wait would remain even if generation became free?
- Which handoff exists only because intent is stored in someone's head?
- Which review is a real decision, and which is compensation for weak inputs?
- Which burden is currently absorbed by a few experienced people and therefore absent from the process map?

Second pass: future-state canvas

Do not automate every current step. Remove, combine, or move work before adding an agent.

```
FUTURE WORK UNIT: _____________________________________________
VALUE HYPOTHESIS: _____________________________________________

Step / gate           Human decision       Agent action         Evidence produced
1. ___________________ ___________________ ___________________ ___________________
2. ___________________ ___________________ ___________________ ___________________
3. ___________________ ___________________ ___________________ ___________________
4. ___________________ ___________________ ___________________ ___________________
5. ___________________ ___________________ ___________________ ___________________
6. ___________________ ___________________ ___________________ ___________________

Versioned intent artifact: _____________________________________
Inherited standards: ___________________________________________
Pipeline gates: ________________________________________________
Named evaluator and authority: _________________________________
Production feedback routed to: _________________________________
Break-glass path and audit evidence: ___________________________
Stop path: _____________________  Rollback path: _______________
```

Pave it for yourselves first

The platform or enablement team should be the first production user of the path it intends to offer. This is not a demonstration. It is a way to discover the missing defaults, documentation, telemetry, and recovery steps before another team inherits them. Complete this mini-charter separately from the product pilot charter.

```
PLATFORM PILOT ZERO MINI-CHARTER
Path owner: ____________________________________________________
Internal work unit used first: _________________________________
Internal-service outcome and baseline: ________________________
Boundary and exclusions: ______________________________________
Risk class and human approval point: __________________________
What must be learned before inviting another team: _____________
First outside customer and why: ________________________________
What remains local until a second use produces reuse evidence: _
Evidence required before promoting a local pattern: ____________
Stop condition and verified fallback: _________________________
Recovery drill and evidence: __________________________________
Decision owner / gate date: ___________________________________
```

The product pilot may be designed while Platform Pilot Zero runs. It does not begin production use until the internal journey, stop exercise, verified fallback, the rollback or disablement drill required by the control profile, and restart evidence have been reviewed at the phase-zero gate.

Exclusion register

The future-state map is incomplete until it says what will not enter the path. Each exclusion carries its retained controls and a review trigger. "Legacy" is not a reason by itself, and an exclusion without an owner becomes an invisible exception.

Field	Entry
System or work class	What is outside the supported path?
Reason and evidence	Why is inclusion currently worse than exclusion?
Retained process and controls	What governs the work instead?
Owner	Who remains accountable?
Risk acceptance or approval	Which authority accepted the boundary?
Review trigger and date	What change forces reconsideration?
Status	Proposed, approved, expired, or retired

Completed Harborline entry. The vendor-managed settlement appliance, receiving two scheduled configuration changes a year, remains outside the generated-change path. The named settlement owner keeps the vendor change procedure, dual human approval, reconciliation, and the recovery drill. Risk approves the exclusion for six months. A vendor API release, a change-volume increase, a control failure, or a missed review date reopens the decision. This entry does not lower the appliance's controls. It makes the model's edge visible.

Exclusions are not forever. They prevent ambition from silently widening the pilot after the controls and capacity were designed.

Harborline illustration

Platform Pilot Zero: the owner's journey

Harborline's platform group chooses one repeated service-account request class. Today, engineers copy an old ticket, identity operations corrects permission errors, and someone reconstructs evidence after approval. The group creates one versioned request spine, pre-approved permission boundary, registered workflow, policy test, human identity approval, expiry record, and production audit link. It runs a real request through that path before inviting the payments team.

```
Internal outcome: reduce service-account request time without weakening
                  permissions or traceability
Boundary: one repeated non-production service-account class
Excluded: production roles, emergency access, and novel permission patterns
Stop condition: missing approval, expiry, or provenance evidence
Fallback: disable the paved route and use the verified manual request path
Exercise: complete one journey, trigger the stop, reconcile the record, and
          approve restart evidence
Promotion condition: completed loop and recovery drill, with every workaround
                     recorded in the paved-path backlog
First outside customer: payments product team
```

This is the prerequisite Harborline uses to decide whether the shared path is ready to host product work. It is not evidence that the payments hypothesis will succeed.

Product pilot: payment retry consolidation

Current state: three teams maintain retry behavior in separate services. Policy is partly in tickets, partly in code, and partly in the memory of a ledger reviewer. A payment engineer proposes a change, waits for ledger review, and discovers idempotency constraints after implementation. Gateway evidence and finance outcome data live in separate systems. Review and rework are the bottleneck, not typing the retry loop.

Future state: one versioned payment-retry spine names the business outcome, duplicate-entry invariant, interfaces, and non-goals. The platform-owned path generates against inherited payments standards, runs contract and policy tests, records provenance, and requires the named ledger evaluator before the highest risk changes proceed. Gateway and finance signals return to the same scorecard. The payments product team is the first outside adopter after Platform Pilot Zero closes.

Design choice	Harborline answer
Work removed	Repeated local interpretation of the retry policy
Human decision preserved	Whether the duplicate-entry invariant is satisfied
Agent action bounded	Generate and test implementation inside the approved spine
New evidence	Spine version, standards version, tests, reviewer, release, outcome
Excluded in pilot	Message-schema redesign, settlement, customer communications
Stop path	Disable the agent-assisted route and block new generated changes
Rollback path	Restore the last verified retry policy through the same pipeline

Gate 2: redesign

```
[ ] Current state includes real waits, rework, exceptions, and hidden labor.
[ ] Future state removes or moves work before automating it.
[ ] Human decisions and agent actions are explicit.
[ ] Intent, standards, gates, evaluation, and feedback form one loop.
[ ] Platform Pilot Zero and the first outside customer have separate charters.
[ ] Pilot zero has a completed loop, tested stop, verified fallback, the
    profile-required rollback or disablement drill, and restart evidence.
[ ] Exclusions, stop path, and rollback path are visible.

Decision: [continue] [revise the flow] [choose different work]
Unresolved design risk: _____________________________________________
Owner and due date: _____________________________________________
```

Only now is the room ready to ask what operating system, capabilities, and organization the redesigned work needs.

Design the Operating Model

Turn the future-state flow into explicit capabilities, ownership, interfaces, and one testable organization hypothesis.

> *A workflow becomes an operating model when the room can say who owns each capability, how decisions travel, and what evidence changes the design.*

Timebox: 45 minutes

Decision: whether the operating model is complete enough to test without a premature reorganization

Outputs: Operating Model Canvas, organization-design hypothesis, decision record

Use Chapters 11 and 12 of the main book, plus the Operating Model Canvas, Platform Capability Charter, operating ownership matrix, and Organization Design Decision Record in the Operating Kit.

Operating Model Canvas

Complete this from the future-state work redesign. Short answers expose gaps. Long answers often hide them.

```
OUTCOME AND BOUNDARY
Outcome: _____
Work unit: _____
Pilot boundary and non-goals: _____

FLOW AND ARTIFACTS
Entry condition: _____
Versioned intent artifact: _____
Paved path: _____
Evaluation evidence: _____
Production feedback: _____

DECISIONS AND OWNERSHIP
Outcome owner: _____
Operating owner: _____
Decisions humans retain: _____
Actions agents may take: _____
Decision policy owner: _____
Agent / workflow owner and backup: _____

CAPABILITIES AND INTERFACES
Team-owned capabilities: _____
Platform-owned capabilities: _____
Security, risk, privacy, or legal interface: _____
People and finance interface: _____
External vendor dependency: _____

LEARNING AND CHANGE
Signals reviewed weekly: _____
Decision made at the first gate: _____
Design assumption most likely to be wrong: _____
What would cause the model to change: _____
```

Platform Capability Charter

Create a charter for each shared capability the pilot depends on. One charter is enough for the workshop if it covers the most important dependency.

```
Capability: _____
Customer and job to be done: _____
Owner and funded team: _____
Service offered: _____
Supported work and explicit exclusions: _____
Interface and adoption path: _____
Service level or reliability expectation: _____
Required controls and evidence: _____
Telemetry and customer feedback: _____
First internal use: _____
First external adopter: _____
Promotion criteria from local pattern to paved path: _____
Review date: _____
```

The charter prevents a platform group from becoming a collection of experts waiting for tickets. It states a customer-facing capability, an adoption path, and the evidence that will improve it.

Operating ownership matrix

The four operating responsibilities sit inside this wider matrix. Use names, not committee labels. **A** means the one accountable decision owner. **R** means the people doing the work. A row may have several responsible or consulted parties, but exactly one accountable owner.

Artifact or decision	One accountable owner	Responsible	Consulted	Informed
Value hypothesis and business measure	Product or outcome owner	Product and finance partner	Design, engineering, risk	Sponsor and platform
Spine ready for cross-functional review	**Spec owner**	Product, design, and engineering authors	Evaluator and risk by tier	Pipeline architect
Technical constraints and interfaces	Engineering owner	Engineering and architecture	Product, design, security, platform	Spec owner
Standard created, inherited, or promoted	**Standards keeper**	Domain experts and platform contributors	Security, risk, evaluator	Affected teams
Agent configuration and paved path	**Pipeline architect**	Platform engineering	Security, risk, standards keeper, evaluator	Product teams
Risk classification and control profile	Security or risk owner	Risk and control specialists	Product, engineering, pipeline architect	Sponsor
Verification decision	**Evaluator**	Qualified reviewers and automated harness	Spec owner, security or risk by tier	Engineering owner
Agent registration or retirement	Operating owner	Agent owner and platform	Security, risk, pipeline architect	Affected teams
Cohort staffing and protected capacity	Engineering director	Managers and role holders	Product, people leader, operating owner	Sponsor
Organization design hypothesis	Engineering executive	People leader and operating owner	Product, platform, security or risk	Affected organization

Artifact or decision	One accountable owner	Responsible	Consulted	Informed
Phase-gate decision	Executive sponsor	Operating owner assembles evidence	Outcome owner, evaluator, risk, people leader	Cohort and governance forum
Workforce commitments and manager conversation	People leader	Managers	Engineering executive and affected people	Sponsor

Adapt the titles to local accountability. Do not let "shared ownership" create a row with two accountable parties or no name at all.

Organization-design hypothesis

Do not redraw the permanent org chart. State the smallest structural hypothesis the pilot can test.

```
We believe the redesigned work needs _____________________________
owned by _____________________________ and connected to ___________
through ____________________________________________________ .

For the pilot we will test this through [protected hats] [temporary pod]
[platform service] [community of practice] [other: ___________].

We will keep the arrangement if _____________________________________
and change it if ____________________________________________ .
Decision date: _____________________  Decision owner: ___________
```

Organization Design Decision Record

```
ODDR ID / date: ____________________________________________________
Decision: __________________________________________________________
Context and work evidence: _________________________________________
Options considered: ________________________________________________
Choice and why: ____________________________________________________
Capabilities affected: _____________________________________________
People affected and how they will be involved: ____________________
Risks and reversible parts: ________________________________________
Evidence that will confirm or challenge the decision: ____________
Review owner and date: _____________________________________________
Supersedes / superseded by: ________________________________________
```

Version this record beside the pilot artifacts. The point is not bureaucracy. It is to make structural choices reviewable when the evidence changes.

Harborline illustration

Canvas field	Harborline answer
Outcome	Reduce failed-retry cost without increasing duplicate ledger entries
Work unit	One change to retry policy, from approved spine through outcome review
Operating owner	Payments engineering director
Outcome owner	Finance operations lead
Team capability	Payment-domain intent, implementation, evaluation, and on-call
Platform capability	Versioned spine template, approved agent path, gates, provenance, telemetry
Control interface	Ledger risk partner approves the invariant and highest-tier decision rights
Learning loop	Weekly delivery health, monthly business value, incident-linked standards updates
Org hypothesis	Protected roles in the payments team plus one platform capability pod are sufficient for the pilot
Change trigger	Review demand exceeds staffed capacity or control decisions queue outside the team

Completed Harborline capability charter

```
Capability: governed agentic delivery path
Customer and job: platform and product teams move a bounded work unit from an
                  approved spine to verified release with one evidence chain
Accountable owner: platform engineering director
Funded capability: protected platform owner and backup, pipeline architect,
                   evaluator support, and a versioned backlog
Service: spine template, registered specialist configuration, policy and
         evaluation gates, provenance, telemetry, support, and recovery drill
Boundary: Platform Pilot Zero plus the named payments cohort
Exclusions: emergency access, novel production permissions, settlement,
            message-schema redesign, and customer communications
First internal use: repeated non-production service-account request class
First outside adopter: payments product team's retry-policy changes
Promotion criteria: Platform Pilot Zero loop and recovery drill complete; outside cohort
                    meets health and control thresholds; capacity stays inside
                    its approved maximum; workarounds become backlog items
Review: phase-zero gate, then monthly operating review
```

Completed Harborline organization-design decision record

```
ODDR: HL-OD-001 / illustrative workshop date
Decision: test protected operating hats in the payments team plus one funded
          platform capability; do not create a permanent AI department
Context: retry behavior is fragmented; evaluation is specialized; the shared
         path needs an accountable owner before a product cohort can use it
Options: unfunded council; centralized delivery team; protected team hats plus
         platform capability; immediate permanent reorganization
Choice: protected hats plus capability, because it preserves product ownership,
        funds the shared path, and remains reversible at the gate
People involved: current role holders review load, authority, learning, and
                 mobility effects with their managers before day 1
Challenge evidence: evaluator utilization above 85%; unresolved control queues;
                    unsupported workarounds; failure of the path to travel
Review owner: payments engineering VP at the day-90 gate
Supersedes: none
```

Harborline does not create an enterprise AI department. It charters one shared capability, protects named hats in the delivery team, and records what the pilot must prove before those arrangements become permanent.

How this collapses in a smaller organization

The capabilities do not disappear when the company is small. The boxes do.

Larger organization	Smaller organization
Platform capability team	One platform owner with a funded backlog
Separate spec owner and product lead	Product or engineering lead holds both hats, with explicit review
Evaluator pool by domain	Named senior reviewer and backup with protected hours
Dedicated standards keeper	Rotating responsibility tied to incidents and retrospectives
Risk, privacy, legal, and security functions	Named external or fractional advisers plus one accountable internal owner
Formal monthly operating review	One recurring leadership decision meeting using the same scorecard

Combining hats is honest. Pretending the work no longer needs an owner is not.

Gate 3: make the model testable

```
[ ] Every capability in the future-state flow has an owner.
[ ] Team and platform boundaries describe services and decisions, not slogans.
[ ] Cross-functional control, people, finance, and vendor interfaces are named.
[ ] The first-use and adoption path are explicit.
[ ] The organization hypothesis is reversible and has a decision date.
[ ] Structural choices are captured in an ODDR.

Decision: [continue] [revise ownership] [narrow the pilot]
Most important assumption under test: _____________________________
Owner: _____________________________  Review date: _____________
```

The next session turns capabilities into real capacity. Titles alone cannot run the model.

Staff the Capabilities

Turn new responsibilities into named authority, protected time, review capacity, and a visible path for people to grow.

> *An operating model written in nouns can sound complete while every load-bearing responsibility remains somebody's side job. This session puts names, hours, and career consequences on the design.*

Timebox: 60 minutes

Decision: whether the pilot has enough qualified capacity and clear authority to run safely

Outputs: capability plan, role map, evaluator-capacity plan, people actions

Use Chapter 13, *Roles, Careers, and Apprenticeship*, for the argument behind this session. This session owns the working forms.

Capability plan

Start with the future-state work, not existing titles.

Capability required	Current strength	Gap	Build, borrow, buy, or narrow	Owner	Evidence of readiness
Intent and specification					
Domain evaluation					
Standards stewardship					
Pipeline and agent path					
Outcome measurement					

Capability required	Current strength	Gap	Build, borrow, buy, or narrow	Owner	Evidence of readiness
Security, risk, privacy, legal					
People transition and apprenticeship					

Do not use hiring as the default answer. Some gaps require protected time, training, a narrower pilot, or a clearer interface with an existing function.

Four operating role cards

The four load-bearing responsibilities may be full-time roles or protected hats. They do not replace product, design, architecture, security, risk, or engineering ownership; the operating ownership matrix in the org-model session shows how the whole group fits together.

Role	Accountable for	May decide	Must not become	Capacity signal
Spec owner	Spine quality, versions, open questions, and non-goals	Whether the spine is ready for cross-functional review	A proxy product owner inventing business intent alone	Specification-gap findings and time waiting for answers
Evaluator	Evidence-based judgment against spine and standards	Accept, reject, or return work inside assigned domain authority	A central queue that relieves teams of quality ownership	Queue time, active review time, re-review demand, and structural rework
Standards keeper	Current team or division rulebook and promotion proposals	Local rule changes inside the approved layer	A style police function or undocumented architect veto	Age of incident-linked updates and repeated defect classes
Pipeline architect	Gate integrity, speed, agent configuration, telemetry, and break-glass path	Technical implementation of approved controls	An approval board for product decisions	Gate duration, bypass frequency, and failed control checks

For each local card, complete the blank form. A title without protected time is not a role. It is a wish.

```
Role or responsibility: _____
Purpose: _____
Primary owner / backup: _____
Inputs consumed: _____
Outputs produced: _____
Required judgment or authority: _____
Decisions and limits: _____
Success measures: _____
Protected hours or percentage: _____
Assignment review date: _____
```

Role and authority map

The four roles sit inside a wider operating group. Record people, authority, protected time, and backup for every load-bearing responsibility. A title without those fields is not staffed.

Responsibility	Named person	Authority or decision boundary	Protected hours per week	Backup
Outcome and value				
Product or domain intent				
Spec ownership				
Evaluation				
Standards stewardship				
Pipeline and agent path				

Responsibility	Named person	Authority or decision boundary	Protected hours per week	Backup
Engineering delivery and production				
Platform capability				
Security, risk, privacy, or legal				
People and career actions				
Financial validation				

Prompt: *Which row will quietly fall back to the most experienced engineer when the pilot gets busy?* Fix that row before moving on.

Evaluator-capacity worksheet

Calculate by domain and risk class. Initial review is the base term. Re-review is a multiplier on that base. A work unit with zero re-review rounds still consumes its first review.

```
DOMAIN / RISK CLASS: _____________________________________________
Expected work units per week                                U = ____
First-review active hours per unit                          F = ____
Expected re-review rounds per entering unit                 R = ____
Re-review effort as a share of first review                 M = 0.__

Required review hours = U x F x [1 + (R x M)]              D = ____

Named evaluator protected hours                             H = ____
Coverage factor for leave, incidents, and mismatch          C = 0.__
Usable capacity = H x C                                     K = ____
Utilization = D / K                                         Z = ____%

Local decision threshold and direction: Z must be [below] ____%
Decision: [staffed] [narrow intake] [add qualified capacity] [hold]
```

Repeat the calculation for every domain or risk class, then add required hours and usable hours only after checking that evaluators are actually interchangeable. Forty available security-review hours cannot cover a forty-hour ledger-review gap if the reviewers lack ledger authority.

Domain / risk	Inputs: U, F, R, M	Demand	Usable capacity	Utilization and action
Total, only where fungible				

Apprenticeship and career actions

The pilot changes learning as well as delivery. Do not wait for a future talent program to address it.

```
Judgment novices should observe: _____________________________
Artifact they will help author: _____________________________
Evaluation task they can practice safely: ___________________
Paired review or calibration cadence: ______________________
Named mentor and protected hours: ___________________________
Evidence of growing judgment: ______________________________

Work likely to thin: ________________________________________
Work likely to grow: ________________________________________
People affected before the first gate: ______________________
Manager conversation owner and date: _______________________
Training, mobility, or transition commitment: _______________
```

The facilitator does not accept “communicate later.” If the pilot changes the work, the people action belongs in the pilot plan now.

Manager conversation guide

Managers repeat this conversation at the start of a pilot and at every phase gate.

```
What the evidence says is changing in our work:
_____________________________________________________________

Work that is becoming cheaper or needs fewer hands:
_____________________________________________________________

Work that is growing, and the skills it requires:
_____________________________________________________________

The internal path, training time, and compensation treatment:
_____________________________________________________________

What we do not know yet, and when we will know more:
_____________________________________________________________

Who can challenge this plan or ask for an individual conversation:
_____________________________________________________________

Challenge forum, owner, and response deadline:
_____________________________________________________________

Unresolved objections and the leadership response, recorded in the gate record:
_____________________________________________________________

Commitments made, owner, and review date:
_____________________________________________________________
```

Do not turn this into reassurance copy. The point is to make the real bargain specific enough that the people living it can question it, plan around it, and see whether leadership kept it.

Harborline illustration

Harborline assigns protected hats in the payments group rather than creating new titles for the pilot.

Responsibility	Harborline assignment	Protected capacity or authority
Outcome owner	Finance operations lead	Judges monthly value result
Spec owner	Payments product lead	8 hours weekly; approves spine readiness
Evaluator	Ledger senior engineer plus payments principal	40 usable hours combined after coverage factor
Standards keeper	Payments staff engineer	4 hours weekly; incident-linked updates
Pipeline architect	Platform engineer	12 hours weekly; owns approved path and telemetry
Risk partner	Payments control lead	Approves classification, exceptions, and gate evidence

Responsibility	Harborline assignment	Protected capacity or authority
People owner	Engineering director	Runs role and mobility conversations before scale

Capacity illustration

These are the authoritative Harborline capacity values.

Domain / risk	Base demand	Re-review factor	Required	Usable qualified hours	Segment utilization
Payment behavior, Tier 1	6 x 2.5h = 15.0h	1 + (0.4 x 0.5)	18.0h	18.0h allocated	100%
Ledger interfaces, Tier 1	4 x 2.0h = 8.0h	1 + (0.5 x 0.5)	10.0h	12.0h allocated	83.3%
Notification support, Tier 2	8 x 0.8h = 6.4h	1 + (0.5 x 0.5)	8.0h	10.0h allocated	80%
De-duplicated total	**29.4h base**		**36.0h required**	**40.0h usable after coverage**	**90%**

Harborline allocates each usable hour to one segment, so cross-qualified time is not counted twice. Payment behavior is already at full segment utilization, and the de-duplicated portfolio total is 90 percent, above Harborline's illustrative maximum of 85 percent. Harborline holds new cohort intake until it staffs qualified backup capacity or demand falls. Capacity is not fungible merely because the spreadsheet can add it.

Gate 4: staff

```
[ ] Every load-bearing capability has a named owner and backup.
[ ] Decision authority is explicit, not implied by title.
[ ] Protected time is visible in the staffing plan.
[ ] Evaluation demand includes initial review and expected re-review.
[ ] Demand is segmented where expertise or authority is not fungible.
[ ] Apprenticeship and manager actions are funded inside the pilot.

Decision: [staffed] [narrow scope] [add capacity] [hold]
Capacity action: _____________________________________________
People action: _____________________________________________
Owners and dates: _____________________________________________
```

With the people and capacity visible, the room can design controls that the operating model is actually able to perform.

PART III: MAKE IT GOVERNABLE

Set the Risk and Control Profile

Classify the work, make decision rights executable, and prove that the pilot can stop and recover before it starts.

> *Governance becomes useful when it changes the path before release, not when it explains the path after an incident.*

Timebox: 45 minutes

Decision: whether the pilot's decision rights, evidence, and recovery paths match its actual risk

Outputs: risk classification, control profile, agent record, exception path, stop and rollback test

Use Chapter 14, *Governance That Ships*, for the argument behind this session. This session owns the working forms. Local legal, regulatory, privacy, security, model-risk, and employment decisions remain with the functions authorized to make them.

Risk-classification worksheet

Classification happens before tool selection and before generation. Score the work against the five dimensions; the highest applicable condition sets the provisional tier, and the named risk owner confirms it under local policy. The rubric supports that conversation. It does not replace counsel, security, privacy, safety, or regulated-model classification where those authorities apply.

Dimension	Tier 3: limited impact	Tier 2: material impact	Tier 1: high or regulated impact
Consequence	Internal inconvenience; easy correction	Customer or operational disruption	Money, identity, health, safety, or regulated obligation
Data	Public or low-sensitivity internal data	Confidential business or limited personal data	Restricted, financial, health, identity, or regulated data
Change class	Docs, tests, or isolated internal tooling	Customer-facing behavior without high-impact transaction	Ledger, authorization, entitlement, safety control, or regulated decision
Reversibility	Immediate and low-cost	Tested rollback with bounded state repair	Irreversible effect, state reconciliation, or material customer remediation

Dimension	Tier 3: limited impact	Tier 2: material impact	Tier 1: high or regulated impact
Blast radius	One person or isolated environment	One service or bounded customer group	Shared platform, many customers, systemic dependency, or external reporting

Dimension	Observed condition and evidence	Provisional tier	Authority or obligation triggered
Consequence			
Data			
Change class			
Reversibility			
Blast radius			

```
Locally assigned risk tier / class: _____________________________
Policy owner approving the classification: _____________________
Obligations that must become tests or gates: ___________________
Actions that always require a named human: _____________________
Actions excluded from the pilot: _______________________________
Reclassification trigger: ______________________________________
```

Do not average the dimensions. One Tier 1 condition may determine the class even when every other dimension is lower. Record autonomy, tool access, vendor dependency, and concentration in the control profile even when local policy does not use them to set the tier.

Control-profile worksheet

```
Work class and provisional tier: _____
Triggered dimensions: _____
Obligations: _____
Allowed autonomy and tool access: _____
Human approval point: _____
Required evaluations and evidence: _____
Data and vendor restrictions: _____
Release path: _____
Stop condition: _____
Rollback or disable path: _____
Owner and next review date: _____
```

Control layer	Required control	Evidence produced	Owner	Failure action
Intent				
Standards and policy				
Agent configuration and tool access				
Test and evaluation harness				
Human approval				
Release and provenance				
Runtime monitoring				

Control layer	Required control	Evidence produced	Owner	Failure action
Break-glass				
Stop and rollback				
Periodic review				

Each evidence item needs a source, retention rule, and person authorized to interpret it. “Logged” is not a control if nobody can retrieve or read the log.

These starting points calibrate the profile by tier:

Control	Tier 3 starting point	Tier 2 starting point	Tier 1 starting point
Human decision	Named owner for production effect	Human confirmation before material release	Named human owns verification, merge, and production decision
Evaluation	Standard harness	Domain harness plus sampled adversarial cases	Domain and obligation-specific harness; independent qualified judgment
Provenance	Tool, config, and approver recorded	Complete release evidence required	Complete evidence required; missing evidence stops the path
Recovery	Disable or revert documented	Recovery tested inside approved window	Recovery and state-reconciliation path tested; stop authority named
Exceptions	Owner and expiry	Approver, compensating control, and expiry	Risk approval, compensating control, short expiry, and gate visibility

Decision-rights worksheet

List actions, not broad job descriptions. This illustrative starting table is not permission to copy its autonomy levels into production: **H** means a human must own the action, **G** means an agent may act inside declared guardrails, and **C** means an agent recommends and a named human confirms.

Action	Tier 3: limited impact	Tier 2: material impact	Tier 1: high or regulated impact
Open pull request	G	G	G
Merge to main	G	C	H
Generate and run tests	G	G	G
Modify CI or gate configuration	C	H	H
Apply schema migration	C	C	H
Deploy to production	G	C	H
Send customer communication	C	H	H
Change spine or standards file	C	C	H
Rotate or access secrets	H	H	H

"G" never means unconstrained. It means the action is allowed only while its tests, environment restrictions, spend limits, telemetry, and rollback conditions hold. Record the approved local version below and publish it as configuration the pipeline reads.

Action	Agent may act inside guardrails	Agent recommends, human confirms	Human must own	Preconditions and evidence
Open a pull request				
Change the intent or standards artifact				
Modify pipeline or gate configuration				

Action	Agent may act inside guardrails	Agent recommends, human confirms	Human must own	Preconditions and evidence
Merge				
Deploy				
Apply schema or data change				
Send customer communication				
Access or rotate secrets				
Invoke break-glass				
Other local action: _____				

Store the approved table where the path can enforce it. A policy remembered by a reviewer is a queue, not an executable control.

Agent register

One record per agent or agentic workflow in the operating environment. Registration should happen as a side effect of the paved path, not through a quarterly memory exercise. Capture at least these fields in the room so no pilot begins anonymously.

```
Agent or workflow ID: _____
Purpose and boundary: _____
Named owner / backup: _____
Systems and data touched: _____
Provider, model, and configuration version: _____
Allowed actions and tool access: _____
Required human approval: _____
Telemetry and evaluation evidence: _____
Kill or stop mechanism, owner, last test: _____
Review and expiry dates: _____
```

In production, the register carries the full schema. Harborline's completed entry shows the shape:

```
agent_id: payments-retry-reviewer-v3
owner: akumar
backup_owner: lchen
purpose: findings-only review of payment-retry changes against spine and standards
touches: [payments-core, ledger-api]
environment: production-delivery
risk_tier: 1
data_classification: confidential
provider: approved-provider-alias
model: approved-model-alias
autonomy: findings-only
human_approval: required-before-merge
spec: repo://agents/payments-retry-reviewer/spine.md
config_version: 3.4
guardrails: read-only tools; no secrets in context; findings-only output
telemetry: dashboard://agents/payments-retry-reviewer
evals: ci://agent-harness/payments-retry-reviewer
kill_switch: gateway policy revocation, tested 2026-06-12
last_reviewed: 2026-06-30
last_eval_passed: 2026-06-28
next_review_due: 2026-09-30
exception_expiry: null
incident_contact: payments-delivery-oncall
```

In the Harborline packet, the readiness inventory finds `retry-helper-local-v1`, a temporary script with no owner or stop path. Harborline does not bless it by adding a row. The operating owner first decides whether to retire it or bring it through specification, evaluation, registration, and the tier-one control profile. Inventory is the start of governance, not the end.

Exception register

Exception	Reason	Approver	Compensating control	Owner	Expiry	Close or renew evidence

Exception	Reason	Approver	Compensating control	Owner	Expiry	Close or renew evidence

No standing exception enters the pilot unnamed or without an expiry.

Release gate checklist

Wire the checklist into the pipeline and produce evidence automatically where the control can be automated. A failed gate means fix the artifact or implementation, or use the approved exception path. It never means quietly weaken the control.

```
[ ] Spine referenced by this release matches the approved version
[ ] Contract tests pass against declared interfaces
[ ] Policy tests preserve required constraints and obligations
[ ] Audit events match the approved event shape
[ ] Accessibility checks pass for touched surfaces
[ ] Provenance records approver, tools, spine, standards, and config versions
[ ] Dependency and vulnerability checks pass under the approved policy
[ ] Rollback or disable evidence for this change class is current
[ ] Any exception records owner, approver, reason, expiry, and compensating control
```

Stop is not rollback

The room must be able to perform both.

- **Stop** prevents the agent or path from taking more action.
- **Rollback** restores a previously verified product or configuration state.
- **Break-glass** allows an authorized human to bypass the normal path under logged, time-bounded conditions.

```
STOP
Trigger: _____________________________________________________
Mechanism: ___________________________________________________
Authorized operator and backup: ______________________________
Maximum time to stop: ________________________________________
Evidence preserved: __________________________________________
Last tested / next test: _____________________________________

ROLLBACK
Trigger: _____________________________________________________
Last verified state: _________________________________________
Mechanism and pipeline: ______________________________________
Data repair or reconciliation: _______________________________
Decision owner: ______________________________________________
Recovery evidence: ___________________________________________
Last tested / next test: _____________________________________
```

Intent-chain post-mortem

Use this for an escaped defect or major structural rework. Walk every link. The defect often lives in the one the meeting wants to skip.

```
Defect summary: what shipped, what it did, and who it reached.

1. SPEC. Did the spine name the scenario that failed?            [yes/no + evidence]
2. STANDARDS. Did a rule exist that would have prevented it?      [yes/no + evidence]
3. PROMPT/CONFIG. Did the generation frame carry the constraint?  [yes/no + evidence]
4. ARTIFACT. Does the implementation conform to the spine?        [yes/no + evidence]
5. HARNESS. Did tests exercise the actual failure mode?           [yes/no + evidence]
6. RUNTIME. Did production evidence reveal a condition the harness missed?
                                                                  [yes/no + evidence]

Primary classification:
[ ] Specification gap       [ ] Plausible refactor
[ ] Softened invariant      [ ] Polished hallucination
[ ] Evaluation gap          [ ] Reviewer absence

Upstream finding owner, not the ticket engineer: _______________
Class-level fix in template, standard, harness, config, or capacity:
_____________________________________________________________
Decision and completion evidence: ___________________________
```

Harborline illustration

Harborline classifies payment-retry changes in its highest local risk tier because the path can affect money movement and duplicate ledger entries. The agent may draft implementation and run tests. A named human owns changes to the spine, standards, gate configuration, merge, and production release.

Control layer	Harborline decision
Intent	Approved spine names the duplicate-entry invariant and non-goals
Harness	Contract, concurrency, idempotency, and policy tests gate the path
Approval	Ledger evaluator and release owner are recorded
Provenance	Spine, standards, configuration, tests, approvers, and release are linked
Runtime	Duplicate-entry and failed-retry signals route to named owners
Exclusion	Schema, settlement, and customer-communication changes stay outside the pilot

Completed Harborline gateway-alias stop record

This is the authoritative record of the stop branch the book's roadmap chapter narrates. Earlier in the quarter, a gateway mapping error leaves the required provider and model alias unresolved for two generated changes. Tier 1 provenance falls below its 100 percent requirement, so this is a stop rather than a hold.

```
Path decision: STOP new generated retry work on the affected configuration
Fallback: pre-pilot reviewed delivery workflow
Immediate action: disable the configuration and quarantine unmerged work
Released-change assessment: behavior, tests, and approvals hold for the two
                            released changes
Recovery disposition: NO PRODUCTION ROLLBACK REQUIRED
Repair: correct the gateway mapping, reconcile records, and add a provenance
        regression test
Restart evidence: complete records plus control-owner review
Restart authority: executive sponsor on control-owner recommendation
```

Separate illustrative configuration-recovery drill

The workshop also practices a different scenario. After reviewer configuration v4 is released, two Tier 1 records omit the required standards-version field. The pipeline architect stops new Tier 1 releases on that route and returns work to the pre-pilot reviewed fallback. The operating owner directs configuration rollback to evaluated v3; production software rolls back only if re-evaluation finds impact. The two affected changes are quarantined for evidence reconstruction and qualified re-evaluation. The path may restart only after v3's evaluation pass is confirmed, missing evidence is reconciled, the regression test covers the provenance field, and the risk owner accepts the recovery record.

This drill is illustrative, not another reported Harborline event and not a fourth gate outcome. A duplicate-entry invariant failure would likewise stop the path, with product rollback and ledger reconciliation determined under the applicable incident policy and named authorities.

Gate 5: control

```
[ ] Risk class is approved by the authorized local owner.
[ ] Obligations appear in intent, standards, tests, gates, or human decisions.
[ ] Decision rights are action-specific and enforceable.
[ ] Every agent or workflow has an owner, evidence, and stop mechanism.
[ ] Exceptions have approvers, compensating controls, and expiry dates.
[ ] Stop, rollback, and break-glass are distinct and testable.

Decision: [control profile accepted] [revise] [narrow] [stop]
Required control action: _____________________________________
Owner: _____________________________  Due: _____________________
```

The path is now governable on paper. The next session defines the evidence that will show whether it works in practice.

Define the Evidence

Build one metrics dictionary, three useful views, and thresholds that force a decision instead of decorating a dashboard.

> *The workshop has made a series of hypotheses. This session decides what evidence could prove them wrong.*

Timebox: 45 minutes

Decision: whether the pilot can produce credible evidence for health, value, control, and people outcomes

Outputs: metrics dictionary, scorecard, ledger sample, Platform Adoption Shadow entries, threshold actions

Use Chapter 15, *The Scorecard: Measuring Evaluation Capacity, Not Activity*, for the argument behind this session. This session owns the working forms.

Begin with decisions, not available data

Write the decisions the scorecard must support.

```
Weekly operating decision: _____________________________________
Monthly outcome decision: _____________________________________
Phase-gate decision: _____________________________________
People or capacity decision: _____________________________________
```

If a measure cannot change one of those decisions, remove it from the headline scorecard. It may still be useful diagnostic evidence.

Metrics dictionary

Create one record per measure before building the view.

```
Metric ID and name: _____
Decision supported: _____
Definition and formula: _____
Unit of work / unit of measure: _____
Population and segmentation: _____
Source events and systems: _____
Exclusions: _____
Baseline and baseline window: _____
Desired direction: [higher] [lower] [inside range] [none]
Decision threshold and direction: _____
Evidence window and minimum sample: _____
Data owner / decision owner: _____
Validation or reconciliation method: _____
Likely gaming or failure mode: _____
Last definition change / approver: _____
```

Use **decision threshold** plus a direction. Do not use *floor* for both a minimum quality measure and a maximum utilization measure.

The canonical program dictionary begins with these measures. Local teams may add measures, but they do not silently redefine these IDs.

Metric ID	Definition	Direction	Decision threshold and action	Owner and source
adoption.gate_qualified_teams	Teams active after readiness and charter approval	Context, not a success target	No unqualified intake; remove or remediate any unqualified team	Program lead; cohort register
adoption.paved_path_coverage	In-scope releases using the approved path or a recorded exception divided by all in-scope releases	Higher	Meet locally approved minimum; hold expansion and reconcile missing work	Platform owner; release events and exception register
health.first_pass_alignment	Applicable declared criteria and constraints satisfied at first structured review	Higher	Improve from approved baseline after minimum sample; hold and analyze if missed	Evaluation lead; evaluation record

Metric ID	Definition	Direction	Decision threshold and action	Owner and source
health.structural_rework_rate	Reviewed work units needing material logic, interface, data, security, or integration change before verification	Lower	No worse than approved baseline; hold and classify misses if exceeded	Engineering lead; review record
health.review_latency	Median and slow-tail elapsed time from review-ready to decision	Lower inside service range	Both statistics inside charter service level; narrow intake or add capacity if missed	Evaluation lead; workflow events
health.defect_class_frequency	Escaped defects and major rework classified by the intent-chain taxonomy	Fewer repeated severe classes	No unresolved severe class and a named class-level repair for every material finding	Engineering owner; incident and review records
economics.cost_per_verified_outcome	Fully loaded delivery and evaluation cost per verified outcome under the declared unit	Lower, or justified by greater value	No increase without approved value gain	Outcome owner and finance partner; cost record
economics.declared_value	Business measure named in the pilot charter	Charter direction	Move in declared direction after evidence window; stop or redesign if hypothesis fails	Outcome owner; system of record

Metric ID	Definition	Direction	Decision threshold and action	Owner and source
risk.provenance_completeness	Releases with every required evidence field divided by in-scope releases	Higher	Equal to locally approved requirement; missing Tier 1 evidence stops the affected path	Risk owner; release evidence
risk.break_glass_review_completion	Break-glass events reconciled and reviewed inside the required window divided by all break-glass events	Higher	Equal to approved requirement; overdue access or reconciliation stops expansion	Control owner; access and incident records
people.evaluator_utilization	Required qualified review hours divided by usable qualified hours, segmented by domain and risk	Lower	At most charter percentage; hold intake when exceeded	Engineering director; capacity worksheet
people.learning_and_mobility	Funded apprenticeship assignments, role conversions, and mobility commitments completed against plan	Meet declared plan	Hold scale when the people commitments required by the gate are unfunded or overdue	People leader; staffing and learning records

One-page pilot scorecard

Keep the categories stable. Local measures can change through the dictionary.

Category	Measure	Baseline / current	Decision threshold and direction	Window / sample	Owner / action
Adoption	Gate-qualified work using the path				
Adoption	Paved-path coverage or recorded exception				
Health	First-pass alignment				
Health	Structural rework				
Health	Review latency, median and slow tail				
Health	Material defect-class findings repaired				
Economics	Cost per verified outcome				
Economics	Pilot business result				
Risk	Complete required provenance				
Risk	Break-glass review completion				

Category	Measure	Baseline / current	Decision threshold and direction	Window / sample	Owner / action
People	Evaluator demand / usable capacity				
People	Funded role and apprenticeship commitments				

The adoption row counts only work that passes the agreed gate. License activation, generated volume, tokens, and suggestion acceptance are inputs or telemetry, not evidence of a better operating model.

Three views, one dictionary

Do not create three definitions of the same measure.

View	Audience and cadence	Questions it answers
Weekly operating view	Pilot team and operating owner	Is the path healthy? Where is work waiting? Must intake change now?
Monthly outcome view	Sponsor, outcome owner, finance, control, people	Is value appearing? Are controls intact? What is changing for people?
Phase-gate view	Named decision owner and cross-functional approvers	Advance, hold, or stop? What fallback and recovery disposition follows?

Each view selects fields from the same dictionary and evidence sources. If the weekly and executive numbers disagree, reconciliation is a gate issue.

Five standing rules keep every view honest:

1. One missed decision-binding threshold prevents an overall green status.
2. Adoption never overrides health, risk, or qualified capacity.
3. Missing required data means the gate is not clear, not that the measure is green.
4. Every miss names an action: hold intake, stop the affected path with a fallback and recovery disposition, or approve a time-bound exception through the proper authority.
5. Definitions, sources, owners, and versions travel with the scorecard even when the display hides them.

Operating Ledger Audit

Dashboards can show adoption while work routes around the path. Audit one delivered work unit backward from its verified outcome to the original request. The audit reconciles the evidence chain and the fully loaded cost that invoices and activity dashboards omit.

```
Work unit and verified outcome: _____________________________
Evidence window: ____________________________________________

FULL COST
Context and specification preparation: ______________________
Inference, tools, and vendor services: ______________________
Human first review and re-review: ___________________________
Rework before release: ______________________________________
Platform and control allocation: ____________________________
Incident, exception, or recovery cost: ______________________
Institutional-memory and documentation work: ________________
Total cost / verified outcome: ______________________________

EVIDENCE CHAIN
[ ] Request links to approved Context Spine and outcome.
[ ] Spine links to risk profile, standards, and configuration.
[ ] Implementation links to harness and qualified review.
[ ] Release links to provenance, exceptions, and approver.
[ ] Runtime observation links to the declared outcome.
[ ] Break-glass, fallback, and recovery records reconcile.

Official path versus actual path: ___________________________
Hidden wait, workaround, or labor: _________________________
Decision or repair this audit buys: ________________________
Owner and review date: ______________________________________
```

This audit follows a selected work unit; it is not a demand to inspect every line manually or relabel a dashboard sample. Record every discovered workaround in the Platform Adoption Shadow that follows. The audit discovers the gap; the shadow keeps it actionable.

Platform Adoption Shadow

The designed path and the path people actually use are rarely identical. Track that gap as product evidence. Do not use the shadow to punish teams for revealing friction.

Field	Question
Workflow or work class	Where did the observed behavior occur?
Official path	What does the operating design expect?
Actual behavior	What workaround, manual handoff, or local tool appeared?
Trigger or reason	Why was the official path bypassed?

Field	Question
Missing link	Which intent, evidence, owner, or control became disconnected?
Consequence	Delay, rework, control exposure, cost, or useful local innovation?
Evidence	Observation, pipeline event, interview, register audit, or incident?
Disposition	Improve path, support locally, register exception, exclude, or retire?
Owner and review date	Who closes or consciously accepts the gap, and when?

Harborline illustration

All values below are invented for the composite case.

Completed Harborline dictionary record

```
Metric ID: health.first_pass_alignment
Version: 1
Decision: whether delivery health supports current scope or the next cohort
Definition: applicable declared criteria and constraints satisfied at the first
            structured review
Formula: satisfied applicable criteria / all applicable declared criteria
Unit and population: reviewed payment-retry work units in the chartered cohort
Start event: work unit enters first structured review
Stop event: first review decision is recorded before structural correction
Source: structured evaluation record linked by work-unit ID
Baseline: 61 percent in the approved pre-pilot baseline window
Direction: higher is better
Threshold: improve from the approved baseline after the minimum sample
Window and sample: four weeks and at least 20 reviewed work units
Owner: evaluation lead
Validation: monthly sample reconciled to intent versions and review decisions
Pass action: eligible to advance on this measure
Breach action: hold expansion and run intent-chain analysis
Likely gaming: omit difficult criteria or record re-review as first review
Definition approver: operating owner and evaluation lead
```

Completed Harborline Operating Ledger Audit

```
Work unit: one verified payment-retry policy change from the current window
Outcome: failed-retry cost and duplicate-entry observation linked to work-unit ID
Full cost: model and gateway spend plus spine preparation, ledger first review
           and re-review, platform allocation, provenance-stop exercise,
           record reconciliation, and finance outcome-validation time
Evidence chain: request, spine, Tier 1 profile, evaluated configuration,
                release evidence, runtime signals, and finance outcome reconcile
Actual-path finding: historical records show retry-helper-local-v1 handled work
                     outside the official path before the pilot
Disposition evidence: the helper was retired before product-pilot day 1; its
                      cleanup and support labor remain in fully loaded cost
Decision bought: record the historical shadow cost and keep the verified path;
                 do not reopen or bless the helper
Owner and review: platform owner at the monthly operating review
```

Completed Harborline shadow entry

Field	Entry
Workflow	Retry policy update for one payments service
Official path	Approved spine, shared payments standard, registered reviewer, gated release
Actual behavior	Engineer runs a local retry helper and pastes output into a normal pull request
Trigger or reason	Shared path cannot yet reproduce one service's local test fixture
Missing link	Tool identity, config version, evaluation evidence, and stop path
Consequence	Faster local generation, but the release record cannot explain the generated behavior
Evidence	Work walk and register comparison during the redesign session
Disposition	Add fixture support to the platform backlog; retire the local helper until it passes registration and evaluation
Owner and date	Pipeline architect, before the pilot's first production release

Completed Harborline scorecard

These are the authoritative Harborline scorecard values; the book narrates the decision they force.

Category	Measure	Baseline	Current	Threshold and direction	Action
Adoption	Gate-qualified teams active	0	3	No unqualified intake	Continue current pilot only
Adoption	Paved-path coverage	68%	96%	At least 95% or a recorded exception	Advance on this measure
Health	First-pass alignment	61%	78%	Improve after minimum sample	Advance on this measure
Health	Structural rework	42%	24%	No worse than approved baseline	Advance on this measure
Health	Review latency, median / slow tail	18h / 51h	16h / 39h	At most 18h / 45h	Advance on this measure
Health	Material defect-class findings repaired	2 of 4	4 of 4	No unresolved severe class; every material finding has a class-level repair	Advance on this measure
Economics	Cost per verified outcome	1.00 index	0.82	No increase without approved value gain	Advance on this measure
Economics	Monthly failed-retry cost	18.4k	12.1k	Improve after full evidence window	Observe one more month
Risk	Complete required provenance	72%	100%	Equal to approved requirement, 100% for this Tier 1 pilot	Advance on this measure
Risk	Break-glass review completion	1 of 2	1 of 1	100% inside required window	Advance on this measure
People	Evaluator utilization	31 / 40h	36 / 40h	At most 85%	Hold new intake
People	Funded role and apprenticeship commitments	1 of 4	4 of 4	All gate-required commitments funded and assigned	Advance on this measure

Decision: hold expansion. Delivery health and control integrity improved, but evaluator utilization is 90 percent and the business result needs one more complete window. The scorecard does not call the pilot a failure. It prevents a promising pilot from scaling past its review capacity.

Monthly operating review

Send one page per product line before the meeting.

1. What decision is requested: advance, hold, stop, or exception? If stop, what fallback and recovery disposition applies?
2. What changed in alignment, structural rework, review latency, value, and control integrity?
3. Where is work waiting, and is the wait a clarity problem, path problem, or qualified-capacity problem?
4. Which decision right, agent record, standard, control profile, or harness changed this month?
5. What did the defect taxonomy teach that should change the shared path?
6. What does the adoption shadow reveal about the difference between policy and practice?
7. What changed for people: work thinning, work growing, mobility, apprenticeship, or load?
8. Which assumption was wrong, and what will change before the next review?

Record decisions, dissent, owners, and dates in the meeting. Narrative without a decision is status theater in a better font.

Threshold action table

Condition	Path decision	Action or recovery disposition
Delivery health misses threshold	Hold	Freeze intake; trace the intent chain and repair the class-level cause
Control integrity misses threshold	Stop	Use the verified fallback; preserve evidence, remediate, and require a signed restart decision
A rollback trigger occurs	Stop	Restore the verified state and begin the applicable incident or recovery process
Evaluator utilization exceeds maximum	Hold	Narrow scope, reduce demand, or add qualified capacity before new intake
Value remains unproven after evidence window	Stop	Return to the fallback and redesign before seeking a new authorization
People commitments are not funded	Hold	Assign funded capacity, owner, and date before adding cohorts

Local owners approve the actual thresholds and exceptions. The default actions keep a missed threshold from becoming a footnote.

Gate 6: instrument

```
[ ] Every headline measure supports a named decision.
[ ] Definitions include unit, population, source, exclusions, and owner.
[ ] Thresholds include direction, window, sample, and required action.
[ ] Weekly, monthly, and phase-gate views use one dictionary.
[ ] Operating Ledger Audit can detect shadow paths and missing evidence.
[ ] Value, control, people, and capacity sit beside delivery health.

Decision: [instrumented] [repair evidence] [revise measures] [hold]
Missing measure or source: _____________________________________
Owner: _____________________________  Due: ____________________
```

The final working session turns this evidence model into a 90-day plan and an explicit gate.

PART IV: RUN THE PILOT

Charter, Gate, and Scale

Turn the workshop into a 90-day learning system with one accountable decision owner and no automatic path to expansion.

> *The workshop earns a pilot, not a rollout. The pilot earns its next cohort only when the evidence, controls, capacity, and people commitments say it should.*

Timebox: 45 minutes plus a 15-minute readback

Decision: approve, conditionally approve, or decline the 90-day pilot

Outputs: Platform Pilot Zero charter, product-pilot charter, 30/60/90 plan, operating cadence, gate record

Use Chapter 16, *The Roadmap: From One Team to the Portfolio*, for the argument behind this session. This session owns the working forms.

Platform Pilot Zero charter

Complete this prerequisite before authorizing the product pilot to enter production. Keep it small enough to finish, real enough to expose the path, and separate enough that its evidence cannot be confused with product value.

```
PILOT ZERO NAME: _____________________________________________
Platform owner: _________________  Decision owner: ____________
Repeated internal work class: ________________________________
Internal-service outcome and baseline: _______________________
Boundary / exclusions: _______________________________________
Risk class / human approval: _________________________________
Working-loop evidence required: ______________________________
Stop condition / verified fallback: __________________________
Recovery and restart drill: _________________________________
Workarounds-to-backlog owner: _______________________________
First outside adopter: ______________________________________
Gate date: __________________________________________________
```

Platform Pilot Zero closes only when one real work unit completes the loop, the stop path and the rollback or disablement drill required by its control profile pass, the record chain reconciles, and every workaround has an owner in the paved-path backlog. A hold keeps the product pilot out of production while the prerequisite is repaired. A stop returns platform work to its verified fallback and requires a new authorization before another candidate begins.

Product-pilot charter

```
PILOT NAME: _____________________________________________
Sponsor: _____________________  Operating owner: _____________
Outcome owner: _____________  Gate decision owner: __________

Value hypothesis: _____________________________________________
Business measure and baseline: _____________________________
Bounded work unit: _____________________________________________
Non-goals and exclusions: _____________________________________
Risk class and obligations: ___________________________________
Named spec owner / evaluator / standards keeper / pipeline architect:
_____________________________________________________________
Evaluator demand / usable capacity: ___________________________
Evidence window and minimum sample: ___________________________

Advance conditions: _____________________________________________
Hold conditions: _____________________________________________
Stop conditions: _____________________________________________
Fallback path: _____________________________________________
Recovery disposition and last verified state: _______________
First gate date: _____________________________________________
```

No condition may say “leadership judgment” without naming the leader and the evidence that informs the judgment. The product charter may be drafted while Platform Pilot Zero runs; its production start is conditional on the Platform Pilot Zero gate.

Harborline's completed charter summary shows the shape:

```
Pilot: Payment retry consolidation
Value hypothesis: Consolidating retry behavior behind one reviewed spine and
  gated path will reduce monthly failed-retry cost without increasing duplicate entries.
Business baseline: 18.4k USD monthly failed-retry cost
Health baselines: 61% first-pass alignment; 42% structural rework;
  18h median and 51h slow-tail review latency
Risk profile: Tier 1; money and ledger integrity; named evaluator and human release owner
Evidence window: four weeks and at least 20 reviewed work units
Capacity threshold: evaluator utilization at most 85% of usable staffed hours
Stop: duplicate-entry invariant failure, unexplained release, or severe unresolved control
failure
Rollback: revert shared retry and reviewer configurations; stop affected release path;
  reconcile impacted ledger state before restart
Decision owner: executive sponsor on operating-owner recommendation
Work suitability: 13 of 14, with no critical zero; dependency complexity scored 1
Status: illustrative charter, pending local approval
```

The scorecard in the evidence session reports how these baselines moved.

Phase zero: close the owner's journey

Work	Owner	Evidence of completion	Due
Approve the Platform Pilot Zero mini-charter			
Baseline the internal-service outcome and path health			
Complete one real work unit through the full path			
Reconcile intent, approval, release, and outcome evidence			
Exercise stop, verified fallback, the profile-required rollback or disablement drill, and restart			
Turn every workaround into an owned platform backlog item			
Decide advance, hold, or stop at the phase-zero gate			

The first 90 days

Days 1 to 30: establish the loop

Work	Owner	Evidence of completion	Due
Approve the intent spine and non-goals			

Work	Owner	Evidence of completion	Due
Baseline health, value, control, and people measures			
Configure and test the paved path			
Approve risk class and decision rights			
Staff evaluator capacity and calibrate reviewers			
Test stop, rollback, and break-glass			
Confirm Platform Pilot Zero gate evidence remains valid			

Days 31 to 60: run bounded production work

Work	Owner	Evidence of completion	Due
Admit only qualified work through the path			
Review weekly delivery health and capacity			
Run intent-chain reviews on structural rework or defects			

Work	Owner	Evidence of completion	Due
Update standards and harness from findings			
Audit one delivered work unit through the Operating Ledger			
Hold manager, apprenticeship, and mobility conversations			

Days 61 to 90: complete evidence and decide

Work	Owner	Evidence of completion	Due
Complete the agreed outcome window			
Reconcile scorecard sources and definitions			
Re-test stop and rollback path			
Review exceptions, incidents, load, and learning			
Revisit the organization-design hypothesis			
Prepare and read the phase-gate record			

Operating cadence

Cadence	Participants	Required inputs	Decision
Twice-weekly path review	Delivery, platform, evaluator	Queue, alignment, rework, latency, exceptions	Intake and immediate repairs
Weekly pilot review	Operating owner and capability owners	Weekly scorecard and evidence gaps	Continue, narrow, hold, or stop
Monthly outcome review	Sponsor, outcome, finance, control, people	Monthly view and completed Operating Ledger Audit	Funding, value, people, and control action
Day-90 phase gate	Named decision owner and required approvers	Complete gate packet	Advance, hold, or stop; record fallback and recovery disposition

Meetings are not the operating model. Decisions, updated artifacts, and named actions are.

Phase-gate decision record

```
Pilot / cohort: _____________________  Gate date: __________
Decision owner: _____________________  Recommender: _________
Evidence window: ____________________  Minimum sample met: [yes/no]

Readiness conditions closed:                          [yes/no + evidence]
Delivery-health thresholds met:                       [yes/no + evidence]
Value hypothesis supported:                           [yes/no + evidence]
Control-integrity thresholds met:                     [yes/no + evidence]
Evaluator capacity inside approved maximum:           [yes/no + evidence]
People and apprenticeship commitments funded:         [yes/no + evidence]
Operating Ledger Audit complete and trustworthy:      [yes/no + evidence]
Organization-design hypothesis supported or revised:  [yes/no + evidence]

Path decision: [advance] [hold] [stop]
Scope approved for next period: _____________________________________
Conditions and actions: _____________________________________________
Fallback if stopped: ________________________________________________
Recovery disposition: [no production rollback] [rollback required]
Repair, reconciliation, disablement, or rollback evidence: _____
Dissent or minority view: ___________________________________________
Owners and dates: ___________________________________________________
Next gate: __________________________________________________________
```

Advance is an active decision. Time passing is not approval.

Harborline illustration

Before day 1 of the product pilot, Harborline closes Platform Pilot Zero on a repeated non-production service-account request. One real request moves from a versioned spine through policy tests, identity approval, registration, release, and audit evidence. The team triggers a missing-expiry stop, disables the paved route, uses the verified manual fallback, reconciles the record, adds a regression test, and records approval to restart. That disablement is the recovery drill required by this control profile. The phase-zero decision owner advances the shared path to its first outside adopter: the payments product team. None of that evidence is counted as payment-pilot value.

At day 90, Harborline's path shows better first-pass alignment, lower structural rework, complete required provenance, and a promising reduction in failed-retry cost. Its Operating Ledger Audit is complete. The stop and profile-required recovery tests pass. The org hypothesis is only partly supported: protected hats work, but payment evaluation is already over its approved maximum-utilization threshold.

```
Decision: HOLD EXPANSION
Current pilot: continue inside existing scope
New cohort: not approved

Conditions to reconsider:
1. Qualified payment-evaluation backup completes calibration.
2. Payment evaluator utilization remains at or below 85 percent for four weeks.
3. One more monthly value window supports the failed-retry result.
4. People leader confirms apprenticeship capacity for the next cohort.

Stop and recovery conditions still in force:
Missing required provenance stops the affected path and rolls a changed
configuration back to its evaluated version. A duplicate-entry invariant
failure stops the path, rolls back the affected release, and requires ledger
reconciliation.

Decision owner: Payments engineering VP
Next gate: 30 days after the capacity action closes
```

The sponsor honors the hold. That act teaches the organization more about the operating model than a speech about responsible scale.

Scaling sequence after an advance

Use the same gates at each level.

1. **Internal first use:** platform or enablement team tests the path on its own real work and decides whether it is ready for an outside adopter.
2. **First outside adopter:** one willing team tests whether the path travels.
3. **Small cohort:** a few representative teams expose domain and capacity differences.
4. **Portfolio service:** only patterns supported by repeated evidence become defaults, funded capabilities, or structural changes.

At each step, retain an exclusion register and permit the evidence to narrow the ambition. Every exclusion keeps its reopening evidence beside it, the way Harborline's settlement entry names a vendor API release, a change-volume increase, a control failure, or a missed review date. A no with a reopening condition stays a decision. Scale the path, not a slide deck.

Sponsor commitments

```
[ ] I will not count unqualified activity as adoption.
[ ] I will honor holds and stops, including when they delay my stated target.
[ ] I will fund evaluation, platform, control, and apprenticeship capacity.
[ ] I will require people effects to be discussed before scale decisions.
[ ] I will keep the organization design reversible until the evidence supports it.
[ ] I will name what happens to efficiency gains and affected roles.
[ ] I will publish transition commitments and fund mobility and training before roles disappear.
[ ] I will record unresolved workforce objections, and my response, in the gate decision record.

Sponsor: _____  Date: _____
```

Gate 7: authorize

```
Decision: [approve pilot] [approve with conditions] [decline]
Conditions that must close before day 1: _____
First action: _____  Owner: _____
Artifact repository: _____
First operating review: _____  First gate: _____
```

The facilitator now uses the kit that follows to close the room, publish the packet, and make the first commitments hard to forget.

FACILITATOR KIT

The Facilitator Kit

A compact runbook for preparing the room, protecting the decisions, publishing the packet, and bringing the evidence back to the next gate.

> *The workbook is useful only if the artifacts survive the meeting. Facilitation ends when the decisions are versioned, the actions are owned, and the next gate is on the calendar.*

Use this kit with the session worksheets, not instead of them. The session cards, prompts, and registers here apply in every working mode, from the 30-minute diagnostic to the phase gate.

Before the room

Facilitator preflight

```
[ ] Sponsor's one-sentence decision is specific and bounded.
[ ] Candidate team and work are named, but not pre-approved.
[ ] Outcome owner and required control partners can attend the gates.
[ ] Source pack arrived three working days in advance.
[ ] Missing evidence is visible in the agenda.
[ ] Versioned destination for the packet is ready.
[ ] Decision log, parking lot, action register, and timer are visible.
[ ] Scribe knows the artifact names and chapter sequence.
[ ] Remote participants can edit, vote, and see the timer equally.
[ ] Sponsor understands that advance, hold, and stop are valid path decisions,
    with a separate fallback and recovery disposition.
```

Source-pack triage

Source	Present and current	Owner if missing	Decision affected
Recent work items and intent artifacts			
Current workflow and exception path			
Pipeline coverage and elapsed time			

Source	Present and current	Owner if missing	Decision affected
Review demand, capacity, loops, and latency			
Written standards			
Agent and automation inventory			
Outcome baseline and source			
Incident, control, and break-glass evidence			
People, role, and staffing constraints			

If a missing source can reverse the decision, do not resolve it with discussion. Convert it into a condition with an owner and date.

During the room

Open with this contract

```
Today we are designing one evidence-producing pilot, not approving a general
rollout. We will distinguish facts, assumptions, and choices. Unknowns will
remain visible. Every gate can advance, hold, or stop. A hold or stop may
require redesign, fallback, repair, reconciliation, disablement, or rollback.
The sponsor will read the final decision aloud.
```

Session card

Use the same card at the start of each chapter.

```
Session: _____________________________________________________
Decision due at the end: _____________________________________
Timebox: _____________________  Decision owner: _____________
Inputs we have: ______________________________________________
Inputs missing: ______________________________________________
Artifact we will complete: ___________________________________
```

Prompts that improve the work

When the room offers opinion

- What did we observe?
- Which source would change this answer?
- Is that current state, desired state, or a design choice?

When a senior voice closes the question early

- Who has evidence closest to the work?
- What would make the opposite answer true?
- What dissent should the decision record preserve?

When scope expands

- Does this belong inside the pilot boundary?
- What control or capacity assumption changes if we include it?
- Should it enter the exclusion register or the next-cohort backlog?

When automation becomes the answer to every step

- Which step should disappear instead?
- Which judgment remains human and why?
- What evidence will reveal a confidently wrong result?

When the room reaches for the org chart

- Which capability or decision in the redesigned flow requires this box?
- Can a protected hat or temporary interface test the hypothesis first?
- What evidence would make the structural change permanent?

When a threshold is debated

- Which decision does it trigger?
- In which direction is better?
- What window and minimum sample make it credible?
- Who may approve an exception, and when does it expire?

Decision log

ID / time	Decision	Evidence and assumptions	Owner	Dissent	Review trigger or date

Action register

Action	Why it matters	Owner	Due	Evidence of closure	Status

Assumption and parking-lot register

Item	Assumption or parking lot	Decision affected	Test or next forum	Owner	Date

Item	Assumption or parking lot	Decision affected	Test or next forum	Owner	Date

An assumption has a test. A parking-lot item has a destination. Anything else is a polite way to lose a problem.

Close the room

Decision readback

The sponsor reads this aloud. The scribe records corrections in the artifacts, not in a separate set of minutes.

```
We decided to [approve] [approve with conditions] [decline] the pilot.

The outcome is: _____
The boundary is: _____
The operating owner is: _____
The risk and control owner is: _____
The staffed evaluator capacity is: _____
The evidence window is: _____

We advance when: _____
We hold when: _____
We stop when: _____
If stopped, the verified fallback is: _____
Recovery disposition: [no production rollback] [rollback required]
Repair / reconciliation / disablement / rollback evidence: _____

The first action is: _____
The first operating review is: _____
The first phase gate is: _____
```

Ask every required decision-maker, in turn: *Can you execute the decision as read?* Record a no as a condition or dissent, not as silence.

Workshop readout

Keep the readout to one page. Link to the packet rather than pasting its content.

```
Decision and date: _____
Pilot and outcome: _____
Scope / non-goals: _____
Why this team and work: _____
Operating-model hypothesis: _____
Risk class and critical controls: _____
Capacity decision: _____
Headline measures and threshold actions: _____
Advance / hold / stop conditions: _____
Fallback and recovery disposition: _____
Open conditions, owners, dates: _____
Packet location and version: _____
Next gate: _____
```

Publish the packet

```
[ ] Worksheets were converted from workshop notes into readable records.
[ ] Evidence links resolve for authorized readers.
[ ] Decisions, conditions, dissent, owners, and dates are preserved.
[ ] Artifact versions and approvers are visible.
[ ] Agent, exception, and exclusion registers are updated.
[ ] Calendar holds exist for operating reviews and the phase gate.
[ ] People affected by the pilot receive a direct manager conversation.
[ ] No slide deck is the sole record of a decision.
```

Publish within two working days. If the group cannot recognize the published decision, reconvene the decision owner rather than rewriting history.

Context Spine template

One file per unit of work, living in the repository beside the code it governs. Keep the five-field body stable: problem, acceptance criteria, constraints, interfaces, and non-goals. Extend the identity header when risk or operating context requires it. The filled values below come from the Harborline payment-retry composite; replace them with your own.

```
---
id: spine-2026-0147
title: Payment retry consolidation
owner: jsmith
outcome_id: retry-failure-cost
status: in-review        # draft | in-review | approved | superseded
version: 1.2
systems: [payments-core, ledger-api]
risk_tier: 1
regulated_obligations: [reg-e-error-resolution]   # or []
approver_of_record: mchen
value_measure: monthly cost of failed retries
baseline: 18400 USD
last_updated: 2026-07-05
open_questions:
  - Does the retry ceiling vary by payment rail?
---

## Problem statement
What is true today, why it is a problem, and for whom. Three sentences.

## Acceptance criteria
Each criterion testable, each one a future failing test.
- Retries use exponential backoff with jitter for transient 5xx errors.
- Duplicate submissions under concurrency produce exactly one ledger entry.

## Constraints and invariants
What must hold no matter how clean the diff looks.
- Idempotency guarantees on all consumer endpoints are preserved.
- No new connection pressure on the ledger during partial outages.

## Interfaces
Contracts touched, with references to their definitions. What changes, what must not.

## Non-goals
The cheapest sentences in the spec and the most expensive to omit.
- No changes to the message schema or dead-letter routing.
```

Canonical artifact crosswalk

The workbook sessions own the canonical fields, blank forms, and completed examples for all twenty-four Operating Kit artifacts. This crosswalk maps each session to the artifacts it completes; the book's Operating Kit appendix is the reader-facing map of the same twenty-four.

Workbook session	Canonical Operating Kit artifacts
Choose the Team and the Work	Readiness Worksheet; Pilot Charter and Work-Suitability Rubric, used separately for Platform Pilot Zero and the product pilot
Redesign the Work	Context Spine Template; Work Redesign Canvas; Exclusion Register; Platform Adoption Shadow

Workbook session	Canonical Operating Kit artifacts
Design the Operating Model	Operating Model Canvas; Platform Capability Charter; Operating Ownership Matrix; Organization Design Decision Record
Staff the Capabilities	Four Operating Role Cards; Evaluator-Capacity Worksheet; Manager Conversation Guide
Set the Risk and Control Profile	Decision-Rights Table; Risk Classification and Control Profile; Release Gate Checklist; Agent Register; Intent-Chain Post-Mortem
Define the Evidence	Metrics Dictionary; Canonical Scorecard and Role-Specific Views; Monthly Operating Review; Operating Ledger Audit
Charter, Gate, and Scale	Pilot Charter; Phase-Gate Decision Record, used first for Platform Pilot Zero and then for the product pilot
Facilitate and Follow Through	Facilitator Guide

When a local team adds or removes a field, capture the reason and approver in the Organization Design Decision Record or the appropriate policy record.

Harborline completed-packet index

Harborline is filled progressively through the workbook. This index shows how the illustration connects; it is not a claim of measured enterprise results.

Packet item	Filled Harborline result
Readiness	7 of 10, with no critical zero; temporary helper must be registered or retired and evaluation backup protected
Platform Pilot Zero	One repeated non-production service-account request completes the loop; a missing-expiry stop, profile-required disablement drill, verified manual fallback, reconciliation, regression test, and restart approval close before product production use
Work suitability	Payment-retry consolidation passes; schema and settlement work excluded
Work redesign	One spine, one gated retry path, ledger evaluation, finance feedback
Capability charter	Platform director owns a funded governed-delivery capability with a named internal journey, first outside adopter, boundary, and promotion criteria
Organization decision	ODDR HL-OD-001 tests protected team hats plus one platform capability before permanent boxes

Packet item	Filled Harborline result
Capacity	Segment allocations are 18/18 payment, 10/12 ledger, and 8/10 Tier 2 hours; de-duplicated total is 36/40, 90 percent, so new intake is held
Control profile	Highest local tier; human-owned merge and release; tested stop and rollback
Metrics dictionary	`health.first_pass_alignment` is versioned with formula, events, source, baseline, minimum sample, gaming risk, and pass / breach actions
Scorecard	Alignment improves after the minimum sample; latency is at most 18h / 45h; Tier 1 provenance is 100%; capacity triggers hold
Operating Ledger Audit	One delivered work unit reconciles full cost and the evidence chain; the already-retired helper's historical shadow labor remains in cost
Gate record	Continue current scope, hold cohort expansion, close capacity and evidence conditions
Gateway-alias stop record	Provenance failure stops the affected configuration, uses the pre-pilot fallback, and requires no production rollback after released-change assessment; signed restart follows repair and reconciliation
Configuration-recovery drill	A separate illustrative scenario stops Tier 1 releases and rolls reviewer configuration v4 back to evaluated v3; production rollback remains a separate impact decision

Between gates

The facilitator's role becomes lighter but does not disappear.

```
WEEKLY
[ ] Decision and action registers are current.
[ ] Threshold misses have named actions.
[ ] Scope has not expanded outside the exclusion register.

MONTHLY
[ ] Source definitions reconcile across views.
[ ] Platform Adoption Shadow is sampled.
[ ] People effects and commitments are reviewed.
[ ] Organization-design assumptions are updated in writing.

BEFORE THE GATE
[ ] Evidence window and minimum sample are complete.
[ ] Stop and rollback paths were re-tested.
[ ] Required decision-makers received the packet early.
[ ] Advance, hold, and stop remain real path decisions.
[ ] Every stop names its fallback, recovery disposition, and restart evidence.
```

The final test is simple: could a new decision-maker reconstruct what was tried, why it was allowed, what happened, and why the organization chose its next move? If yes, the packet is doing its job.

NOTES, ASSUMPTIONS, AND USE

Notes, Assumptions, and Use

This workbook operationalizes Part IV of the second edition of *The Bottleneck Is Never the Stack*, and it is the canonical home of the Operating Kit's artifact fields, forms, and completed examples. It does not add a second theory of the operating model. Read the book's notes and sources for the research, regulatory references, and argument behind the exercises.

What is illustrative

The book's evidence legend applies here: observed practice, external reported evidence, composite illustration, or planning assumption. Everything filled in below the blank forms belongs to the last two categories.

Harborline is a composite created to show a connected packet. It is not a disguised case study. Its work, roles, readiness score, capacity values, scorecard results, thresholds, and gate decisions are invented examples. They demonstrate how the artifacts connect; they do not establish a benchmark or predict a result.

The workbook's default readiness screen, 7 of 10 with no zero in pipeline coverage, evaluation capacity, or ownership, is a starting decision rule. The 85 percent maximum evaluator-utilization threshold is also illustrative. Local owners should approve thresholds, directions, evidence windows, minimum samples, and exception authority in the Metrics Dictionary and risk policy before use.

What must be local

The authorized people in each organization must decide:

- system and work risk classification;
- legal, regulatory, privacy, security, model-risk, and employment obligations;
- which actions an agent may take and which a human must own;
- retention, provenance, vendor, data, and audit requirements;
- scorecard definitions and decision thresholds;
- staffing, career, mobility, and transition commitments; and
- advance, hold, and stop authority; fallback, recovery, rollback, restart, and exception authority.

Provider terms, product controls, models, training-use policies, retention, indemnification, and regulation change. Validate the current facts with counsel, procurement, privacy, security, risk, finance, and people leaders before turning an illustrative form into policy.

How to adapt the packet

Keep the artifact names and decision sequence aligned with the main book. Add local fields when they support a real decision or obligation. Remove a field only with a named owner and recorded rationale. Version the packet, record definition changes, and use an Organization Design Decision Record for material changes to the operating model.

This is operating guidance, not legal, regulatory, accounting, or employment advice. The local decisions belong to the people with the authority and context to make them.