

Process Re-engineering Sprint — Sample Deliverable

Prepared for Coastal Claims Group

Engagement Reference: sample

Generated on August 27, 2026

DRAFT — not yet certified

Table of Contents

Executive Summary	3
Research & Market Context	4
Findings & Analysis	6
Process mapping (as-is)	10
Bottleneck & cost analysis	14
To-be design	17
Implementation playbook	21
Risks & Mitigations	25
90-Day Action Plan	29
Appendix: Methodology & Assumptions	32

Executive Summary

Coastal Claims Group is a 41-person Florida/Georgia claims administration business facing a specific, urgent problem: an average 21-day claim cycle time — against a target of under 10 days — that is driving client complaints and putting two of the company's largest contracts at risk. The constraints on the fix are as important as the problem itself: no headcount can be added, the core Guidewire system cannot be replaced this year, human judgment cannot be removed from claim assessment or approval, and every change must preserve the documented timelines and audit trails Florida and Georgia regulators require. This engagement mapped the current process end-to-end, quantified where time and money are actually being lost, designed a to-be process that works within those constraints, and built a 90-day plan to implement it before hurricane season begins.

The analysis produced five headline findings. **First, document collection is the dominant bottleneck by every measure** — it consumes 8 of 21 days (38% of total cycle time) and roughly 37% of processing-labor cost, driven by an 18% rate of incomplete initial submissions that forces a second request round adding 4–5 days each time. **Second, this bottleneck is structurally self-inflicted, not inherent to the work:** there is no standardized document checklist and no automated follow-up, so incomplete submissions are caught only after a processor manually reviews what arrives, and the company's own 2022 attempt to fix this with a dedicated "document chase" employee proves the concept works but also proves that a person-dependent fix doesn't survive staff turnover. **Third, QA and manager approval operate as universal, sequential gates rather than risk-differentiated checkpoints** — every claim, regardless of complexity, passes through the same 3-day paper-based QA review and 2-day approval step, even though QA only finds issues on 12% of claims, meaning the other 88% wait in a queue that changes nothing for them. **Fourth, the process runs on three disconnected systems** — Guidewire, Outlook, and an informally maintained Excel tracker — creating both unmeasured handoff delays (accounting for part of the discrepancy between the client's own stage-by-stage estimate and the stated 21-day average) and a latent audit-trail weakness, since assignment timing and document-request status currently live outside the system of record. **Fifth, the math works:** recoverable time across all identified bottlenecks totals 7.8–10.3 days, which brackets the under-10-day target almost exactly — meaning the goal is achievable through disciplined execution of known fixes, not through unproven or speculative interventions, though reliably hitting it during the 40–60% hurricane-season volume surge will require the redesign to perform at the stronger end of these estimates.

Five recommendations follow directly from these findings, sequenced by impact and implementation risk. **First, replace manual document requests with a standardized, claim-type-specific checklist and an automated, system-triggered reminder cadence** — built into Guidewire's existing workflow capability so the fix survives staff turnover, directly repairing the failure mode of the 2022 pilot. **Second, eliminate the manual assignment queue** by routing new claims automatically based on real-time caseload data already in Guidewire, retiring the disconnected Excel tracker in the process. **Third, digitize the QA checklist and apply risk-tiered review** — full manual review preserved for high-value or flagged claims, faster sampling-based review for the remaining majority — which both speeds throughput and strengthens the audit trail through automatic timestamping, though this step requires compliance counsel sign-off before go-live given its regulatory sensitivity. **Fourth, parallelize manager approval and template decision communications** for the lower-risk claim tier, closing a final logging gap where communication method isn't always recorded back into Guidewire. **Fifth, govern the rollout in three sequenced phases over 90 days** with hard go/no-go gates — particularly

requiring compliance validation before the QA tiering goes live and requiring the Phase 1 pilot to show measurable improvement before scaling — so the program can be corrected early rather than discovered broken during hurricane season.

If Coastal executes this plan, the expected impact is a reduction in average cycle time from 21 days to an estimated 9–12 days, achieved without adding headcount, without replacing Guidewire, and without reducing human oversight of claim decisions — directly satisfying the constraints leadership set. The financial case is straightforward: implementation requires roughly 100–150 hours of existing staff time over 12 weeks and no new system spend, against client-contract revenue exposure plausibly in the seven figures if the two threatened relationships are lost. Success depends most on disciplined adoption by the 22-person processor team and timely compliance sign-off on QA tiering — both addressed explicitly in the risk mitigation plan — and on completing the rollout by end of May so the redesign is proven under normal volume before facing its first real test during the June–November surge.

Research & Market Context

Industry and Market Context

Coastal Claims Group operates in the property & casualty (P&C) claims administration segment of the insurance services industry, serving carriers and policyholders across Florida and Georgia — two of the highest-volume, highest-volatility P&C markets in the country due to hurricane exposure. This regional context matters for benchmarking: Florida in particular has among the most stringent state-mandated claims-handling requirements in the U.S. (Fla. Stat. § 626.9541 and related unfair claims practices provisions generally require acknowledgment of claims within 14 days and payment or denial decisions within 90 days of receiving proof-of-loss documentation, with additional prompt-pay rules for specific lines). Georgia imposes comparable, though slightly less prescriptive, timelines under its Unfair Claims Settlement Practices Act. These regulatory floors are wide relative to Coastal's current 21-day average and even wider relative to its 10-day target, confirming that the objective is a competitive/service differentiator rather than a compliance necessity — but it also means any re-engineered process must preserve documented timestamps and audit trails at every step, since regulators and carrier clients audit against these statutory windows.

The claims administration sub-sector — particularly third-party administrators (TPAs) and brokerage-affiliated claims units like Coastal — is under structural pressure from two directions simultaneously: rising claim volumes driven by increased storm frequency/severity in the Southeast, and rising client (carrier) expectations for speed and transparency, often contractually enforced through service-level agreements (SLAs). Industry surveys (e.g., LexisNexis Risk Solutions claims studies, J.D. Power P&C claims satisfaction research) consistently show that cycle time and communication frequency — not payout amount — are the top two drivers of policyholder and carrier-client satisfaction. This aligns directly with Coastal's stated problem: client complaints and contract risk are tied to speed and rework, not to claim outcomes.

Assumption (labeled): Based on general industry sources rather than client-provided data, typical non-catastrophe P&C claims cycle times for comparable regional TPAs/brokerages range from 7–15 days for straightforward property claims, with best-in-class digitally-enabled operators achieving 5–7 days. Coastal's 21-day average places it meaningfully above typical market performance even before accounting for hurricane-season surges, reinforcing the urgency of the engagement.

Benchmarks and Typical Performance Ranges

To anchor the bottleneck analysis and to-be design, it is useful to compare Coastal's current-state metrics against typical industry ranges. The following benchmarks are drawn from general claims-operations literature (IAIABC claims processing studies, McKinsey and Deloitte insurance operations research, and TPA industry association data) and are explicitly flagged as **assumptions** where Coastal-specific data was not available in intake:

- **Document collection cycle time:** Industry benchmark for first-pass document completeness in P&C claims is typically 65–75% (i.e., 25–35% of claims require a follow-up request). Coastal's 18% incomplete-submission rate is actually *better* than this typical range — meaning the real issue is less about the *rate* of incomplete submissions and more about the *lack of automated follow-up*, which extends the tail of that 18% by 4–5 days per affected claim. **(Assumption: industry incomplete-submission benchmark of 25–35%, used for comparative context only.)**
- **Straight-through processing (STP) rates:** Leading claims operations that use automated intake triage and document-checklist validation report STP rates (claims requiring no manual rework loop) of 70–85%. Coastal's current QA rework rate of 12% plus its document rework rate of 18% suggests an effective "first-pass clean" rate of roughly 70% (accounting for overlap) — within range, but with clear room to close the gap toward 85%+ through upstream validation rather than downstream catching.
- **Queue/assignment wait time:** Manual, round-robin assignment processes in comparable-sized operations typically produce 1–3 day queue delays, consistent with Coastal's reported 1.5 days. Automated or rules-based assignment/triage engines in similar-sized claims teams (20–30 processors) typically reduce this to same-day or next-business-day assignment — a benchmark of **under 4 hours** is commonly cited for digitally-routed intake. **(Assumption, based on general claims-ops automation case studies.)**
- **QA review turnaround:** Manual, paper-based QA checklists in mid-sized claims units typically take 2–4 days per review cycle; digitized checklist tools with parallel (rather than sequential) review integrated into the core system typically cut this to 0.5–1.5 days. Coastal's 3-day QA step is squarely within the "manual/paper" range, not the digitized range.
- **Cost per claim:** At \$215K/month fully loaded cost against ~340 claims/month, Coastal's approximate cost per claim is **\$632** (rising further during hurricane-season surges if volume increases 40–60% without proportional cost increases, which would actually *improve* unit cost if throughput holds — a relevant consideration for the to-be design). **Industry benchmark for comparable regional TPA claims-handling cost per claim is estimated at \$450–\$700**, placing Coastal within a reasonable but not best-in-class range. **(Assumption, general industry TPA cost benchmark.)**

Collectively, these benchmarks suggest Coastal's inefficiency is concentrated in *manual coordination overhead* (assignment, document chasing, paper QA handoffs) rather than in staffing capacity or fundamental claim complexity — a critical distinction supporting a headcount-neutral, process-redesign approach rather than a staffing or systems-replacement intervention.

Comparable Practices from Leading Firms

Several practices are well-documented among higher-performing claims organizations and TPAs, and are directly relevant to Coastal's redesign given the "no new systems, headcount-neutral" constraint:

- **Automated document-chase workflows layered on top of existing core systems.**

Leading TPAs commonly use lightweight automation (scheduled email/SMS reminders, portal upload nudges) triggered by the core claims system's status field, without replacing the core system itself. This directly addresses the lesson from Coastal's failed "document chase person" pilot: the fix that stuck in comparable firms was *automating* the chase function so it doesn't depend on a single employee's ongoing manual effort.

- **Rules-based, system-driven assignment (skills/workload balancing) rather than manual round-robin.** Many mid-sized claims operations replace team-lead-managed spreadsheets with simple queue-management logic (even within existing tools like Guidewire's workflow/business rules engine, or a connected lightweight scheduling layer) to eliminate the 1–2 day unassigned-queue gap entirely.

- **Parallel and risk-tiered QA, not universal sequential QA.** Best-practice claims units triage QA intensity by claim complexity/value, applying lighter-touch or sampling-based review to low-complexity, low-value claims and reserving full manual QA for higher-risk files — preserving audit rigor where regulatorily required while removing QA as a universal bottleneck.

- **Digitized checklists and sign-offs replacing paper-based QA**, maintaining the audit trail (often *more* robust than paper, since timestamps are system-logged) while cutting turnaround from days to hours.

- **Pre-submission document checklists communicated at first contact**, reducing incomplete-submission rates by setting clearer expectations upfront — a low-cost intervention several regional brokerages have used to cut second-request rates by 30–50%. (**Assumption: improvement range based on general process-improvement case studies, not Coastal-specific data.**)

Frameworks Anchoring This Analysis

This engagement will apply three complementary frameworks. **Lean process mapping and value-stream analysis** will structure the as-is documentation and bottleneck identification, distinguishing value-added time (document review, assessment drafting) from non-value-added wait time (queue delays, rework loops) — critical given that document collection alone accounts for roughly 38% of total cycle time. **Theory of Constraints** will guide bottleneck prioritization, focusing redesign effort on the single largest constraint (document collection) before optimizing downstream steps, avoiding the common failure mode of locally optimizing QA or approval steps while the true bottleneck persists. Finally, a **regulatory-compliance overlay** — mapping every redesigned step against Florida/Georgia statutory timelines and audit-trail requirements — will run throughout the to-be design and implementation playbook, ensuring speed gains never come at the expense of documented compliance, consistent with the engagement's explicit constraints.

Findings & Analysis

Root Cause Analysis

Coastal's 21-day average cycle time is not the product of one failure but of four compounding, interlocking root causes — and the intake data makes it possible to isolate and size each one precisely.

Root cause 1: Manual, disconnected coordination mechanisms. The as-is process runs on three uncoordinated tools — Guidewire for system-of-record, Outlook for all policyholder and

internal communication, and a shared Excel tracker "some team leads keep" for assignment status that is explicitly "not connected to Guidewire." This is the structural source of the 1.5-day assignment queue wait: assignment is "manual, round-robin by team lead," meaning a human must notice a new claim, check availability off a spreadsheet, and hand it off — with no system trigger. It is also the reason the 2022 document-chase fix "didn't stick": the intervention was a *person*, not a *process or system rule*, so when that person left, "the responsibility got redistributed back to processors" and the underlying workflow reverted to its original, unreliable state. Any redesign that again relies on a dedicated human chaser without a system-level trigger will fail the same way.

Root cause 2: No upstream quality control on document requests. Step 3 — "processor requests supporting documents from policyholder (email, no automated reminders)" — has no checklist standardization, no reminder cadence, and no validation before submission. This produces the single largest cost driver in the entire process: 18% of the 340 monthly claims (roughly 61 claims/month) require a second document round, each adding 4–5 days. There is no mechanism to catch incompleteness before the processor manually reviews it, so the cost of the error is paid in full cycle time rather than prevented at the point of request.

Root cause 3: QA and approval are universal, sequential gates rather than risk-differentiated checkpoints. Every claim — regardless of complexity or value — passes through the same paper-based QA checklist (3 days) and the same manager approval step (2 days), in strict sequence after processor drafting. QA "catches errors/omissions on roughly 12% of claims," meaning 88% of claims sit in the same 3-day queue for a check that, for them, changes nothing. This is a classic case of a control designed for the exception being applied to the rule.

Root cause 4: Compliance and audit-trail requirements are currently satisfied through slow, manual artifacts. The paper QA checklist "gets scanned and attached to the file" — a workaround that preserves the audit trail Florida and Georgia regulators require, but does so at the cost of speed, since paper-based sign-off cannot be parallelized or automated. This is important context: the constraint is not the *existence* of QA and audit documentation, but the *medium* through which it's currently captured.

Quantifying the Gap

Coastal's own breakdown of the 21-day average gives a clear, defensible baseline for where time is lost and what closing the gap to under 10 days actually requires.

Process Stage	Current Avg. Days	% of Total Cycle	Root Cause Driving Delay	Est. Days Recoverable
Assignment queue wait	1.5	7%	Manual round-robin, no system trigger	1.0–1.3
Document collection (incl. rework)	8.0	38%	No upfront checklist, no automated reminders, 18% second-request rate	4.5–5.5
Processor review/drafting	4.0	19%	Rework triggered by incomplete docs; some inherent to task	0.5–1.0

QA review	3.0	14%	Universal sequential paper-based review; 12% rework rate	1.5–2.0
Manager approval	2.0	10%	Sequential after QA; no parallelization	0.5–1.0
Decision communication	1.0	5%	Manual email/mail drafting	0.3–0.5
Total	21.0 (actually 19.5 per stated breakdown, treated as ~21 with rounding/variance)	100%		8.3–11.3

(Note: the client's stage breakdown sums to 19.5 days against a stated 21-day average — a modest discrepancy likely reflecting variance across claim types, hurricane-season surges, or unmeasured handoff time between stages not captured in the six line items. This gap itself is a finding: the as-is process mapping should explicitly measure inter-stage handoff time, since even 1.5 unaccounted days suggests additional friction between steps not visible in the current tracking tools.)

The arithmetic case for the target is straightforward: to move from ~20–21 days to under 10, Coastal needs to recover roughly 10–11 days, or just over half of total cycle time. The table above shows that the recoverable range (8.3–11.3 days) brackets this target almost exactly — meaning the objective is achievable through the identified bottlenecks alone, without requiring speculative or unproven interventions. This is a materially different conclusion than "we need to cut every step by 50%"; it shows that concentrated effort on two stages — document collection and QA — can carry most of the burden.

Document collection is the dominant lever. At 8 of 21 days (38% of total cycle time), it is nearly three times the size of the next-largest stage. Using Coastal's own figures: 18% of ~340 claims/month = ~61 claims/month carrying an extra 4–5 days each purely from second-request rework. If a pre-submission checklist and automated reminder cadence reduce the second-request rate from 18% to an industry-informed 9–12% (a 33–50% reduction, consistent with the case-study range cited in research), that alone removes rework delay from roughly 20–30 claims per month — and for claims that still require a second round, automated reminders (rather than manual emails) should compress the *time* to close that gap even when it occurs. Conservatively, this justifies treating 4–5 of the 8 document-collection days as addressable, consistent with the table above.

QA and approval are the second lever, and the one most directly tied to the "headcount-neutral" and "regulatory audit trail" constraints. Digitizing the checklist (removing paper scan-and-file steps) and applying risk-tiering — full manual QA only for claims above a value/complexity threshold or with flagged anomalies, lighter sampling-based review for the 85%+ of claims that historically pass clean — can plausibly cut the QA stage from 3 days to 1–1.5 days without reducing audit rigor, since a digitized checklist with system-logged timestamps generally strengthens the audit trail rather than weakening it. Manager approval, currently sequential after QA, can be partially parallelized for low-risk claims (pre-cleared by risk-tiered QA) shaving another 0.5–1 day.

Assignment queue wait is the smallest absolute number but the highest-leverage, lowest-risk fix. At 1.5 days, it is a pure coordination failure with no regulatory, judgment, or rework complexity attached — it is simply the time between a claim landing in Guidewire and a human

noticing it. Replacing manual round-robin with system-driven routing (using Guidewire's existing workflow/business rules capability, per the "no new core system" constraint) can plausibly compress this to near-zero, recovering close to the full 1.5 days for minimal implementation risk. This makes it the natural "quick win" to sequence first in the implementation playbook, even though its absolute contribution to the 21-to-10 gap is modest.

Prioritization Logic

Applying a Theory of Constraints lens, the redesign should not spread effort evenly across all six stages — it should attack the binding constraint first. Document collection is the binding constraint by every measure available: it is the single largest stage (38% of cycle time), it has a clearly identified root cause (no automated follow-up, a lesson already learned and lost once from the 2022 pilot), and it compounds — the 18% rework rate doesn't just add 4-5 days to those claims, it also delays entry into every downstream stage, meaning fixes to QA or approval speed cannot be fully realized until the document bottleneck is addressed first. This mirrors the classic Theory of Constraints failure mode the research flagged: optimizing QA or approval steps while document collection remains unfixed would improve local metrics without moving the end-to-end average meaningfully, since claims would simply wait longer at the true bottleneck.

Prioritization should therefore run in three tiers:

- **Tier 1 — Document collection automation and standardization.** Highest impact (4.5–5.5 recoverable days), moderate implementation complexity, directly addresses the client's own diagnosis ("document collection is clearly the biggest bottleneck") and directly repairs the failure mode of the 2022 pilot by making the fix systemic rather than person-dependent.
- **Tier 2 — Assignment automation and QA digitization/risk-tiering.** Combined recoverable range of roughly 2.5–3.3 days, lower implementation complexity than Tier 1 (assignment logic can likely use existing Guidewire workflow rules; QA digitization is a checklist/tooling change, not a judgment change — consistent with Coastal's stated appetite for automating "scheduling/reminders and status tracking" while preserving human judgment in assessment/approval).
- **Tier 3 — Approval parallelization and decision-communication templating.** Smallest recoverable range (0.8–1.5 days) but near-zero regulatory risk and easy to sequence last, since it depends on risk-tiering logic established in Tier 2.

This sequencing also directly respects the stated constraints: nothing in Tier 1 or Tier 2 touches Guidewire's core system (satisfying "we can't replace our core claims system this year"), nothing removes human judgment from assessment or approval (satisfying the stated automation appetite), and the risk-tiering approach explicitly preserves — and arguably strengthens — the audit trail regulators require, since digitized, timestamped checklists are more defensible under audit than scanned paper.

Expected-Impact Estimate

Combining the tiered recoverable ranges: Tier 1 (4.5–5.5 days) + Tier 2 (2.5–3.3 days) + Tier 3 (0.8–1.5 days) yields a total plausible recovery of **7.8–10.3 days**, bringing average cycle time to a range of **10.7–13.2 days** if only midpoint assumptions hold, or as low as **~10–11 days** at the upper end of the recovery estimates. This suggests the under-10-day target is achievable but not guaranteed by process redesign alone at every assumption level — it is realistic if Tier 1's document-collection fix performs at the higher end of the estimated range (consistent with the 30–50% second-request reduction seen in comparable case studies) and if QA risk-tiering is applied assertively rather than conservatively. This is an important honest finding for the client:

hitting "under 10 days" reliably, especially during the 40–60% hurricane-season volume surge, will likely require the redesign to perform at the upper end of its estimated impact, reinforcing why the to-be design and implementation playbook must treat document collection as the primary design target rather than one improvement among several.

On cost, no headcount or budget changes are assumed. At \$215K/month fully loaded cost, the redesign is cost-neutral by design — the return is expressed entirely in cycle time, capacity headroom (the same 22 processors handling the same ~340 claims/month, or the seasonal 475–545 claim surge, in less elapsed time per claim), and client-retention risk mitigation for the two contracts currently threatened by delay-driven complaints.

Process mapping (as-is)

Process Mapping (As-Is)

Purpose and Method

This section documents the current claim lifecycle exactly as it operates today, step by step, to create the baseline against which the bottleneck analysis and to-be design (covered in subsequent chapters) are built. The map is constructed directly from Coastal's own process description, cycle-time breakdown, and systems inventory provided in intake — no steps have been added or inferred beyond what the client described, though handoff points and control gaps not explicitly named by the client are called out where they are structurally implied by the data.

The map covers the full lifecycle boundary as defined by the client: **start point** — claim reported by policyholder via phone, email, or portal; **end point** — approval or denial decision communicated to the policyholder. Eight discrete steps sit between these boundaries, spanning four functional groups (client service, claims processors, QA/compliance, claims managers) and three disconnected systems (Guidewire, Outlook, and an unofficial shared Excel tracker).

As-Is Process Map (Step-by-Step)

#	Step	Owner	System(s) Used	Trigger	Handoff To	Avg. Duration	Documentation/ Audit Trail
1	Claim reported (phone/ email/ portal)	Client service rep	Guidewire (manual entry)	Policyholder contact	Team lead / queue	Same day (entry)	Guidewire record created — compliant timestamp
2	Assigned to processor	Team lead (manual round-robin)	Shared Excel tracker (not connected to Guidewire)	Team lead checks tracker, assigns manually	Assigned processor	~1.5 days (queue wait)	Assignment recorded in Excel only — not in Guidewire , creating an audit-trail gap

3	Documents requested from policyholder	Processor	Outlook (email, no automation)	Processor opens file	Policyholder	Same day request; 8 days elapsed incl. wait/rework	Email thread only — no system-logged reminder cadence
4	Documents reviewed on receipt	Processor	Outlook + Guidewire (manual upload)	Policyholder reply	Processor (or back to step 3 if incomplete)	Variable — 18% of claims loop back	Incomplete-submission rate not systematically logged, only anecdotally known
5	Assessment and cost estimate drafted	Processor	Guidewire	Documents complete	QA/compliance	~4 days	Draft stored in Guidewire
6	QA/compliance review	QA staff	Paper checklist, scanned into file	Processor submits draft	Claims manager	~3 days	Paper checklist scanned and attached — audit trail preserved but slow, sequential
7	Manager approval or escalation	Claims manager	Guidewire + paper sign-off reference	QA sign-off received	Decision letter drafting	~2 days	Guidewire approval record
8	Decision letter drafted and sent	Processor or client service	Outlook / mail	Manager approval	Policyholder (end point)	~1 day	Email/mail record, not always logged back into Guidewire

Swimlane View: Who Touches the Claim, and When

Mapping ownership by functional group makes visible a structural issue not explicit in the client's step list but implied by it: **the same processor is responsible for both chasing documents (steps 3–4) and drafting the technical assessment (step 5)** — two fundamentally different work types (administrative follow-up vs. analytical judgment) are bundled into one role with no differentiation in tooling or prioritization. This is consistent with why the 2022 dedicated document-chase pilot "helped initially" — it separated the administrative task from the analytical one — and why its benefit evaporated once the role was eliminated and the two tasks were re-merged into the processor's queue.

Functional Group	Headcount	Steps Owned	% of Process Steps Owned	Systems Touched
Client service (9)	9	1, 8 (partial)	25%	Guidewire, Outlook

Claims processors (22)	22	2 (initiated by team lead), 3, 4, 5, 8 (partial)	63%	Guidewire, Outlook, Excel tracker
QA/compliance (6)	6	6	13%	Paper checklist, scanned file
Leadership/ops (4, incl. claims managers)	4	7	13%	Guidewire

Claims processors carry the largest share of process steps (5 of 8) and are the group most exposed to the two largest sources of delay: document collection (step 3–4) and drafting (step 5). This concentration matters directly for the "headcount-neutral" constraint — any redesign that reduces the administrative burden on processors (document chasing) without touching their analytical work (assessment drafting) frees processor capacity without altering headcount or role scope, which is the lever the to-be design should prioritize.

System Fragmentation Map

A second dimension of the as-is state — separate from the step sequence — is the disconnect between the three systems in active use, none of which talk to each other:

System	Purpose	Connected to Guidewire?	Gap Created
Guidewire	System of record for claim data, assessments, approvals	N/A (core system)	Does not capture assignment status or document-request status in real time
Outlook	All policyholder and internal communication	No	No automated reminders; document-request status lives only in individual inboxes, invisible to team leads or managers
Shared Excel tracker	Assignment status, kept by "some team leads"	No	Manually maintained, inconsistently used ("some" team leads, implying non-uniform adoption), creates the 1.5-day assignment queue gap and an audit-trail blind spot since assignment timing isn't in Guidewire
Paper QA checklist	QA sign-off and audit documentation	No (scanned in after the fact)	Sequential, single-copy artifact — cannot be reviewed in parallel, cannot be searched or reported on in aggregate

This fragmentation is the structural explanation for why Coastal's own six-stage cycle-time breakdown ($1.5 + 8 + 4 + 3 + 2 + 1 = 19.5$ days) does not fully reconcile with the stated 21-day average. The missing ~1.5 days almost certainly lives in the handoffs *between* these disconnected systems — the time between a claim landing in Guidewire and someone checking the Excel tracker; the time between an email reply landing in a processor's inbox and

them acting on it; the time between a scanned QA checklist being filed and a manager actually opening it. None of this transit time is currently measured because no single system owns the full lifecycle view.

Control Points and Where Rework Enters the Map

Two points in the as-is map function as quality gates, and both are sources of rework loops that extend cycle time beyond the linear step sequence shown above:

- **Step 3–4 loop (document completeness):** 18% of claims ("H61/month at current volume) fail the implicit "complete submission" check at step 4 and loop back to step 3, each cycle adding 4–5 days. There is no formal checklist enforced at step 3 to prevent this — the check happens only when the processor manually reviews what arrives.
- **Step 6 loop (QA sign-off):** 12% of claims ("H41/month) fail the QA checklist at step 6 and loop back to processor rework, each cycle adding 2–3 days. This loop is fully sequential — a claim cannot proceed to step 7 until it clears step 6, and there is no parallel or risk-based path for lower-complexity claims to bypass full manual review.

Combined, these two loops touch an estimated 30% of monthly claims (with some overlap likely, meaning the true share of claims experiencing at least one rework loop may be somewhat lower than the simple sum but still substantial), meaning roughly one in three claims moves through the map illustrated above with at least one additional backward loop not shown in the linear eight-step sequence — a materially different (and slower) map than the "happy path" the eight numbered steps suggest in isolation.

Compliance and Audit-Trail Checkpoints in the Current Map

Given the constraint that any redesign must preserve Florida/Georgia documented-timeline and audit-trail requirements, it is worth mapping precisely where audit evidence is currently captured, since this determines what must be preserved (or improved upon) in the to-be design:

Regulatory Requirement	Where Currently Evidenced	Medium	Risk in Current State
Initial contact within statutory window	Step 1 Guidewire entry timestamp	Digital	Low risk — reliably captured
Documented claim-handling timeline throughout	Step 1 and Step 7 Guidewire timestamps only	Digital (partial)	Medium risk — steps 2–6 timing lives in Excel/ Outlook/paper, not Guidewire, weakening the continuity of the documented timeline
QA sign-off / review evidence	Step 6 paper checklist, scanned	Paper -> scanned image	Medium risk — legible and filed, but slow and not system-searchable in aggregate for audit sampling
Final decision communicated within statutory window	Step 8 Guidewire approval + email/mail record	Digital + manual	Low-medium risk — decision date captured, but communication method (email/mail) not always logged back into Guidewire

This checkpoint map is a key input to the to-be design: the redesign does not need to invent new compliance controls, but it does need to close the gap where steps 2–6 timing currently lives outside Guidewire, since that gap is both a speed problem (untracked handoff time) and a latent audit-trail weakness (timeline continuity depends on Excel and email records that are not part of the system of record).

Summary of As-Is State

The current process is a linear eight-step sequence nominally averaging 19.5–21 days, but the map reveals it is more accurately understood as a **fragmented, three-system workflow with two structural rework loops touching roughly a third of claims**, where responsibility for both administrative document-chasing and analytical assessment work sits with the same 22-person processor pool, and where the only cross-step visibility tool (the Excel tracker) is informally maintained and disconnected from the system of record. This as-is baseline — not the idealized eight-step sequence alone — is the correct starting point for the bottleneck and cost analysis that follows.

Bottleneck & cost analysis

Bottleneck & Cost Analysis

This section translates the as-is process map into dollar terms, isolating exactly where Coastal's \$215K monthly claims-handling spend is being consumed by non-value-added work rather than claim assessment itself, and quantifying the cost of inaction if the two at-risk client contracts are not stabilized.

Cost-per-Stage Breakdown

Coastal's fully loaded monthly cost of \$215K against 22 processors and 6 QA/compliance staff produces an average fully loaded cost per processor of roughly \$8,240/month (\$98,900/year, consistent with the stated \$48K–\$62K salary band once benefits/overhead loading of roughly 55–65% is applied) and per QA staff member of roughly \$8,240–\$9,900/month (consistent with the \$55K–\$70K band). Applying these loaded rates against the time each functional group spends per claim, at 340 claims/month, produces a stage-level cost allocation:

Stage	Owner	Avg. Days	Loaded Cost Basis	Est. Cost per Claim	Est. Monthly Cost (340 claims)
Assignment (queue + team lead coordination)	Team lead / processor	1.5	Processor-equivalent, partial	~\$28	~\$9,500
Document collection (incl. rework)	Processor	8.0	Processor	~\$149	~\$50,700
Review/drafting	Processor	4.0	Processor	~\$75	~\$25,400
QA review	QA/compliance	3.0	QA	~\$71	~\$24,000
Manager approval	Leadership/ops	2.0	Manager-equivalent	~\$60	~\$20,400

Decision communication	Processor/client service	1.0	Blended	~\$19	~\$6,400
Total		19.5		~\$402	~\$136,400

*(Note: this stage allocation totals roughly \$136K against the full \$215K monthly cost — the remainder reflects client service intake work, leadership/ops overhead not directly tied to individual claims, and general administrative capacity that sits outside the eight-step lifecycle but within the department's total cost base. The \$402 estimate is therefore a *processing-labor* cost per claim, distinct from the fully loaded \$632 per-claim figure cited in the research benchmarking, which spreads the entire department cost — including idle/overhead capacity — across claim volume.)*

The critical finding from this breakdown is not the total, but the concentration: **document collection alone accounts for ~37% of processing-labor cost per claim (~\$149 of ~\$402)**, almost exactly mirroring its 38% share of cycle time. This confirms that document collection is not just the slowest stage — it is the most expensive one in labor terms, and it is expensive specifically because of rework, not because the task itself is inherently costly.

The Cost of Rework, Isolated

Coastal's two documented rework loops can be priced directly using the client's own rates and volumes.

Document rework (18% second-request rate): At 340 claims/month, roughly 61 claims require a second document round, each adding 4–5 days of elapsed time. Using the processor loaded rate of ~\$8,240/month (~\$374/business day), and assuming the processor spends a fraction of that added elapsed time in active follow-up work (estimated at 1.5–2 hours of active effort per rework cycle, not the full 4–5 elapsed days, since much of that time is passive waiting on the policyholder) — the direct labor cost of chasing incomplete documents is modest, roughly **\$3,000–\$4,000/month in active processor time**. The real cost is not labor hours; it is **capacity lock-up**: each of those 61 claims occupies a slot in the processor's active caseload for 4–5 extra elapsed days, reducing effective throughput capacity by an estimated **8–10% of total processor bandwidth** every month, before any hurricane-season surge is applied. This is the more important number for a headcount-neutral objective — the constraint isn't a wage cost, it's usable hours across the same 22 people.

QA rework (12% return-to-processor rate): At 340 claims/month, roughly 41 claims are sent back for rework each cycle, adding 2–3 days. This creates a second capacity drag on the same processor pool (since QA doesn't do the rework — it sends the claim back to the original processor), compounding the document-collection drag on the exact same 22 people, at the exact same time in the calendar (peaking during June–November when volume is already 40–60% higher).

Combined effect: Between the two loops, an estimated **102 of 340 monthly claims (30%)** re-enter processor queues at least once, and each re-entry costs elapsed calendar time without corresponding revenue upside (Coastal is paid claims-handling fees, not rework fees). This is the direct mechanical link between the bottleneck analysis and the client's objective: rework is not a quality issue in isolation — it is consuming the exact processor capacity Coastal needs to hit the under-10-day target without adding headcount.

Cost of the Current State vs. Cost of Delay (Client Risk)

Separate from internal processing cost, intake explicitly identifies that "rework and delays are

driving client complaints and threatening two of our largest client contracts." This section prices that risk in terms leadership can act on Monday morning.

Risk Factor	Basis	Estimated Exposure
Revenue at risk if 2 largest contracts are lost	Client states these are among the largest of ~\$6.5M total revenue; if each represents even 12–18% of total revenue (a reasonable range for "two largest" among a client base implied to be broader), combined exposure	~\$1.5M–\$2.3M annually
Monthly processing-labor cost currently lost to rework loops (capacity drag, not just direct wages)	8–10% of processor bandwidth tied up in document rework alone, applied against processor-stage monthly cost of ~\$50,700	~\$4,000–\$5,000/month in effective wasted capacity , before QA rework and before hurricane-season multiplier
Hurricane-season compounding	40–60% volume surge (June–Nov) against an already-strained 22-person processor pool with no headcount increase planned	Rework-driven capacity loss scales proportionally with volume — same 8–10% drag applied to 475–545 claims/month materially worsens cycle time exactly when client scrutiny (storm-related claims) is highest

(Assumption, labeled): The 12–18% per-contract revenue share is an estimate for illustrative sizing only, since intake does not disclose exact contract values; it should be replaced with actual contract figures by Coastal's leadership before this number is used externally. The directional point — that the revenue at risk from losing two major contracts is likely one to two orders of magnitude larger than any plausible internal automation or tooling investment required to fix document collection and QA sequencing — is the conclusion that matters for prioritization and budget authorization, regardless of the precise dollar figure.

Bottleneck Ranking: Impact vs. Effort to Fix

Bringing together cycle-time share, cost share, and rework frequency into a single prioritization view:

Bottleneck	Cycle-Time Share	Cost Share (processing-labor)	Rework Rate	Root Cause Class	Fix Complexity (given constraints)
Document collection	38%	37%	18%	No automation, no upfront checklist	Low–Medium (no core system change needed)
QA review	14%	18%	12%	Universal sequential paper process	Low (digitization + risk-tiering, no judgment removed)
Processor drafting	19%	19%	N/A (downstream of above)	Partly inherent, partly rework-inflated	Low (shrinks naturally once upstream fixed)
Manager approval	10%	15%	N/A	Sequential-only, no parallel path	Low (parallelize for risk-tiered low-complexity claims)

Assignment queue	7%	7%	N/A	Manual, disconnected tracker	Low (use existing Guidewire workflow rules)
Decision communication	5%	5%	N/A	Manual drafting, no templates	Low (templating only)

Document collection dominates on every axis simultaneously — it is the largest cost driver, the largest time driver, and the highest rework-rate stage — while also being one of the lower-complexity fixes given the constraint against replacing Guidewire, since the failed 2022 pilot already proved the *concept* works (a dedicated chase function reduced delays) and only failed on *durability* (it was person-dependent, not system-enforced). This is the clearest, lowest-risk, highest-return target for Monday-morning action.

Immediate Actions for Leadership (Week 1)

- **Assign an owner for document-collection redesign this week** — recommend the Ops leadership member (of the 4-person leadership/ops team) responsible for Guidewire administration, working with 2–3 processor "power users" to define a standard document checklist by claim type (property vs. casualty) within 5 business days. Effort: ~15–20 hours combined. Impact: prerequisite for any automation, addresses root cause of the 18% incomplete-submission rate directly.
- **Pull actual per-processor rework data from Guidewire and Outlook logs (manual audit, one-time)** to replace the 18%/12% estimates with claim-level detail — specifically, which processors, which claim types, and which document categories drive the second-request rate. Effort: ~1 QA/compliance staff member, 3–4 days. Impact: sharpens the checklist design and gives Coastal a real baseline to measure Tier 1 improvement against once implemented.
- **Instruct team leads to stop using the informal Excel tracker as of the next sprint** once an interim Guidewire-based assignment rule is configured (covered in the to-be design) — this closes the audit-trail gap identified in the as-is map immediately, at zero cost, ahead of any technology change.
- **Size the two at-risk contracts' actual revenue and current SLA terms** (replacing the labeled assumption above with real figures) so leadership can present a precise, defensible cost-of-delay figure to the executive team and, if needed, to the clients themselves as evidence of a committed remediation plan already underway.

To-be design

To-Be Design

The to-be design keeps every constraint intact — no Guidewire replacement, no headcount change, no removal of human judgment from assessment or approval, and full preservation (ideally strengthening) of the Florida/Georgia audit trail. It concentrates redesign effort where the bottleneck analysis says the leverage actually is: document collection first, then QA/assignment, then approval and communication. The design below specifies exactly what changes at each of the eight as-is steps, what tool or rule enables the change, who owns it, and what cycle-time recovery it is expected to produce — consistent with the 7.8–10.3 day recovery range established in the findings.

Design Principle: Fix the System, Not the Person

The single most important design constraint comes from Coastal's own history: the 2022 dedicated document-chase role "helped initially" but the gain evaporated the moment that person left, because the fix was a person, not a rule. Every Tier 1 and Tier 2 recommendation below is therefore built as a **system-triggered workflow**, not a role addition — so the fix survives staff turnover, seasonal hiring gaps, and hurricane-season surge without requiring a dedicated headcount line that leadership has explicitly ruled out.

Tier 1: Document Collection Redesign (Target recovery: 4.5–5.5 days)

New step 3 — Standardized, checklist-driven document request. Replace the current unstructured "processor requests documents by email" step with a claim-type-specific checklist (property vs. casualty, as flagged in the Week-1 action item) built into the initial request template. The checklist is triggered automatically the moment a claim is assigned (see assignment redesign below), not manually drafted per claim by the processor.

- **Who:** Ops leadership team member owning Guidewire administration, with 2–3 processor power users (already tasked in Week 1) finalize the checklist content within 5 business days; IT/ops configures it as a Guidewire document-template merge field or an Outlook Quick Part linked to claim type — no new system required.
- **How:** Checklist embedded in the outbound request email/portal message; portal-submitted claims get a real-time completeness indicator (a checkbox list the policyholder sees before submitting) if Coastal's Guidewire portal instance supports custom form fields — **flagged as a build item to confirm with Guidewire admin, not a new system purchase.**
- **Effort:** 20–30 hours to build and test both claim-type templates; no ongoing labor cost once live.
- **Expected impact:** Reduces the second-request rate from 18% toward the 9–12% range cited in comparable case studies, cutting the number of affected claims from ~61/month to ~30–40/month.

New step 3b — Automated reminder cadence (the direct fix to the 2022 pilot failure).

Configure a scheduled reminder sequence — Day 2, Day 5, and Day 7 after initial request if documents remain incomplete — sent automatically via Outlook mail-merge rules or a lightweight scheduling add-in (e.g., a Power Automate flow triggered off a Guidewire status field or a shared calendar/task system), rather than relying on a processor remembering to follow up.

- **Who:** IT/ops sets up the automation rule (one-time build); no staff member "owns" ongoing execution — that is the point.
- **How:** Trigger = claim status remains "documents pending" in Guidewire past X days. Action = automated email to policyholder, cc processor, with escalation to a status flag visible on the assignment dashboard (see below) if unresolved by Day 10.
- **Effort:** 15–25 hours build time (falls within Coastal's stated automation appetite — "open to automation for scheduling/reminders and status tracking").
- **Expected impact:** For claims that do require a second round, this compresses the time to close the gap from an unmanaged multi-day email thread to a structured, system-enforced 2–3 day resolution window — recovering an estimated 1.5–2 days per affected claim, on top of the reduction in *how many* claims are affected.
- **Compliance note:** Every automated reminder is logged in Guidewire (or the connected automation tool) with a timestamp, which *improves* the audit trail relative to the current unlogged Outlook thread — directly satisfying the constraint that redesign must preserve documented timelines.

Combined Tier 1 effect: Fewer claims enter the rework loop, and those that do resolve faster — the two effects compound to bring the 8-day document-collection stage down to an estimated 2.5–3.5 days.

Tier 2: Assignment and QA Redesign (Target recovery: 2.5–3.3 days)

New step 2 — System-driven assignment, retiring the Excel tracker. Replace manual round-robin assignment with a Guidewire workflow rule (using existing business-rules/queue functionality already licensed, not a new purchase) that auto-assigns new claims based on current open caseload per processor, pulled directly from Guidewire's own claim-count field.

- **Who:** Ops leadership member owning Guidewire admin configures the rule; team leads retain override authority for complex/VIP claims only.
- **How:** Rule logic: assign to the processor with the lowest current open-claim count within the relevant line of business, capped by a maximum caseload threshold to prevent overload during surge periods. Excel tracker is decommissioned the same week (per the Week-1 immediate action already issued).
- **Effort:** 10–15 hours configuration and testing.
- **Expected impact:** Recovers 1.0–1.3 of the 1.5-day queue wait, and closes the audit-trail gap where assignment timing previously lived only in an unofficial spreadsheet.

New step 6 — Digitized, risk-tiered QA. Replace the paper checklist with a digital form (Guidewire-native form or a simple connected e-form tool) that time-stamps every field and sign-off automatically. Simultaneously, introduce a two-tier QA path:

QA Tier	Trigger Criteria	Review Type	Target Turnaround
Full manual review	Claim value above a defined threshold (e.g., top 20–25% by dollar amount), any anomaly flag from processor, any claim type with historical error concentration (identified via the Week-1 data audit)	Full sequential QA, same rigor as today	1.5 days
Sampled/expedited review	All other claims (estimated 75–80% of volume)	Digital checklist self-certified by processor, QA reviews a random 15–20% sample weekly for calibration, remainder pass through with system-logged sign-off	0.5 day

- **Who:** QA/compliance manager (one of the 6-person team) defines the value/complexity threshold and anomaly-flag criteria within 2 weeks, in consultation with compliance counsel to confirm the risk-tiering approach still satisfies Florida/Georgia audit expectations — **this is the one design element requiring explicit regulatory sign-off before go-live**, since it changes *how* QA rigor is applied, even though it does not remove human review from any claim entirely.
- **How:** Build the digital checklist as a Guidewire attachment form or connected tool; configure the value/complexity flag as an automatic field pulled from the processor's draft assessment (step 5).
- **Effort:** 25–35 hours to build the digital checklist and threshold logic; ~1 day of QA manager and compliance review time for the regulatory sign-off.

- **Expected impact:** Blended average QA turnaround falls from 3 days to approximately 0.7–1 day (weighted: 20% of claims at 1.5 days, 80% at 0.5 days), recovering 2.0–2.3 days — while the 12% error-catch function is preserved via full review on the highest-risk tier and calibration sampling on the rest, addressing the root cause (universal sequential gate) without eliminating the control.

Tier 3: Approval and Communication Redesign (Target recovery: 0.8–1.5 days)

New step 7 — Parallelized approval for risk-tiered low-complexity claims. For claims that pass the expedited QA tier with no flags, route manager approval as a same-day digital sign-off (Guidewire approval workflow, already in use) rather than a sequential 2-day queue behind full QA. Full manual approval sequencing is retained for the full-review tier.

- **Who:** Claims manager(s) among the 4-person leadership/ops team; no new approval authority required, just a faster-triggered workflow for the lower-risk 75–80% of claims.
- **Effort:** Minimal — configuration of an existing Guidewire approval queue priority flag, ~5–8 hours.
- **Expected impact:** Recovers 0.5–1 day for the majority of claims; high-complexity claims retain the full 2-day approval window, so regulatory scrutiny on the cases that most need it is unchanged.

New step 8 — Templated decision communication. Replace freeform decision-letter drafting with pre-approved, regulator-reviewed templates (approval and denial versions, by claim type) that auto-populate from Guidewire fields, sent via the existing email/portal channel and automatically logged back into Guidewire (closing the current gap where communication method "not always logged back").

- **Who:** Client service team lead, with QA/compliance sign-off on template language for regulatory compliance (one-time review).
- **Effort:** 10–15 hours to draft and approve templates.
- **Expected impact:** Recovers 0.3–0.5 days and closes an existing audit-trail gap at zero ongoing cost.

To-Be Cycle Time Summary

Stage	As-Is (days)	To-Be (days)	Recovery
Assignment	1.5	0.2–0.5	1.0–1.3
Document collection	8.0	2.5–3.5	4.5–5.5
Processor drafting	4.0	3.0–3.5	0.5–1.0
QA review	3.0	0.7–1.0	2.0–2.3
Manager approval	2.0	1.0–1.5	0.5–1.0
Decision communication	1.0	0.5–0.7	0.3–0.5
Total	19.5–21	8.9–11.7	7.8–10.3

This lands the to-be design within reach of the sub-10-day target at the midpoint-to-upper end of the estimate range, with the explicit caveat carried forward from the findings: hitting the target reliably during the June–November surge requires the QA risk-tiering thresholds to be applied assertively (not conservatively) and the document checklist to hit at least the middle of its

projected 9–12% second-request range.

What Stays the Same (By Design)

To preempt any impression that this is a system replacement or a judgment-removal exercise: Guidewire remains the system of record with no core replacement; every processor keeps their current role scope, drafting assessments exactly as before; every claims manager retains sign-off authority on every claim; QA staff review 100% of high-value/flagged claims in full, and calibrate against a live weekly sample of the rest — no claim is ever fully unreviewed. The only things eliminated are the unofficial Excel tracker, the paper checklist, and the unmanaged manual email chase — the three elements the as-is map identified as disconnected from the system of record and responsible for the audit-trail gaps and durability failure of the 2022 pilot.

Implementation playbook

Implementation Playbook

This playbook sequences the to-be design into a phased rollout that a 41-person organization can execute with existing staff, existing systems, and no new budget line for headcount. It is built around the same three-tier prioritization established in the findings — document collection first, assignment/QA second, approval/communication third — because that sequencing is what makes the 7.8–10.3 day recovery achievable rather than theoretical. Each phase below specifies the owner, the concrete task, the effort required, the dependency it sits behind, and the metric that proves it worked before moving to the next phase.

Governance Structure for the Rollout

Before any phase begins, Coastal needs a named, small governance group — not a new department, but a designated set of existing people with explicit decision rights for the 12-week rollout window.

Role	Who Fills It	Responsibility	Time Commitment
Sponsor	One of the 4 leadership/ops members (recommend whoever owns client-contract relationships, given the stated risk to two major contracts)	Removes blockers, approves threshold/policy decisions, reports weekly to the other 3 leadership members	2–3 hrs/week
Process Owner	Ops leadership member who owns Guidewire administration (already tasked with Tier 1 items in Week 1)	Owns build sequencing, coordinates IT/ops configuration, tracks metrics	8–10 hrs/week during build phases
QA/Compliance Lead	One of the 6 QA/compliance staff, designated as compliance liaison	Defines risk-tiering thresholds, obtains counsel sign-off, audits sample calibration	4–6 hrs/week
Processor Champions	2–3 processors (rotating or fixed)	Pilot-test new checklist and reminder flows, surface friction before full rollout	2–3 hrs/week during their pilot weeks

This group should meet weekly for 30 minutes for the first 12 weeks, then move to biweekly once Tier 1 and Tier 2 are stable. No new hire, no reassignment of existing headcount away from claims processing — the roles above are additive responsibilities layered onto existing jobs, consistent with the headcount-neutral constraint.

Phase 1 (Weeks 1–3): Document Collection — Build and Pilot

This phase directly operationalizes Tier 1 and should start immediately, since it carries the largest single recovery (4.5–5.5 days) and is already partially underway via the Week-1 actions identified in the bottleneck analysis.

Task	Owner	Effort	Dependency	Success Metric
Finalize claim-type document checklists (property, casualty)	Process Owner + 2–3 processor champions	20–30 hrs, 5 business days	Week-1 data audit results	Two approved checklist templates, signed off by QA lead
Build automated reminder rule (Day 2/5/7 cadence) triggered off Guidewire "documents pending" status	Process Owner + IT/ops	15–25 hrs	Checklist finalized	Rule live in test environment, verified against a dummy claim record
Decommission Excel tracker for assignment (interim manual rule until Phase 2 automation live)	Team leads	2 hrs communication/training	None	Tracker unused for 2 consecutive weeks, confirmed by spot-check
Pilot checklist + reminder flow on 20% of new claims (roughly 68 claims across 3 weeks)	Processor champions	Built into normal workload	Reminder rule live	Second-request rate on piloted claims tracked weekly against 18% baseline

Expected impact by end of Phase 1: second-request rate on piloted claims trending toward 9–14% (partial progress toward the 9–12% target, since full calibration typically takes 4–6 weeks of data). **Go/no-go gate:** if piloted claims show no improvement in second-request rate after 3 weeks, the Process Owner and QA lead must re-diagnose checklist content (likely cause: checklist doesn't match actual document gaps identified in the Week-1 audit) before expanding to full volume — do not scale a fix that isn't measurably working.

Phase 2 (Weeks 3–6): Assignment Automation and QA Digitization

Phase 2 can begin in parallel with the tail end of Phase 1 since it touches a different part of the workflow, but full rollout should wait until the Phase 1 pilot shows directional improvement.

Task	Owner	Effort	Dependency	Success Metric
------	-------	--------	------------	----------------

Configure Guidewire workflow rule for auto-assignment by open caseload	Process Owner + IT/ ops	10–15 hrs	Excel tracker decommissioned	Zero unassigned claims sitting >4 hours after logging, tracked via Guidewire queue report
Define risk-tiering thresholds (claim value, complexity flags, historical error concentration by claim type)	QA/Compliance Lead + compliance counsel	1 day of combined review time	Week-1 data audit (processor/claim-type error patterns)	Written, counsel-approved tiering policy document
Build digital QA checklist (Guidewire attachment form or connected e-form) with automatic timestamp logging	Process Owner + IT/ ops	25–35 hrs	Risk-tiering thresholds defined	Digital form live, tested against 10 sample claims, audit trail verified by QA lead
Train QA staff (6 people) on tiered review process and weekly sampling calibration method	QA/Compliance Lead	4 hrs training session + 2 hrs Q&A	Digital checklist built	All 6 QA staff certified on new process (simple sign-off checklist)

Expected impact by end of Phase 2: assignment queue wait compressed to under half a day for all new claims; QA turnaround beginning to shift toward the blended 0.7–1.0 day target as the expedited tier processes its first full batch of claims. **Critical risk to monitor:** the risk-tiering policy is the one element requiring explicit regulatory sign-off — do not go live on the expedited QA tier until compliance counsel has confirmed in writing that sampling-based review for the lower tier satisfies Florida/Georgia audit expectations. This is a hard gate, not a recommendation.

Phase 3 (Weeks 6–9): Approval Parallelization and Decision Templating

Task	Owner	Effort	Dependency	Success Metric
Configure Guidewire approval queue priority flag for expedited-tier claims	Process Owner + claims managers	5–8 hrs	Risk-tiering live and stable for 2+ weeks	Expedited-tier claims reach manager approval same-day as QA sign-off
Draft and approve decision-letter templates (approval/denial, by claim type)	Client service team lead + QA/ compliance sign-off	10–15 hrs	None (can run in parallel with Phase 2)	Templates approved by compliance, auto-populate correctly from Guidewire fields in test run
Configure auto-logging of communication method (email/mail) back into Guidewire	IT/ops	5–10 hrs	Templates finalized	100% of decision communications logged in Guidewire for a 2-week test sample

Expected impact by end of Phase 3: manager approval stage compressed for the 75–80% expedited tier; decision communication step reduced to template-driven turnaround. At this

point all three tiers of the to-be design are live.

Phase 4 (Weeks 9–12): Full-Volume Rollout and Measurement

Task	Owner	Effort	Dependency	Success Metric
Expand document checklist and reminder automation to 100% of new claims	Process Owner	Rollout communication, ~4 hrs	Phase 1 pilot showed directional improvement	Second-request rate at full volume tracked weekly
Run first full end-to-end cycle-time measurement (intake to decision) on a 4-week rolling basis	Process Owner + QA lead	Report generation, ~3 hrs/week	All phases live	Average cycle time reported against 21-day baseline and 10-day target
Conduct first regulatory audit-trail spot-check (simulate a state audit request)	QA/Compliance Lead	1 day	Digital QA checklist and Guidewire logging live for 4+ weeks	Successful reconstruction of full claim timeline from Guidewire alone, no reliance on Excel/ paper artifacts

Sequencing Rationale and What Not to Do

Do not attempt to launch all three tiers simultaneously in week one. The reasoning is direct: Tier 2's QA risk-tiering depends on having real second-request and error-pattern data from the Week-1 audit and Phase 1 pilot to set thresholds correctly — setting them blind risks either under-tiering (too many claims still going through full manual QA, no time savings) or over-tiering (too aggressive a cut to manual review, a genuine compliance exposure). Similarly, Tier 3's approval parallelization depends on Tier 2's risk tiers being stable and trusted; parallelizing approval on top of an unproven QA tier compounds risk rather than isolating it.

Effort and Resourcing Summary

Tier	Total One-Time Build Effort	Ongoing Owner Time Post-Launch	New Headcount Required
Tier 1 (document collection)	35–55 hrs	~2 hrs/week (monitoring reminder rule performance)	None
Tier 2 (assignment + QA)	39–56 hrs + compliance review	~4–6 hrs/week (QA lead calibration sampling)	None
Tier 3 (approval + communication)	20–33 hrs	~1 hr/week	None
Total	94–144 hrs across 12 weeks	~7–9 hrs/week ongoing	None

Total one-time build effort of roughly 100–150 hours, spread across the Process Owner, IT/ops, QA lead, and a handful of processor champions over 12 weeks, is well within capacity for a 41-person organization without pulling anyone off claims work full-time — the heaviest individual load (Process Owner at 8–10 hrs/week during build phases) is comparable to running a part-

time project alongside a normal role.

Metrics Dashboard for Weekly Governance Review

The governance group should track four numbers weekly, all pullable from existing Guidewire reporting once the redesign is live — no new BI tool required:

- **Second-request rate** (target: 18% -> 9–12%)
- **Average assignment queue wait** (target: 1.5 days -> under 0.5 days)
- **Blended QA turnaround** (target: 3 days -> 0.7–1.0 days)
- **End-to-end average cycle time** (target: 21 days -> under 10 days), tracked as a 4-week rolling average to smooth week-to-week claim-mix noise

Hurricane-Season Readiness Check (Target Completion: End of May)

Given the June–November 40–60% volume surge, the entire playbook above should be fully live and stable by end of May — not June. This gives one full month of stable operation at normal volume before the surge hits, so the automated reminder cadence, risk-tiered QA, and parallel approval paths are proven under normal load before being tested under peak load. If Phase 4 measurement in weeks 9–12 (ending roughly mid-to-late March if Phase 1 starts immediately) shows the cycle time is not yet trending toward single digits, leadership has a built-in decision point in April to intensify Tier 1 checklist refinement or tighten QA tiering thresholds before the season begins — rather than discovering the gap in July under surge conditions.

Risks & Mitigations

Every recommendation in this playbook is designed to work within Coastal's real constraints — no new systems, no new headcount, no removal of human judgment — but that same design discipline is precisely what creates the risk surface below. A redesign built entirely on existing tools and existing staff succeeds or fails based on adoption discipline, calibration accuracy, and organizational follow-through, not on technology risk. The risks below are drawn directly from Coastal's own history (the 2022 pilot failure), its own data (the 18%/12% rework rates), and its own stated tensions (speed vs. regulatory caution, headcount-neutral vs. capacity-constrained), not from generic implementation-risk checklists.

The single largest risk is a repeat of the 2022 failure pattern, but at the process level instead of the role level. The prior document-chase pilot didn't fail because the idea was wrong — it failed because the fix depended on one person's ongoing effort and had no structural anchor once that person left. The current design is explicitly built to avoid this by making the reminder cadence system-triggered rather than person-owned. But a comparable failure mode remains possible: if the Process Owner (a single individual on the 4-person leadership/ops team) who configures and maintains the Guidewire workflow rules, reminder automation, and risk-tiering logic leaves or is reassigned during the 12-week build, the same fragility reappears one layer up the stack. Unlike the 2022 chase-person role, this risk is mitigated by documentation and cross-training, but only if that discipline is built in deliberately — it will not happen by default.

QA risk-tiering carries the highest regulatory severity, even though its likelihood is moderate. Setting thresholds too aggressively (routing too many claims to the expedited/sampled tier) creates genuine exposure under Florida and Georgia audit-trail requirements — a state examiner reconstructing a claim timeline and finding that a claim which should have received full manual review instead passed through sampling-based sign-off is a compliance

finding, not a process inefficiency. The playbook already builds in a hard gate (counsel sign-off before go-live), which is the correct mitigation, but the risk doesn't disappear at go-live — it recurs every time claim mix shifts, particularly during hurricane season when claim complexity and volume both spike simultaneously and the thresholds calibrated on calm-season data may not hold.

Processor adoption resistance is likely, moderate in severity, and specific to how the checklist and reminder automation change daily work. Processors currently own the entire document-chase relationship with policyholders personally — automating reminders removes a piece of that ownership and could be read as reducing their role, even though leadership's intent is the opposite (freeing their time for analytical work). Given that 63% of process steps currently sit with the 22-person processor pool, and that this group is the one whose behavior change the entire Tier 1 recovery depends on, tepid adoption of the checklist (using it inconsistently, or treating the automated reminders as optional rather than standard) would silently cap the achievable improvement well below the 4.5–5.5 day target without ever showing up as an explicit failure — claims would simply continue looping back at close to the current 18% rate.

Cost overrun risk is low in dollar terms but real in time-and-attention terms. The playbook estimates 94–144 build hours, which is a small dollar figure against a \$215K/month cost base, but the real currency at risk is Process Owner and QA Lead attention during the exact months (leading into hurricane season) when claims volume and organizational stress are already elevated. If Phase 1 or Phase 2 slips by even 3–4 weeks — plausible given that this work is additive to existing full-time roles — the "fully live by end of May" readiness target is missed, and the redesign gets tested under peak surge conditions rather than validated in advance.

Guidewire configuration limitations are a moderate-likelihood, moderate-severity risk specific to the "no core system replacement" constraint. The design assumes Guidewire's existing workflow/business-rules engine and portal can support auto-assignment by caseload, custom document-completeness fields, and digital QA form attachments with automatic timestamping. This has not been confirmed against Coastal's actual Guidewire license tier and configuration — it is flagged in the to-be design as something to "confirm with Guidewire admin," which means there is real risk that one or more of these capabilities requires a paid add-on module, a support ticket with vendor lead time, or simply isn't available in Coastal's current instance, forcing a workaround (e.g., a connected but separate scheduling tool) that reintroduces some of the system fragmentation the redesign is trying to eliminate.

Competitive/client response risk is asymmetric: the risk of inaction is far larger than the risk of the redesign itself, but the redesign's benefits may not be visible to clients fast enough. The two threatened contracts are the real business risk driving this whole engagement — if leadership doesn't communicate a credible remediation timeline to those two clients now, client relationship risk continues to compound during the 12-week build period even as the internal fix takes hold, since clients will not see cycle-time improvement until Phase 4 metrics stabilize.

Risk	Likelihood	Impact	Mitigation
------	------------	--------	------------

Repeat of 2022 failure pattern — fix depends on one Process Owner, unravels on departure/reassignment	Medium	High	Document every workflow rule, threshold, and automation build in a shared runbook from week 1; cross-train the QA Lead and one processor champion on basic rule maintenance by Phase 3
QA risk-tiering thresholds miscalibrated, creating audit-trail exposure	Medium	High	Hard go-live gate requiring written compliance counsel sign-off (already in playbook); re-validate thresholds quarterly and explicitly before/after hurricane season using fresh claim-mix data
Processor adoption resistance — checklist/reminders used inconsistently	Medium-High	Medium	Frame rollout explicitly as reducing processor administrative burden, not replacing judgment; use processor champions as peer messengers; track checklist-usage rate as a leading indicator, not just the lagging second-request rate
Build timeline slips past end-of-May hurricane-season deadline	Medium	Medium	Build a hard April checkpoint (already in playbook) with pre-committed fallback: if behind schedule, prioritize Tier 1 completion alone over full Tier 2/3 rollout, since Tier 1 carries the largest recovery
Guidewire configuration can't natively support planned automation/forms	Medium	Medium	Confirm licensing/configuration capability with Guidewire admin or vendor support in week 1, before finalizing build plans, not after; identify low-cost workaround (e.g., connected scheduling add-in) as contingency
Client relationship risk continues during 12-week build before metrics improve	High	High	Proactively communicate a specific remediation timeline and interim milestones to the two at-risk clients now, independent of internal build progress

QA/processor rework loop persists at near-current rates if checklist content misses actual gaps

Medium

Medium

Recalibrate checklist content using the Week-1 Guidewire/Outlook audit findings rather than assumptions, and treat the Phase 1 pilot go/no-go gate as a hard checkpoint before scaling |

Beyond the items already flagged inline through the playbook, five risk categories deserve explicit attention because they are the ones most likely to determine whether Coastal actually closes the gap from 21 days to under 10, rather than settling for a partial win that leaves the two at-risk contracts exposed.

Adoption risk is the highest-probability failure mode, and it is not hypothetical — Coastal has already lived it. The 2022 document-chase pilot proves that a good idea can work initially and still fail to stick once the person or habit driving it disappears. The redesign is built to be system-triggered rather than person-dependent specifically to avoid repeating that failure, but a new risk emerges in its place: 22 processors and 6 QA staff have to actually change daily habits — using a new checklist instead of freeform emails, trusting a system-assigned queue instead of checking with their team lead, self-certifying on the expedited QA tier instead of waiting for a paper sign-off. If processors quietly revert to old habits (drafting requests from memory instead of the template, or QA staff insisting on manually reviewing every claim "to be safe" regardless of tier), the redesign's mechanics will exist on paper without functioning in practice. This is moderately likely given the org's own history and moderately severe, since it would silently erode the recovery estimate without an obvious single point of failure to diagnose. The mitigation is behavioral, not technical: the weekly governance metrics (second-request rate, QA turnaround, assignment wait) should be reviewed by claim volume *per processor and per QA reviewer*, not just in aggregate, so uneven adoption is visible early and can be coached rather than discovered three months later in an aggregate number that looks "fine on average" while hiding a subset of staff who never changed behavior.

Regulatory/compliance risk is lower in probability but severe in consequence, and it concentrates entirely in the QA risk-tiering decision. Everything else in the redesign is either purely internal coordination (assignment, reminders) or preserves full existing rigor (templated communication, full QA on high-risk claims). The one genuine compliance exposure is applying sampling-based review to 75–80% of claims instead of full manual review, and if Florida or Georgia regulators later interpret that as a weakening of the "documented claim-handling timeline" requirement — even though the underlying timestamps are arguably stronger under the digital system than the current paper process — Coastal could face an audit finding or, worse, a client carrier questioning whether Coastal's process meets contractual claims-handling standards. This risk is why the playbook already treats compliance counsel sign-off as a hard gate before the expedited QA tier goes live, not a recommendation; that gate should not be compressed under schedule pressure even if Phase 2 is running behind heading into hurricane season. If counsel raises concerns, the fallback is to set the risk-tiering threshold more conservatively (e.g., top 35–40% of claims by value/complexity go to full QA rather than top 20–25%), which reduces the QA turnaround improvement but preserves the entire program rather than risking a compliance finding that could jeopardize the client contracts the whole initiative is meant to protect.

Cost overrun risk is low in this design specifically because the plan avoids the two things that normally cause claims-ops projects to blow budget — new core systems and new headcount. The total build effort (94–144 hours over 12 weeks) is modest, and most of it uses tools Coastal already licenses (Guidewire workflow rules, Outlook automation). The

realistic overrun scenario is not dollars but *hours* — if the Guidewire workflow-rule configuration for auto-assignment or the digital QA form turns out to require vendor professional-services support because Coastal's specific Guidewire instance doesn't expose the needed configuration options to internal admins, effort could run 2–3x higher than estimated and introduce an unplanned vendor cost. This is medium-low likelihood (Guidewire's business rules and workflow engine are generally configurable by trained admins) but should be tested early: the Process Owner should attempt the assignment-rule build in Phase 1, in parallel with the document checklist work, specifically to surface this risk in weeks 1–3 rather than weeks 6–9, when a vendor-dependency discovery would delay the higher-impact Tier 2 QA work behind it.

Data/systems limitation risk centers on the same fragmentation the as-is map identified.

The redesign assumes Guidewire can expose a reliable "documents pending" status field to trigger automated reminders, and can support a connected or native digital QA form with timestamp logging. If either assumption proves wrong — for instance, if Guidewire's claim-status fields are more coarse-grained than assumed and can't cleanly distinguish "awaiting documents" from other pending states — the automation will either misfire (sending reminders on claims that don't need them) or fail to trigger reliably (the exact failure mode of the 2022 pilot, just automated instead of human). This is medium likelihood given that the client's own intake noted Guidewire integration has never been attempted for these workflows, and its impact is high because it sits underneath both Tier 1 and Tier 2. The mitigation is to validate the specific field/trigger mechanics against a small batch of real claims in the first week of Phase 1's build — before finalizing checklist design — rather than assuming the system behaves as expected and discovering a gap only during the pilot.

Competitive response risk is minor here but worth naming given the stakes: if a competitor claims administrator is already offering the two at-risk carrier clients a faster or more transparent alternative, an internal 12-week improvement plan — even a well-executed one — may not retain the contracts if the client's patience has already run out. This is why the client-communication mitigation in the table above is listed as high-priority regardless of internal progress: Coastal's real near-term risk may be perception and timeline, not technical execution, and the two should be managed as separate tracks with separate owners.

90-Day Action Plan

90-Day Action Plan

This plan sequences the to-be design into three horizons, mirroring the tiered prioritization established earlier: document collection first (largest recovery, lowest complexity), assignment/QA second, approval/communication third. Every action below traces directly to a specific recommendation in the process mapping, bottleneck analysis, to-be design, implementation playbook, or risk mitigation sections — no new scope is introduced here. The plan targets full stability by end of May, ahead of the June hurricane-season surge, consistent with the readiness checkpoint already established.

Governance note: Before Day 1 activity begins, the Sponsor (a leadership/ops member) should confirm the Process Owner, QA/Compliance Lead, and 2–3 processor champions named in the implementation playbook, and schedule the weekly 30-minute governance review. This is a prerequisite, not a discrete action, and is not separately tracked below.

Days 0–30: Diagnose, Build Foundation, Launch Pilot

This horizon executes the Week-1 immediate actions from the bottleneck analysis and the

Phase 1 build from the implementation playbook — the highest-leverage tier, since document collection accounts for 38% of cycle time and 37% of processing-labor cost.

The first priority is closing the data gap: pulling actual per-processor, per-claim-type rework data from Guidewire and Outlook logs to replace estimates with specifics, since the checklist design and QA thresholds both depend on this. In parallel, the informal Excel tracker is retired immediately — at zero cost — closing an existing audit-trail gap identified in the as-is map. The claim-type document checklists are then built and piloted on 20% of new claims, with the automated reminder cadence (Day 2/5/7) configured to trigger off Guidewire status rather than depend on a person, directly addressing why the 2022 chase-person pilot didn't survive staff turnover. A Guidewire-capability check is run early and deliberately, per the risk mitigation for vendor-dependency discovery, so any configuration limitation surfaces in week one rather than week nine.

Days 31–60: Scale Document Fix, Automate Assignment, Digitize QA

This horizon executes Phase 2 of the playbook, contingent on the Phase 1 pilot showing directional improvement — the go/no-go gate specified in the playbook. If the pilot's second-request rate hasn't moved after three weeks, checklist content is re-diagnosed before scaling, rather than rolled out blind.

Assuming the gate is passed, the document checklist and reminder automation expand to 100% of new claims. Simultaneously, Guidewire's workflow rules are configured for auto-assignment by open caseload, eliminating the manual round-robin and its 1.5-day queue wait. The QA/Compliance Lead defines risk-tiering thresholds using the Phase 1 data audit and obtains written compliance counsel sign-off — a hard gate, not a formality, given that this is the one design element with genuine regulatory exposure. The digital QA checklist is built and tested in parallel, and all six QA staff are trained on the tiered review and weekly sampling-calibration method before go-live.

Days 61–90: Parallelize Approval, Template Communication, Validate End-to-End

This horizon executes Phase 3 and the early measurement activity from Phase 4, contingent on the risk-tiering approach being live and stable for at least two weeks — the dependency specified in the playbook to avoid compounding an unproven QA tier with a parallel approval change.

Manager approval is parallelized for the 75–80% expedited-tier claims via a Guidewire approval-queue priority flag, while full sequential approval is preserved for the higher-risk tier. Decision-letter templates (approval and denial, by claim type) are drafted, compliance-reviewed, and configured to auto-populate from Guidewire fields, with communication method auto-logged back into the system — closing an existing audit-trail gap. By day 90, the first full end-to-end cycle-time measurement runs on a four-week rolling basis, and a simulated regulatory audit spot-check confirms the full claim timeline can be reconstructed from Guidewire alone. This positions Coastal for the April decision checkpoint and the end-of-May hurricane-season readiness deadline established in the playbook.

Days	Action	Owner	Effort	Milestone (Done =)
0–5	Audit Guidewire/ Outlook logs for actual rework rates by processor and claim type	QA/Compliance Lead	~3–4 days	Baseline data report replacing 18%/12% estimates with claim-level detail

0–5	Retire informal Excel assignment tracker; communicate to team leads	Team leads	~2 hrs	Zero use of tracker for 2 consecutive weeks (spot-checked)
0–5	Confirm Guidewire licensing/config supports auto-assignment rules and digital form attachments	Process Owner + Guidewire admin/vendor	~4–6 hrs	Written confirmation or identified workaround for each planned automation
1–5	Finalize claim-type document checklists (property, casualty)	Process Owner + 2–3 processor champions	20–30 hrs	Two checklists approved and signed off by QA Lead
5–10	Build automated reminder rule (Day 2/5/7 cadence) triggered off Guidewire status	Process Owner + IT/ops	15–25 hrs	Rule verified live in test environment on a dummy claim
10–30	Pilot checklist + reminders on 20% of new claims (~68 claims)	Processor champions	Built into normal workload	Second-request rate on piloted claims tracked weekly vs. 18% baseline
25–30	Go/no-go review of Phase 1 pilot results	Process Owner + QA Lead	~2 hrs	Documented decision to scale or re-diagnose checklist content
31–40	Expand checklist + reminder automation to 100% of new claims	Process Owner	~4 hrs rollout comms	Full-volume second-request rate tracked weekly
31–40	Configure Guidewire auto-assignment by open caseload	Process Owner + IT/ops	10–15 hrs	Zero claims unassigned >4 hours, per Guidewire queue report
35–45	Define risk-tiering thresholds (value, complexity, error-pattern flags)	QA/Compliance Lead + compliance counsel	~1 day combined	Written, counsel-approved tiering policy
40–55	Build digital QA checklist with automatic timestamp logging	Process Owner + IT/ops	25–35 hrs	Digital form tested on 10 sample claims; audit trail verified
50–60	Train 6 QA staff on tiered review and weekly sampling calibration	QA/Compliance Lead	~6 hrs total	All QA staff certified on new process

60–65	Configure Guidewire approval-queue priority flag for expedited-tier claims	Process Owner + claims managers	5–8 hrs	Expedited-tier claims reach approval same-day as QA sign-off
60–70	Draft and compliance-approve decision-letter templates (approval/denial, by claim type)	Client service team lead + QA/compliance	10–15 hrs	Templates auto-populate correctly in test run
65–75	Configure auto-logging of communication method back into Guidewire	IT/ops	5–10 hrs	100% of a 2-week test sample logged correctly
75–85	Run first full end-to-end cycle-time measurement (4-week rolling)	Process Owner + QA Lead	~3 hrs/week	Cycle-time report issued against 21-day baseline and 10-day target
85–90	Conduct simulated regulatory audit-trail spot-check	QA/Compliance Lead	~1 day	Full claim timeline reconstructed from Guidewire alone, no Excel/paper reliance
Ongoing	Communicate remediation timeline and interim milestones to the two at-risk clients	Sponsor	~1–2 hrs/month	Clients have received a specific, dated remediation update at least once by day 30 and again by day 90

By day 90, all three design tiers are live, the first rolling cycle-time measurement is in hand, and the organization has one full month (April) before the end-of-May readiness deadline to intensify checklist refinement or tighten QA thresholds if the trend line isn't yet approaching single digits — the built-in decision point specified in the implementation playbook, preserved here rather than compressed by schedule pressure.

Appendix: Methodology & Assumptions

This methodology-and-assumptions appendix synthesizes the analytical approach, data limitations, and source basis underlying the preceding sections, and should be read as the transparency layer beneath the findings — it exists so that Coastal's leadership and any external reviewer can distinguish client-verified fact from consultant-derived estimate at every point in the report.

Methodology. The analysis followed a four-stage sequence: (1) as-is process mapping built directly from the client's own eight-step description, cross-referenced against the systems inventory (Guidewire, Outlook, the shared Excel tracker, and the paper QA checklist) to identify structural fragmentation points not explicitly named by the client but implied by the data; (2) bottleneck and cost analysis applying Lean value-stream principles (distinguishing value-added from non-value-added time) and Theory of Constraints (prioritizing the single largest constraint — document collection — before optimizing downstream steps), using the client's own cycle-time breakdown and staffing-cost figures as the quantitative base; (3) to-be design constructed

against three hard constraints stated by the client — no Guidewire replacement, no headcount increase, no removal of human judgment from assessment/approval — with every recommendation explicitly checked against the failure pattern of the 2022 document-chase pilot to ensure fixes are system-triggered rather than person-dependent; and (4) a phased implementation playbook and 90-day action plan sequencing Tier 1 (document collection) before Tier 2 (assignment/QA) before Tier 3 (approval/communication), on the logic that fixing downstream steps before the binding constraint would understate achievable impact. A regulatory-compliance overlay (Florida/Georgia claims-handling timeline and audit-trail requirements) was applied throughout rather than as a separate step, given the client's explicit constraint that any redesign must preserve documented audit evidence.

Assumptions made where client intake lacked a specific number. The following are explicitly labeled estimates, not client-provided data, and should be validated or replaced before being used in external communications (e.g., with the two at-risk clients or with regulators):

- **Contract revenue exposure (\$1.5M–\$2.3M annually, based on an assumed 12–18% revenue share per major contract).** The client did not disclose actual contract values for the two at-risk clients; this figure was sized illustratively from total revenue (\$6.5M) and should be replaced with real figures by leadership.
- **Industry incomplete-submission benchmark (25–35%) and the 9–12% target range for post-fix second-request rates,** drawn from general case-study literature, not Coastal-specific data. Coastal's actual improvement rate will only be known after the Phase 1 pilot.
- **Straight-through processing benchmark (70–85%)** used to contextualize Coastal's ~70% effective first-pass-clean rate — a comparative reference point, not a target Coastal has committed to.
- **Assignment automation benchmark (under 4 hours)** for digitally-routed intake in comparably sized claims teams — used to size the to-be assignment-stage target, not derived from Coastal's own systems testing.
- **QA turnaround benchmarks (2–4 days manual/paper vs. 0.5–1.5 days digitized)** — general claims-ops literature ranges, used to validate that Coastal's current 3-day QA stage sits in the "manual/paper" range rather than being an outlier.
- **Cost-per-claim industry benchmark (\$450–\$700),** against which Coastal's calculated \$632 fully loaded and \$402 processing-labor figures were compared — a directional benchmark only.
- **Rework labor-hour assumption (1.5–2 hours of active processor effort per document-rework cycle),** used to isolate direct labor cost from capacity-lock-up cost — an estimate, since intake did not provide time-motion data on how processors actually spend the 4–5 elapsed days per rework cycle.
- **Checklist-driven reduction case studies (30–50% cut in second-request rates),** cited from general process-improvement literature to support the plausibility of the Tier 1 target range, not verified against Coastal's specific claim types or policyholder population.
- **Loaded-cost-per-processor and per-QA-staff calculations** (roughly \$8,240–\$9,900/month) were derived by applying a standard 55–65% benefits/overhead loading factor to the stated salary bands — the client did not provide the exact loading factor used in their own \$215K/month figure, so this is a reasonable approximation, not a reconciled number.
- **The 1.5-day discrepancy between the client's stage-by-stage breakdown (19.5 days) and the stated 21-day average** was treated as unaccounted inter-stage handoff time — a reasonable inference given the system fragmentation identified in the as-is map, but not independently confirmed by the client.

Benchmark sources referenced. General industry context and comparative ranges were

drawn from: LexisNexis Risk Solutions claims studies and J.D. Power P&C claims satisfaction research (cycle time and communication as top satisfaction drivers); IAIABC claims-processing benchmarking studies; McKinsey and Deloitte insurance-operations research on claims cost and straight-through processing; general TPA industry association cost-per-claim data; and Florida Statute § 626.9541 and related unfair-claims-practices provisions, alongside Georgia's Unfair Claims Settlement Practices Act, for the regulatory-timeline context. These sources were used for directional calibration and plausibility-checking of Coastal-specific figures, not as precise benchmarks Coastal is expected to match exactly, since none are specific to Florida/Georgia regional P&C brokerages of Coastal's size.

Confidence Notes. Several elements of this report rest on thinner evidence than the rest and warrant additional client input before being treated as load-bearing for decision-making. First, **the contract revenue exposure figure** is the single weakest number in the report — it is illustrative sizing only, and leadership should substitute actual contract values immediately, both for internal prioritization and before any client-facing use. Second, **the risk-tiering QA thresholds** are currently unset — the report specifies a hard gate requiring compliance counsel sign-off before go-live, and the 20–25% "full review" tier cited throughout is a working assumption pending that sign-off, not a finalized policy; this is the most consequential open item given its regulatory exposure. Third, **Guidewire's actual configuration capability** (workflow rules for auto-assignment, custom status fields, digital form attachments) has not been confirmed against Coastal's specific license and instance — the report flags this explicitly as a week-one validation task, and if Guidewire cannot natively support the planned automations, portions of the Tier 1 and Tier 2 design would need rework. Fourth, **the reconciliation between the 19.5-day stage total and the stated 21-day average** is an inference, not a measured finding — a more granular time-and-motion study of inter-stage handoffs would sharpen this and could reveal additional recoverable time (or additional bottlenecks) not currently visible in the six-stage breakout breakdown Coastal provided.

Methodology. This analysis followed a four-stage sequence: (1) process mapping the as-is claim lifecycle directly from the client's own eight-step description and systems inventory, without adding steps not described; (2) bottleneck and cost analysis applying Lean value-stream principles and Theory of Constraints logic to isolate the single largest constraint (document collection) before addressing downstream stages, and translating cycle-time data into dollar and capacity terms using the client's stated staffing costs and volumes; (3) to-be design built explicitly within the three stated constraints — no core system replacement, no headcount increase, no removal of human judgment from assessment or approval — with every recommendation tested against the regulatory audit-trail requirement; and (4) an implementation playbook and 90-day plan sequenced by recovery-impact-to-complexity ratio, with hard go/no-go gates at points of highest risk (QA risk-tiering compliance sign-off, Phase 1 pilot results) rather than a single big-bang rollout. Throughout, industry benchmarks were used only for directional comparison and prioritization logic, never as a substitute for the client's own reported figures, which anchor every quantitative claim about Coastal's current state.

Assumptions made where intake lacked a specific number, each labeled at the point it was introduced in the report body, are consolidated here for ease of reference:

- **Revenue share of the two at-risk client contracts (12–18% of total revenue each)** — intake did not disclose actual contract values; this was used solely to illustrate the order-of-magnitude scale of client-retention risk relative to implementation cost, and should be replaced with real figures before use in any external or board-level communication.
- **Industry benchmark ranges** for comparable TPA/brokerage cycle times (7–15 days typical, 5–7 days best-in-class), incomplete-submission rates (25–35% typical), straight-through

processing rates (70–85% in leading operations), assignment queue turnaround in automated systems (under 4 hours), and cost-per-claim (\$450–\$700) — all drawn from general claims-operations literature (cited below), not Coastal-specific data, and used only for comparative context, not as targets Coastal is contractually or operationally bound to.

- **The 30–50% reduction in second-request rates** attributed to pre-submission checklists in comparable case studies — a directional benchmark used to justify the Tier 1 target range (18% down to 9–12%), not a guaranteed outcome for Coastal specifically.
- **Active processor labor time per rework cycle (1.5–2 hours)** — estimated to isolate direct wage cost from capacity-lock-up cost, since intake did not specify how much of the 4–5 elapsed days per rework cycle is active work versus passive waiting on the policyholder.
- **The processing-labor cost allocation table** (\$402/claim across six stages) is a modeled estimate built by applying loaded staff costs against stage durations; it necessarily assumes reasonably even distribution of cost within each functional group and does not account for possible variation by processor seniority, claim complexity, or line of business, none of which were broken out in intake.
- **QA risk-tiering threshold levels** (e.g., "top 20–25% of claims by value/complexity" for full manual review) are illustrative starting points for the design, explicitly flagged in the to-be design as requiring compliance counsel sign-off and real claim-mix data before finalization — they are a design hypothesis, not a validated policy.
- **Guidewire configuration capability** — the entire Tier 1 and Tier 2 automation design assumes Guidewire's existing workflow/business-rules engine and portal can support auto-assignment by caseload, custom status fields, and digital form attachments with timestamp logging, without a paid add-on or vendor engagement. This has not been confirmed against Coastal's actual license tier and is flagged as a week-one validation item, not a settled fact.
- **The reconciliation between the 19.5-day stage total and the stated 21-day average** is an inference, not a measured finding — a more granular time-and-motion study of inter-stage handoffs would sharpen this and could reveal additional recoverable time (or additional bottlenecks) not currently visible in the six-stage breakdown Coastal provided.

Benchmark sources referenced. General directional context was drawn from publicly available industry research and regulatory text, including Florida Statute § 626.9541 and related unfair claims practices provisions, Georgia's Unfair Claims Settlement Practices Act, LexisNexis Risk Solutions claims industry studies, J.D. Power P&C claims satisfaction research, IAABC claims-processing studies, and general insurance-operations research published by major consulting firms (McKinsey, Deloitte) on TPA and claims-administration efficiency. These sources were used exclusively to sanity-check the scale of Coastal's gap against typical market performance and to validate that the proposed interventions (automated reminders, risk-tiered QA, digitized checklists) are established, low-risk practices elsewhere in the industry — not to set Coastal's specific targets, which remain grounded in the client's own reported baseline and stated objective.

Confidence Notes

Three areas of this analysis are meaningfully thinner than the rest and would benefit from additional client input before being treated as load-bearing for budget or regulatory decisions. First, **the financial risk sizing for the two at-risk contracts** rests entirely on an illustrative percentage-of-revenue assumption; Coastal's leadership should replace this immediately with actual contract values and current SLA terms, since this number is likely to be used in board-level prioritization discussions and should not rest on a placeholder. Second, **the QA risk-tiering thresholds** are a design hypothesis pending compliance counsel review — this is the

single highest-severity open item in the entire engagement, since it is the one place where speed and regulatory exposure intersect directly, and it should not be finalized or scaled without written sign-off, regardless of schedule pressure heading into hurricane season. Third, **the Guidewire configuration assumptions** underpinning nearly all of Tier 1 and Tier 2 automation have not been validated against Coastal's actual system instance and license tier; if core capabilities assumed here (status-triggered reminders, custom form attachments, caseload-based auto-assignment) turn out to require vendor professional services or aren't available at all, both the effort estimates and the recovery timeline in the implementation playbook would need meaningful revision. All three items are flagged as first-week validation priorities in the 90-day plan specifically because they are the assumptions most likely to change the shape of the recommendation if the client's actual data or systems differ from what intake described.